Chapter 8

Visual object naming in optic aphasia

The general aim of this thesis is to investigate the effects of damage in connectionist networks that build attractors, in order to identify and extend the computational principles that enable them to reproduce the detailed pattern of cognitive impairments in some types of patients with neurological damage. All of the simulations presented thus far, as well as most other work in the area, involve impairments in reading—acquired dyslexia. This is in part because reading has been one of the most intensely studied domains in both cognitive psychology and neuropsychology in recent years (Coltheart, 1987), resulting in a rich, and often counterintuitive, set of empirical findings. In addition, reading is appealing as a domain for computational modeling because the surface forms (i.e. strings of letters and phonemes) are fairly simple. However, reading shares many characteristics with other cognitive tasks, most notably visual object recognition. If connectionist modeling is to be more generally applicable in neuropsychology, then the computational principles that have proven useful in understanding deficits in reading should have interesting consequences in a related domain such as object recognition and naming. The purpose of this chapter is to demonstrate that the principles that explain the reading behavior of deep dyslexics can be extended to account for the characteristics of patients, called "optic aphasics," who have a selective deficit in visual object naming.[1]

Optic aphasics are particularly interesting because their inability to name visually presented objects does not appear to be caused either by impaired visual recognition or by impaired naming. Visual recognition demonstrated by means other than naming (e.g. gesturing), and naming via modalities other than vision (e.g. by tactile input), are relatively intact. Only naming visual input is selectively impaired. This is difficult to interpret in terms of a single general semantic system that subserves all types of recognition and naming (see Figure 8.1).

The rest of the chapter begins with a more detailed description of the characteristics of optic aphasics, distinguishing them from closely related patients, called "visual agnosics," in which recognition actually is impaired. Somewhat reminiscent of deep dyslexic reading errors, both types

[1]This research was done in collaboration with Tim Shallice and is also described in Plaut & Shallice (in press).

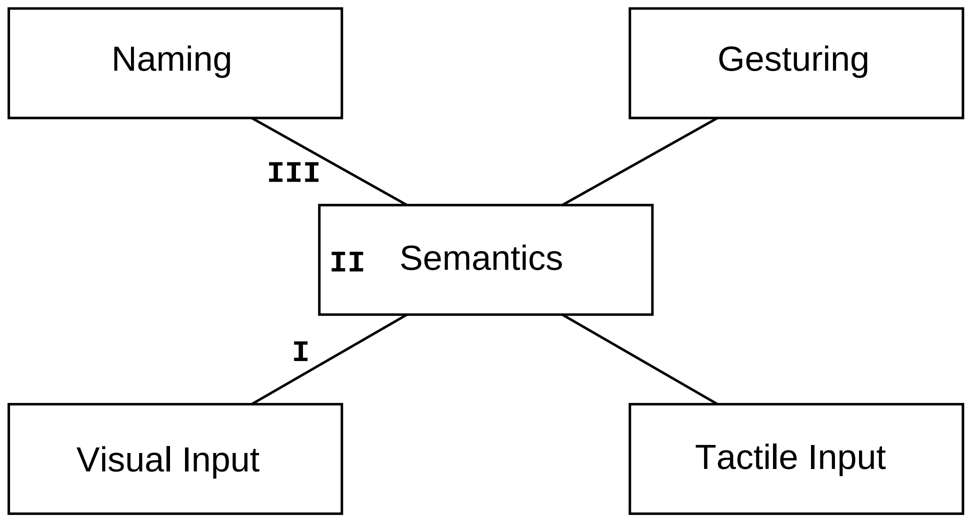

Figure 8.1: A schematic model of naming and gesturing to visual and tactile input. The selective deficit in visual naming shown by optic aphasics is difficult to account for in terms of damage to such a model. Damage between vision and semantics (**I**) appears to be ruled out because gesturing to visual input can be intact; similarly for damage within semantics (**II**); damage between semantics and naming (**III**) appears to be ruled out because tactile naming can be intact.

of patients make errors in object naming that are visually and/or semantically similar to the presented object, although they differ in their typical distribution of error types. Another common feature is the frequent occurrence, particularly in optic aphasia, of *perseverations*—intrusions of responses given to previously presented objects when naming the current object. In our previous connectionist simulations it is impossible for the presentation of one word to affect a subsequent word. Thus, accounting for the error pattern in optic aphasia requires extending the computational formalism to include a mechanism for influences across successive input presentations. We introduce short-term correlational weights between units, for which there is independent computational and empirical motivation. A network with these weights, as well as the conventional slowly modifiable weights, is trained to generate semantic representations of objects when presented with a representation of their visual form. The pattern of errors that the network makes under a variety of types of damage is then compared with that of optic aphasics. Although we do not attempt to simulate the intact recognition in these patients when naming is impaired, we conclude the chapter with a discussion of how this might be incorporated within the current theoretical framework.

8.1 Optic aphasia

It will be easier to understand the nature of impaired visual object *naming* in optic aphasia if we first consider patients with selective deficits in visual object *recognition*.

8.1.1 Visual agnosia

Visual agnosia, first termed “mind blindness” by Monk (1881), is a relatively rare neuropsychological disorder in which patients have an impairment in visual object recognition that cannot be explained by lower level sensory deficits or more generalized mental deterioration (for reviews, see Bauer & Rubens, 1979; Farah, 1990; Humphreys & Riddoch, 1987; Levine, 1982). In contrast to their poor recognition with visual presentation, recognition of objects presented in other modalities, such as with tactile or auditory input, is relatively intact. Lissauer (1890, translated in Shallice & Jackson, 1988) was the first to distinguish two types of visual agnosia. In “apperceptive” agnosia, difficulties in recognition are attributed to a relatively high-level perceptual impairment. In other words, even though basic visual abilities are intact, the patient does not see objects normally and so cannot recognize them. In general, apperceptive agnosics cannot match, copy or even trace shapes, and are overly distracted by irrelevant lines in drawings (Benson & Greenberg, 1969).

In contrast, “associative” agnosics appear to see objects quite normally but still cannot recognize them. The experience of viewing an object is said to be of a “normal percept, stripped of its meaning” (Teuber, 1968). Thus Taylor & Warrington (1971) describe a patient who, on presentation of a second photograph (from a different angle) of an object he previously failed to recognize, stated “I already told you that I don’t know what that is.” This might be conceptualized as a disconnection between vision and semantics, at **I** in Figure 8.1 (Geshwind, 1965). However, more detailed testing has raised the possibility that a more subtle visual impairment may underly the recognition deficits of associative agnosics (Farah, 1990). Although these patients can match, copy and trace objects, they perform these tasks in a highly abnormal, slavish line-by-line fashion. Some cannot distinguish possible from impossible objects (Ratcliff & Newcombe, 1982), match usual and unusual views of objects (Humphreys & Riddoch, 1984; Warrington & Taylor, 1978), or integrate features and parts into a coherent whole (Humphreys & Riddoch, 1987; Levine & Calvanio, 1989). Under this interpretation, patients heretofore described as associative agnosics would actually constitute high-level apperceptive agnosics under a strict interpretation of Lissauer’s distinction.

8.1.2 Optic aphasia

First described by Freund (1889), optic aphasia would appear at first glance to be a “true” form of associative agnosia. Patients cannot name visual stimuli, even though their naming of auditory and tactile stimuli is relatively intact. Visual representations in the form of structural descriptions (Marr & Nishihara, 1978; Palmer, 1977) appear intact. Thus Riddoch & Humphreys (1987) found that their patient J.B. could perform an “object decision” task analogous to lexical decision: distinguish drawings of real objects from those of nonexistent objects composed of the parts of real objects. It is difficult to see how such a task could be accomplished without visual representations that are sufficient for recognition. Furthermore, optic aphasics are much less affected than associative agnosics by manipulations of the visual quality of stimuli (Larrabee et al., 1985). Noting that associative agnosics have been known to resolve into optic aphasics during the course of recovery, some researchers (e.g. Bauer & Rubens, 1979) suggest that optic aphasia is simply a mild version of associative agnosia.

However, a critical difference emerges when the recognition capabilities of optic aphasics are tested in tasks other than naming. Unlike associative agnosics, optic aphasics can often correctly gesture the use of an object they cannot name (Gil et al., 1985; Lhermitte & Beauvois, 1973; Riddoch & Humphreys, 1987), as well as sort objects by semantic category (Assal & Regli, 1980; Coslett & Saffran, 1989b; Riddoch & Humphreys, 1987). In addition, they have little trouble interacting with the world in daily life, whereas agnosics are quite impaired (Humphreys & Riddoch, 1987). Clearly, the recognition capabilities of optic aphasics are much better than their naming performance would suggest. As pointed out previously, this dissociation between preserved visual recognition and impaired naming, with intact naming via other modalities, is difficult to explain in terms of conventional formulations of the recognition and naming process.

Optic aphasics and associative agnosics also differ in the typical pattern of errors they make when naming objects. Associative agnosics make predominantly visual (also called "morphological") errors, misnaming objects as ones with similar shape (e.g. *rubber band* ⇒ "bracelet", Levine, 1978), although some patients also make semantic errors (*violin* ⇒ "trumpet", Davidoff & Wilson, 1985). In contrast, the naming errors of optic aphasics are most typically semantically related to the stimulus, or both visually and semantically related (e.g. *apple* ⇒ "orange") but only rarely are purely visually related (Coslett & Saffran, 1989b; Gil et al., 1985; Lhermitte & Beauvois, 1973; Riddoch & Humphreys, 1987).

In addition, a striking characteristic of optic aphasics, and to a lesser extent of associative agnosics, is the frequent occurrence of perseverative effects in errors. Thus the responses to previously presented objects often interfere with naming the current object. Some authors dismiss these as uninteresting, classifying perseverative errors along with "other" errors, perhaps in part because perseveration is not specific to optic aphasics but is common in many types of aphasic patients with more generalized language deficits (Albert & Sandson, 1986). However, the perseverative effects in optic aphasia are particularly well documented (Lhermitte & Beauvois, 1973) and show interesting interactions with the nature of the other types of errors. For example, not only are the exact responses (correct or incorrect) to previously presented objects often given subsequently, but responses that are *semantically related* to previous objects are also given. Thus J.F. (Lhermitte & Beauvois, 1973) correctly named a picture of a *baby*, but then misnamed a *knife* as a "child." In Lhermitte & Beauvois' terminology, this constitutes a "vertical" semantic error (in time, moving down a list of objects) as distinction from the the standard "horizontal" semantic errors (from stimulus to response listed across the page). Vertical mixed visual-and-semantic errors also occur (e.g. *knife* ⇒ "pen" following correct naming of a *pencil*). Thus semantic and perseverative influences on naming errors appear intimately related.

It should be noted that the error patterns in visual object naming of some patients that are classified as associative agnosics share many of characteristics of optic aphasia. For instance, Lissauer' (1890) original patient G.L. shows a more even balance between visual errors and semantic errors than is typical of either associative agnosics or optic aphasics, as well as a strong tendency to perseverate. As an illustration, here is a portion of an object naming protocol from this patient (adapted from Shallice & Jackson, 1988, p. 174):

Object	Response
light	"Pencil."...given the light to feel and immediately recognizes it as a light.
pen	"That's a light."...asked to touch the object but is frightened to burn himself. He then recognizes it at once as "writing pen."
spectacles	"Lamp." After feeling it: "Spectacles."
handkerchief	"Spectacles." After touch: "Cloth."
carafe of water	"Lamp into which you put a light."...After touching it the patient persists: "It is a lamp."
metal box	"I don't really know." After touching: "Matches."
candle	"A piece of light."
bread roll	"Bread roll."
pen	"Candle snuffers." After touching: "Writing pen."
piece of paper	"Handkerchief." After touching: "Envelope."
door knob	"Snuffers," after some thought: "A candlestick."

Notice the frequent perseverations, including a semantic perseveration (*pen* $\Rightarrow$ ""candle snuffers"" following the presentation of a *candle*. One explanation is that such a patient suffers from a combination of lesions leading to both associative agnosia and optic aphasia to some degree. A more interesting possibility is that each syndrome reflects the effects of single lesions at different points within the mechanism that recognizes and names objects. This would be analogous to explaining the differences among types of deep dyslexics (i.e. input, central, and output, Friedman & Perlman, 1982; Shallice & Warrington, 1980) in terms of damage at different locations along the semantic route for reading, while still accounting for their qualitative similarities in terms of the same computational principles.

A final characteristic of the behavior of optic aphasics in naming objects that we mention for completeness is that, if given unlimited time, they can often "home in" on the correct name. Thus J.F. produced the following responses (from Lhermitte & Beauvois, 1973, p. 706):

Object	Response
basket	(a preceding picture had been named "log") "A kind of andiron made of cane, of osier, a basket."
cup	(*cork screw* had just been correctly named) "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."
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."
window-blind	"a parasol, metallic curtain rods...the cloth roof...surrounding sails...it could be a parasol...there are rods, but isn't it a shelter? A window-blind...not a window-blind, the window-blind is rolled up, there is no device to roll, a sunshade...it should be a window-blind which does not roll."

This last response is particularly interesting because the correct name is explicitly given but rejected at first by the patient. Thus the naming impairment cannot be due to difficulty in generating the names themselves. However, we will have little to say about this characteristic of optic aphasia, since our simulations can generate only single responses to a given object.

8.1.3 Theoretical accounts

A number of different theoretical accounts have been given for the impaired naming with intact recognition of visual stimuli in optic aphasia, although none of them are completely satisfactory. One proposal (Ratcliff & Newcombe, 1982) is that there is a "direct" route from vision to naming, analogous to the (lexical) phonological route that has been proposed in reading (e.g. Sartori et al., 1987; Schwartz et al., 1980). If this route is impaired, naming via the semantic route becomes unstable and yields semantic errors. Unfortunately, there is no independent evidence for such a route as there is in reading (e.g. patients who can name objects but not know what they are) and so this proposal remains completely *ad hoc*.

Another explanation (e.g. Beauvois, 1982) is that semantics is not a unitary entity, but is separated into "visual" and "verbal" components. Visual input can only directly access visual semantics, and naming can only be based on verbal semantics. Visual object naming requires communication from visual to verbal semantics; optic aphasia arises naturally from a disconnection between them. Intact gesturing and categorization can be based on visual semantics, while intact auditory recognition is based on direct access to verbal semantics. Tactile naming would presumably occur via a third semantic component, and so on. The main problem with this account is that a reasonable definition of "visual" semantics is too narrow to account for the range of information that patients appear to have available in recognition tasks. Both visual and verbal semantics would require virtually complete replication of all semantic knowledge about objects, and thus the distinction becomes unexplanatory.

A related proposal (Coslett & Saffran, 1989b) is that semantics is divided not by modality but by *hemisphere*, with naming only supported in the left hemisphere. On this hypothesis, optic aphasia arises when visual input from both hemispheres is disconnected from left-hemisphere semantics, with residual comprehension subserved by right-hemisphere semantics. In essence, this theory parallels the right-hemisphere hypothesis for reading in deep dyslexia (Coltheart, 1980b; 1983; Saffran et al., 1980) and shares many of its strengths and weaknesses (see Coltheart et al., 1987a; Patterson & Besner, 1984a; Shallice, 1988; and the General Discussion). As with the visual-verbal disconnection hypothesis, the main issue is the adequacy of independent constraints on the nature of the separate semantic systems. Hypotheses about *localization* are only germane to theories of cognitive *function* if there is some reliable means of associating the two. With regard to the functional characteristics of right hemisphere semantics, the available evidence is tantalizing but inconclusive.

Yet another account of optic aphasia (Riddoch & Humphreys, 1987) locates the impairment between vision and semantics (location **I** in Figure 8.1) and challenges the claim that recognition is intact in these patients. In this way optic aphasia would amount to a type of "semantic access agnosia." Because our simulations involve a similar type of damage, and because intact visual recognition plays a critical role in the *definition* of optic aphasia as distinct from associative agnosia, this challenge deserves close consideration.

The claims of intact recognition in optic aphasia have been based almost entirely on their performance in tasks involving either gesturing or semantic categorization. Considering gesturing tasks first, in fact adequate gesturing to misnamed objects has been demonstrated in only three cases (Gil et al., 1985; Lhermitte & Beauvois, 1973; Riddoch & Humphreys, 1987)—in two others (Coslett & Saffran, 1989b; Larrabee et al., 1985) gestures corresponded to the (incorrect) named object, and in a third (Assal & Regli, 1980), gesturing and naming visual stimuli were equally impaired. However, gesturing in J.F. (Lhermitte & Beauvois, 1973) appeared to be quite well-preserved—he never made an incorrect gesture to a set of 100 pictures of objects of which he made 31 naming errors, although the authors fail to state whether he failed to gesture at all for some pictures. To explain this performance, Riddoch & Humphreys point out that gesturing is often judged less stringently than naming, and typically requires less precise semantics. Thus W.L.P. (Schwartz et al., 1979), with severely impaired semantics due to progressive dementia, could nonetheless normally manipulate and demonstrate the use of objects she could not name. In fact, her gestures were sufficiently precise for observers to distinguish, for example, a spoon from a fork, or a pipe from a cigarette. Riddoch & Humphreys' patient J.B. was 75% correct at gesturing but only 45.5% correct at naming in a task in which objects were selected to have fully discriminable gestures. Thus gesturing shows some impairment, but is still better than naming. To account for the remaining advantage in gesturing, Riddoch & Humphreys suggest that appropriate gesturing can often be accomplished solely or in part on the basis of non-semantic information, such as shape. However, it is unclear whether shape alone is sufficient to account for the correct gestures.

Riddoch & Humphreys argue that categorization tasks also require less precise semantics than naming. In fact, tasks such as sorting objects into semantic categories are performed well but never perfectly by optic aphasics (Assal & Regli, 1980; Coslett & Saffran, 1989b; Riddoch & Humphreys, 1987). Another problem is that these tasks may have been too easy to reveal slight or moderate impairments. In fact, when Riddoch & Humphreys tested J.B. on a more difficult categorization task, involving discriminations *within* categories, they found he was significantly impaired with visual as compared with auditory presentation. Although they failed to compare J.B.'s performance with that of normal control subjects, there was no evidence that J.B. had any auditory comprehension deficit as he was perfect at naming to auditory definitions.

There are a few other sources of evidence that visual recognition is not normal in optic aphasics. J.F. would occasionally fail to recognize complex objects he could draw and describe quite well, and occasionally pointed to a semantically related object when the object named by the experimenter was not among those present (Lhermitte & Beauvois, 1973). Also, Gil et al.'s (1985) patient performed poorly at sorting objects on the basis of metaphorical associations under visual as compared with auditory presentation. Taken together, there would seem to be reasonable evidence that the semantic representations derived from visual objects by optic aphasics are impaired relative to those of normals. In a sense this observation mirrors the discovery of more subtle visual impairments in associative agnosia that may underly those patients' deficits. It is an open question whether the "semantic access" impairment in optic aphasia can be both severe enough to account for poor naming performance, and yet mild enough to allow for the observed levels of performance on gesturing and categorization tasks. We will return to this possibility at the end of the chapter.

A final, more recent proposal (Farah, 1990) hypothesizes that optic aphasics have two partial lesions, one between vision and semantics, and the other between semantics and naming (locations **I** and **III** in Figure 8.1). Each separate impairment is sufficiently mild to allow reasonable performance on gesturing (for **I**) or tactile naming (for **III**) but tasks that require both pathways—visual naming—would be much more drastically impaired. Because damage is proposed between vision and semantics on this hypothesis, relatively mild recognition impairments as described above might be predicted. The explanation for the disproportionate naming impairment hinges on the notion that a connectionist system with attractors between each of these levels would be sufficiently robust under each partial lesion alone, but would show *superadditive* impairment under the combined lesions. Unfortunately, preliminary simulations exploring this possibility, carried out by the author in collaboration with M. Farah, have been unsuccessful to date.

None of these proposals is completely satisfactory—each involves either *ad hoc* assumptions or insufficiently supported claims. They all focus on explaining the dissociation between impaired naming and intact recognition, but have little to say about other characteristics of optic aphasics, such as the nature of the naming errors they make (an exception to this is Riddoch & Humphreys, 1987). Perhaps a more detailed consideration of these aspects of the syndrome may shed light on other aspects of the syndrome. In the simulations described below, we focus on reproducing the pattern of semantic and perseverative influences on naming errors in optic aphasia, without directly addressing how such a system might also support gesturing and categorization (to whatever extent they are relatively preserved in these patients). Thus they should not be interpreted as a complete simulation of optic aphasia, but merely as demonstrations of principles that might lead to a more complete account. How the work might be extended into a full account is discussed following the presentation of the simulations.

8.2 Short-term correlational weights

The co-occurrence of visual, semantic, and mixed visual-and-semantic errors in object naming would appear to be analogous to the related error types in deep dyslexic reading, suggesting a natural account in terms of a network that maps visual representations onto semantic representations using attractors. However, the perseverative effects in optic aphasia, and their interactions with semantic effects, are less straightforward. In the dyslexia simulations, the network is completely reset before the presentation of each word—there is no opportunity for the response to one stimulus to influence responses to subsequent stimuli. Accounting for the perseverative effects in optic aphasia requires an elaboration of the computational formalism we have used thus far.

There are many possibly ways of introducing effects of the temporal order of stimulus presentation into connectionist networks. The approach we adopt involves introducing short-term weights that depend on the recent correlations between unit states. These weights augment the standard weights that are slowly modified over the course of learning. In particular, each connection is given a short-term correlational weight whose value is an exponentially decaying weighted average over stimulus presentations of the correlation of the states of the units it connects. More formally, if s_i and s_j are the states of units i and j after processing stimulus $n-1$, then the correlational weight c_{ij} on the connection from i to j is set according to

$$c_{ij}^{[n]} = \lambda s_i' s_j' + (1-\lambda) c_{ij}^{[n-1]} \tag{8.1}$$

where $s_i' = 2s_i - 1$ (scaling each unit state to range between ± 1) and λ is the exponential weighting proportion (0.5 in our simulations). In processing the next stimulus n, the summed input $x_j^{(t)}$ to each unit j at iteration t becomes

$$x_j^{(t)} = \sum_i s_i^{(t-1)} \left(w_{ij} + \gamma c_{ij}^{[n]} \right) \tag{8.2}$$

where γ balances the contribution of the short-term weights relative to the long-term (learning) weights (0.05 in our simulations; cf. Equation 10.1, p. 303). The states of units are computed from their summed input according to the standard sigmoid function (see Equation 10.2, p. 303). Notice that the short-term correlational weights do not change over iterations in processing a stimulus, but change only once the network has settled. The effect of the short-term weights is to bias the network towards recently occurring patterns of activity. Although our simulations involve back-propagation networks, it may help to think of the short-term weights as temporarily lowering the energy (improving the "goodness") of the minima corresponding to the previous stimulus—in a DBM this would be precisely true.[2]

There is independent computational and empirical motivation for introducing short-term weights. In the domain of object recognition, the most common use of short-term interactions among units is to temporarily bind together combinations of visual features into a coherent whole (Crick, 1984; von der Malsburg, 1981; 1988; von der Malsburg & Schneider, 1986). The recent discovery of synchronized oscillations in the responses of visual cortical cells to disjoint moving contours of a single object (Eckhorn et al., 1988; Gray et al., 1989) has lead to the development a number of models of synchronized neuronal activity for feature binding involving short-term interactions among units (Atiya & Baldi, 1989; Baldi & Meir, 1990; Eckhorn et al., 1989; Horn et al., 1991; Hummel & Biederman, 1990; Kammen et al., 1990; Konig & Shillen, 1991; Sompolinsky et al., 1989; Sporns et al., 1989; Wilson & Bower, 1990).

Short-term weights have other interesting computational properties. As described in Section 7.2, learning with fast weights can minimize the interference to old knowledge caused by new learning, and to rapidly recover the old knowledge by canceling out this interference (Hinton & Plaut, 1987). Although the fast weights employed by Hinton & Plaut are changed by error on the task rather than by the states of units directly, they would induce similar biases towards previous interpretations if applied in an "on-line" learning paradigm in which both slow and fast weights were updated after every stimulus presentation (McClelland & Rumelhart, 1985) . Hinton (personal communication, described in McClelland & Kawamoto, 1986) demonstrated how to use short-term weights to implement recursion in a network that draws shapes composed of other shapes. The long-term weights hold the knowledge about how to draw shapes. The short-term weights hold context information about what to draw next once the network is finished with drawing the current shape. Thus the short-term weights function like a "stack" that can reinstate the calling context once a drawing "subroutine" returns. Short-term interactions have also been employed for recruitment of units during learning (Feldman, 1982).

Short-term weights are also useful in accounting for empirical phenomena in cognitive psychology. The most obvious of these are repetition and semantic priming effects, both in normals (Collins & Quillian, 1969; Mandler, 1980; Meyer & Schvaneveldt, 1976) and amnesics (for a re-

[2] In many ways the current simulations would have been more natural within a DBM framework. For instance, the interactions between short-term and long-term weights are more easily understood in terms of their effects on an energy surface. In addition, unit states would not have to be normalized to between ± 1 when setting the short-term weights. Unfortunately we must leave such simulations to future research.

view, see Shimamura, 1986). McClelland & Rumelhart (1985) simulate a range of priming effects with immediate changes directly to the slow weights rather than to a separate set of short-term weights, although the same results would also hold in the latter case. Another appropriate domain involves short-term memory and its consolidation into long-term memory (e.g. Gardner-Medwin, 1989). Goebel (1990) suggests how to use fast weights for serial rehearsal in short-term memory. Cleeremans & McClelland (1991) show how fast weights can account for the temporary biases of subjects in learning to respond to structured event sequences. This last work is particularly interesting because it involves specific biases towards recently occurring *associations* between stimuli, above and beyond the bias changes for the individual stimuli themselves. This suggests that the short-term mechanism involves *weights* between units rather than, or in addition to, simple threshold changes for individual units (cf. Morton, 1969).

Thus there is some independent motivation for extending the computational formalism to include short-term correlational weights as a means of introducing temporal interactions between successive stimuli. However, it should be kept in mind that we are extending the formalism in direct response to the observation of perseverations in optic aphasia, and it is in this sense rather *ad hoc*. For this reason, the simple occurrence of perseverations in the network should be viewed as less interesting than the interactions of these perseverative effects with other aspects of the network's behavior, which are not inherent in the extension of the formalism.

In many ways, it would have been more natural to introduce temporal effects by processing each object beginning from the set of states corresponding to the interpretation of the previous object, rather than resetting the network. However, we chose to introduce short-term weights because perseverative effects in optic aphasics can span intervening objects (Lhermitte & Beauvois, 1973), which would be difficult to account for solely in terms of sustained activity across object presentations (see also Joordens & Besner, 1992).

8.3 A simulation of visual object naming in optic aphasia

We develop a network for mapping visual representations of objects onto semantic representations, and compare its behavior under damage with that of optic aphasics. We begin by describing the details of the task the network is to perform. We then described the network architecture and the procedure by which it is trained. Following this, we describe how the behavior of the network under damage is compared with the behavior of optic aphasics in visual object naming tasks.

8.3.1 The task

Forty objects were chosen from four categories of indoor objects: kitchen objects, office objects, furniture, and tools (see Table 8.1). The objects were chosen to have names with at most five letters and two syllables, although these constraints are not relevant to the simulations presented in

Objects in each category			
Kitchen Objects	Office Objects	Furniture	Tools
cup	*pen*	*chair*	*saw*
spoon	*file*	*table*	*nail*
pan	*paper*	*bed*	*plane*
fork	*book*	*sofa*	*ruler*
knife	*disk*	*stool*	*screw*
bowl	*tape*	*rug*	*awl*
can	*stamp*	*radio*	*axe*
plate	*board*	*tele*	*bolt*
dish	*glue*	*divan*	*nut*
glass	*ink*	*desk*	*vice*

Table 8.1: The objects used in the simulations.

this chapter. We first describe their visual (input) representations, and then their semantic (output) representations.

The input representation for objects was designed to coarsely approximate the kind of visual information that would be available for the purposes of object recognition. The representation of each object loosely corresponds to a structural description (Marr & Nishihara, 1978; Palmer, 1977), augmented with information about color, texture, size, and more general visual characteristics of the object. Table 8.2 lists the type of information represented by each of the 44 visual features. The possible values for each of these types of information are encoded as different patterns of activity over the designated feature groups. The first 25 features are devoted to representing the shape of the object in terms of up to three "components," one of which is designated as the main component.[3] The shape of each component is encoded over five units, as shown in Table 8.3. The position and size of the second and third components relative to the main component are described in terms of two and three additional units, respectively (see Table 8.4). The remaining 19 of the 44 features describe more general visual characteristics of the object, as well as color, texture, and absolute size information (see Table 8.5).

Table 8.6 describes the visual representations of each of the 40 objects in terms of the codes listed in the tables for values of each type of information. Figure 8.2 shows the actual assignment of each of the 44 visual features to each object. The objects are listed together by category and, although it is somewhat difficult to see directly in the figure, there seems to be a tendency for objects within a semantic category to be somewhat *visually* similar. This can be seen more clearly in the similarity matrix for the visual representations of objects, shown in Figure 8.3. In this figure,

[3]These might be thought of as loosely corresponding to Biederman's (1987) "geons."

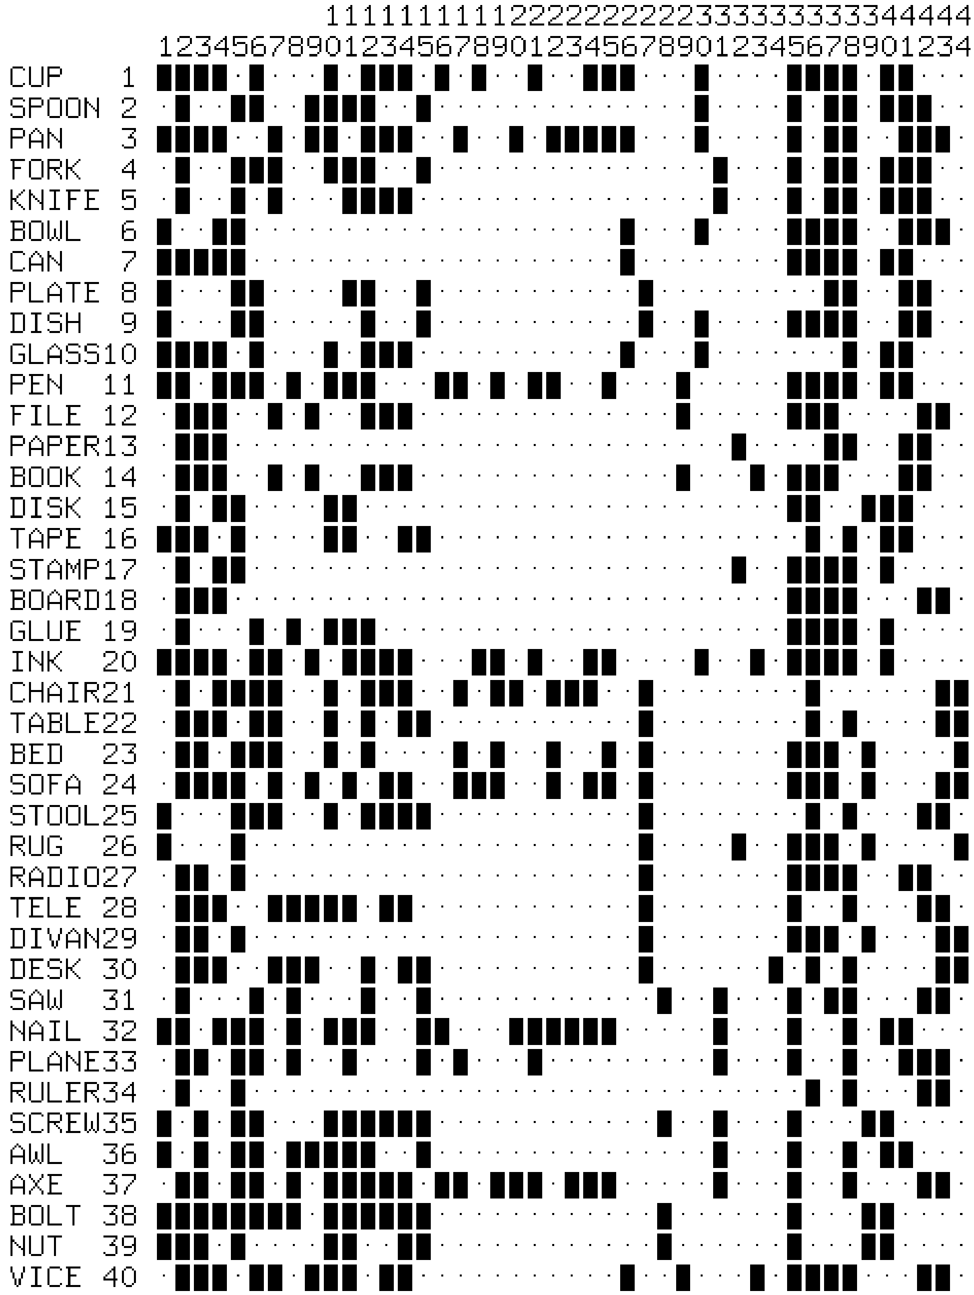

Figure 8.2: Assignment of visual features to objects.

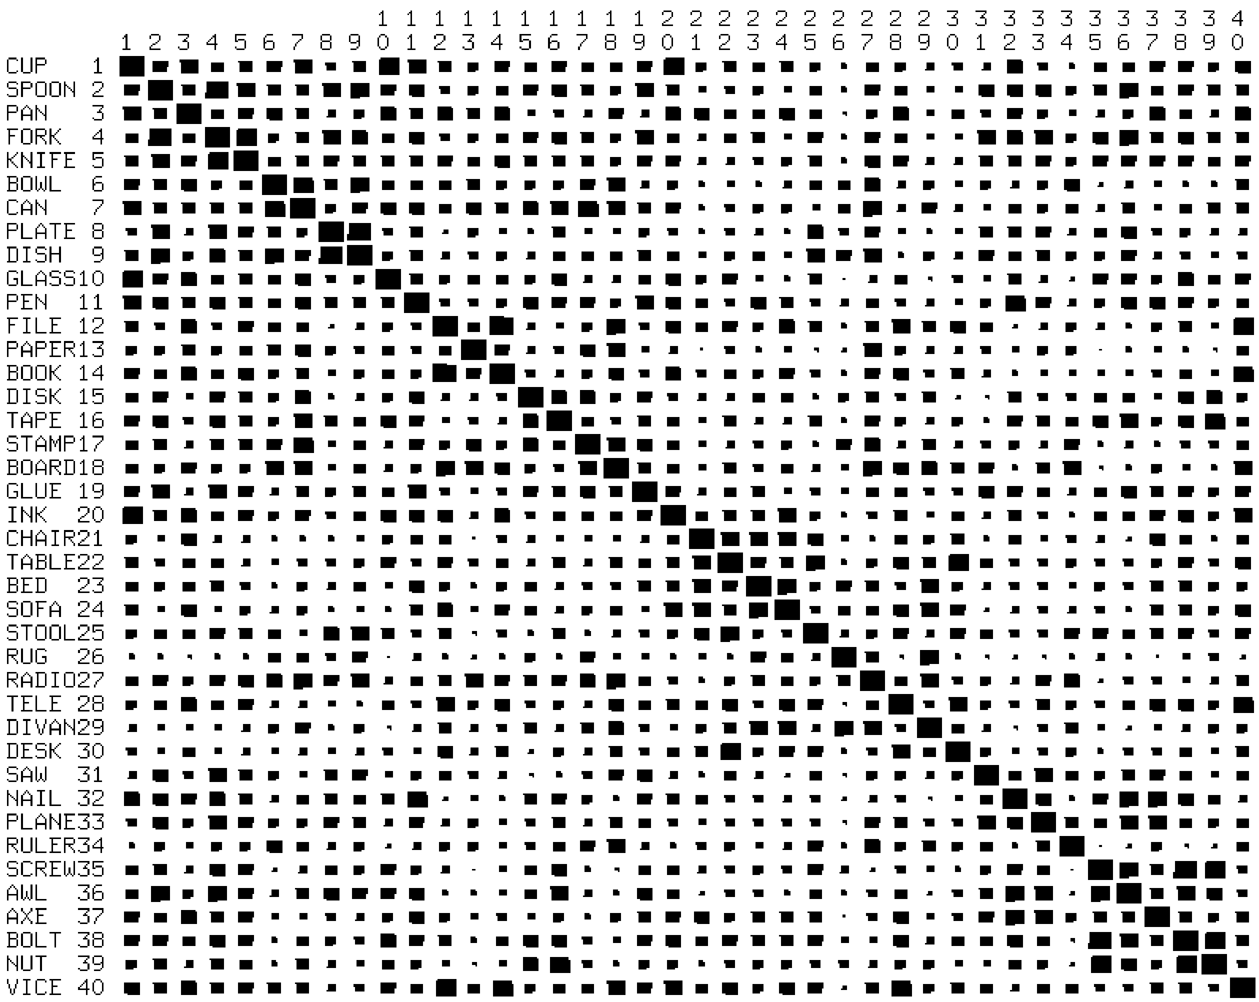

Figure 8.3: The similarity matrix for the visual representations of objects. The size of each blob represents the proximity between the representations of two particular objects.

Types of Visual Features		
Features	Description	
1–5	Main Component	Shape
6–10	Second Component	Shape
11–12		Relative Position
13–15		Relative Size
16–20	Third Component	Shape
21–22		Relative Position
23–25		Relative Size
26–34	General Characteristics	
35–37	Color	
38–39	Texture	
40–44	Absolute Size	

Table 8.2: The type of information represented by each visual feature.

the visual similarity within each semantic category corresponds to one of four 10-by-10 blocks along the diagonal. Notice that visual similarity is relatively high within some semantic categories, such as among some kitchen objects, items of furniture, and tools. We will return to the issue of the relationship between visual and semantic similarity after we present the details of the semantic representations.

Much like in previous simulations, the semantics of each of the 40 objects is represented in terms of a set of semantic features, listed in Table 8.7. Twenty-eight of the 86 features represent the visual semantics of the object. The first 14 of these are identical to the general visual characteristics, color, and texture encoding used in the visual (input) representations. The next three are a condensed version of the absolute size encoding, and the remaining 11 summarize the shape of the object. Following this, there are features for consistency (2), material of which the the object is made (8), where the object is found (10), its general (10) and specific (22) function, and the general actions associated with it (7). We assume that information about more specific actions associated with an object would be given a non-semantic, possibly motoric, representation, in the same way that the semantic representation of an object does not contain detailed visual information.

Figure 8.4 shows the assignment of semantic features to each of the 40 objects. Notice that some features are quite indicative of an object's category. For example, unlike most objects, all items of furniture have a horizontal main component (feature 2), which makes sense given their common function. Features relating to general function (49-57) tend to distinguish between but not within categories. Figure 8.5 shows the similarity matrix for the semantic representations. Within-category similarities are reflected in four blocks along the main diagonal of the matrix. "Kitchen objects" are quite similar to each other, whereas "office objects" and "furniture" are less

Component Shape						
Features					Code	Description
1	1	1	1	1	cy	cylinder
1	1	1	1	0	cyh	cylinder—hollow
1	1	1	0	1	cys	cylinder—short
1	1	0	1	1	cyl	cylinder—long
1	1	0	1	0	t	top
1	1	0	0	1	lp	legs/prongs
1	1	0	0	0	slp	single leg/prong
1	0	1	1	1	sph	sphere
1	0	1	0	1	tap	taper-to-point
1	0	1	0	0	cu	curve
1	0	0	1	1	hsp	half-sphere
1	0	0	0	1	pc	plane—circular
1	0	0	0	0	r	rim
0	1	1	1	1	br	box—rectangular
0	1	1	1	0	pp	parallel planes
0	1	1	0	1	bt	box—thin
0	1	1	0	0	bl	box—long
0	1	0	1	1	ps	plane—square
0	1	0	1	0	pr	plane—rectangle
0	1	0	0	1	lf	long/flat
0	1	0	0	0	lft	long/flat/thin
0	0	1	1	0	liq	liquid
0	0	0	0	1	ho	hole
0	0	0	0	0		(no 2nd or 3rd component)

Table 8.3: The encoding used to describe the visual shape of each of the three (possible) components of each object. The meanings of the features are roughly (1) contains curves, (2) sides contain parallel lines, (3) sizes along all three dimensions of the same order of magnitude, (4) more equal in dimensions than shapes with similar values for preceding features, and (5) more regular than shapes with similar values for preceding features. The “Code” letters will be used to describe the assignment of visual features to objects.

Relative Position			
Features		Code	Description
1	1	ee	extension at end
1	0	em	extension at middle
0	1	ae	attachment at end
0	0	am	attachment at middle
0	0		(no 2nd or 3rd component)

Relative Size				
Features			Code	Description
1	1	1	l	longer
1	1	0	e	equal
0	1	1	s	smaller
0	0	1	ms	much smaller (1/4 to 1/2)
0	0	0	vms	very much smaller
0	0	0		(no 2nd or 3rd component)

Table 8.4: The encoding of the position and size of the second and third components relative to the main component.

General Characteristics	
Code	**Description**
dv	direction of main component—vertical
dh	direction of main component—horizontal
ss	screw/sawtooth
ifl	internal flexibility between components
con	concave
sh	sharp
dis	distortable
int	interior visible
dr	rectangle/handle apparent on surface

Color				
Features			**Code**	**Description**
1	1	1	va	various
1	1	0	brn	brown
1	0	1	si	silver
1	0	0	gr	grey
0	0	1	wh	white
0	0	0	tr	transparent

Texture			
Features		**Code**	**Description**
1	0	sm	smooth
0	1	ro	rough
0	0	ei	either

Absolute Size						
Features					**Code**	**Description**
1	0	0	0	0	s<3i	less than 3 inches
1	1	0	0	0	s3-6i	3 to 6 inches
1	1	1	0	0	s3-12i	3 to 12 inches
0	1	1	0	0	s6-12i	6 to 12 inches
0	1	1	1	0	s6i-2f	6 inches to 2 feet
0	0	1	1	0	s1-2f	1 to 2 feet
0	0	0	1	1	s2-6f	2 to 6 feet
0	0	0	0	1	s>6f	greater than 6 feet

Table 8.5: The coding for general visual characteristics, color, texture and absolute size. Each "general" characteristic is represented by a separate feature.

Assignment of Visual Features to Objects											
Object	Main	Second			Third			General	Color	Texture	Size
cup	cyh	pc	ae	e	cu	am	s	dv con	va	sm	s3-6i
spoon	lf	hsp	ee	ms				con	si	sm	s3-12i
pan	cyh	ps	ae	e	lf	ae	l	dv con	si	sm	s6i-2f
fork	lf	lp	ee	ms				sh	si	sm	s3-12i
knife	lf	lft	ee	e				sh	si	sm	s3-12i
bowl	hsp							dv con	va	sm	s6i-2f
can	cy							dv	va	sm	s3-6i
plate	pc	r	ee	ms				dh	wh	sm	s6-12i
dish	pc	r	ae	ms				dh con	va	sm	s6-12i
glass	cyh	pc	ae	e				dv con	tr	sm	s3-6i
pen	cyl	tap	ee	vms	t	ee	ms	ifl	va	sm	s3-6i
file	pr	pr	ae	e				ifl	va	ei	s1-2f
paper	pr							dis	wh	sm	s6-12i
book	pr	pr	ae	e				ifl int	va	ei	s6-12i
disk	ps	ho	am	vms					bl	ro	s3-6i
tape	cys	ho	am	s					brn	sm	s3-6i
stamp	ps							dis	va	sm	s<3i
board	pr								va	sm	s1-2f
glue	lft	tap	ee	vms					va	sm	s<3i
ink	cyh	t	ee	e	liq	am	s	con int	va	sm	s<3i
chair	ps	lp	ae	e	ps	ae	e	dh	brn	ei	s2-6f
table	pr	lp	ae	s				dh	brn	sm	s2-6f
bed	bt	lp	ae	vms	pr	ae	ms	dh	va	ro	s>6f
sofa	br	pr	am	e	pp	ae	s	dh	va	ro	s2-6f
stool	pc	lp	ae	l				dh	brn	sm	s1-2f
rug	pc							dh dis	va	ro	s>6f
radio	bt							dh	va	sm	s6-12i
tele	pr	br	am	e				dh	gr	sm	s1-2f
divan	bt							dh	va	ro	s2-6f
desk	pr	pp	ae	s				dh dr	brn	sm	s2-6f
saw	lft	cu						ss sh	si	sm	s1-2f
nail	cyl	tap	ee	ms	pc	ee	l	sh	gr	sm	s3-6i
plane	bl	cu	am	ms	lft	am	vms	sh	gr	sm	s6i-2f
ruler	lf								brn	sm	s1-2f
screw	tap	pc	ee	l				ss sh	gr	ro	s<3i
awl	cyl	tap	ee	e	sph	ee	ms	sh	gr	sm	s3-6i
axe	bt	tap	ee	e	cyl	am	l	sh	gr	sm	s1-2f
bolt	cy	cys	ee	l				ss	gr	ro	s<3i
nut	cys	ho	am	s				ss	gr	ro	s<3i
vice	pp	cyl	am	e				dv ifl int	va	sm	s1-2f

Table 8.6: A description of the visual representation of each object in terms of the codes used in previous Tables to describe values of each type of visual information.

Semantic features

Visual Characteristics	
1	main component vertical
2	main component horizontal
3	screw/sawtooth
4	internal flexibility
5	concave
6	sharp
7	distortable
8	interior visible
9	rectangle/handle apparent
10	color (visual coding)
11	color (visual coding)
12	color (visual coding)
13	smooth
14	rough
15	size less than 6 inches
16	size 6 inches to 2 feet
17	size greater than 2 feet
18	main-shape 1D
19	main-shape 2D
20	main-shape 3D
21	rectangular cross-section
22	circular cross-section
23	has legs
24	has other appendage
25	simple
26	complex
27	liquid
28	has hole

Consistency	
29	hard
30	soft

Made of	
31	metal
32	pottery
33	wood
34	cloth
35	glass
36	plastic
37	paper
38	other substance

Where found	
39	home
40	office
41	outdoors
42	kitchen/dining room
43	living room/study
44	bedroom
45	work-room
46	on ground
47	on surface
48	otherwise supported

General Function	
49	cooking
50	eating
51	drinking
52	leisure
53	rest
54	carpentry
55	work-office
56	work-home
57	aesthetic

Specific Function	
58	chopping/cutting
59	holding in place
60	writing
61	information-holding
62	measuring
63	reading
64	sticking
65	assigning-value
66	holding food/drink
67	sitting
68	lying
69	use-with-liquid
70	use-with-solid
71	sleeping
72	for comfort
73	for listening
74	for viewing
75	manipulating another artefact
76	is manipulated by another artefact
77	functioning with another object
78	functioning alone
79	container

General Action	
80	use with one arm
81	use with two arms
82	use with hand (little arm movement)
83	use involves mouth
84	easily breakable
85	placed in lap/held in front of body
86	characteristic action of whole body

Table 8.7: The semantic features used to described objects.

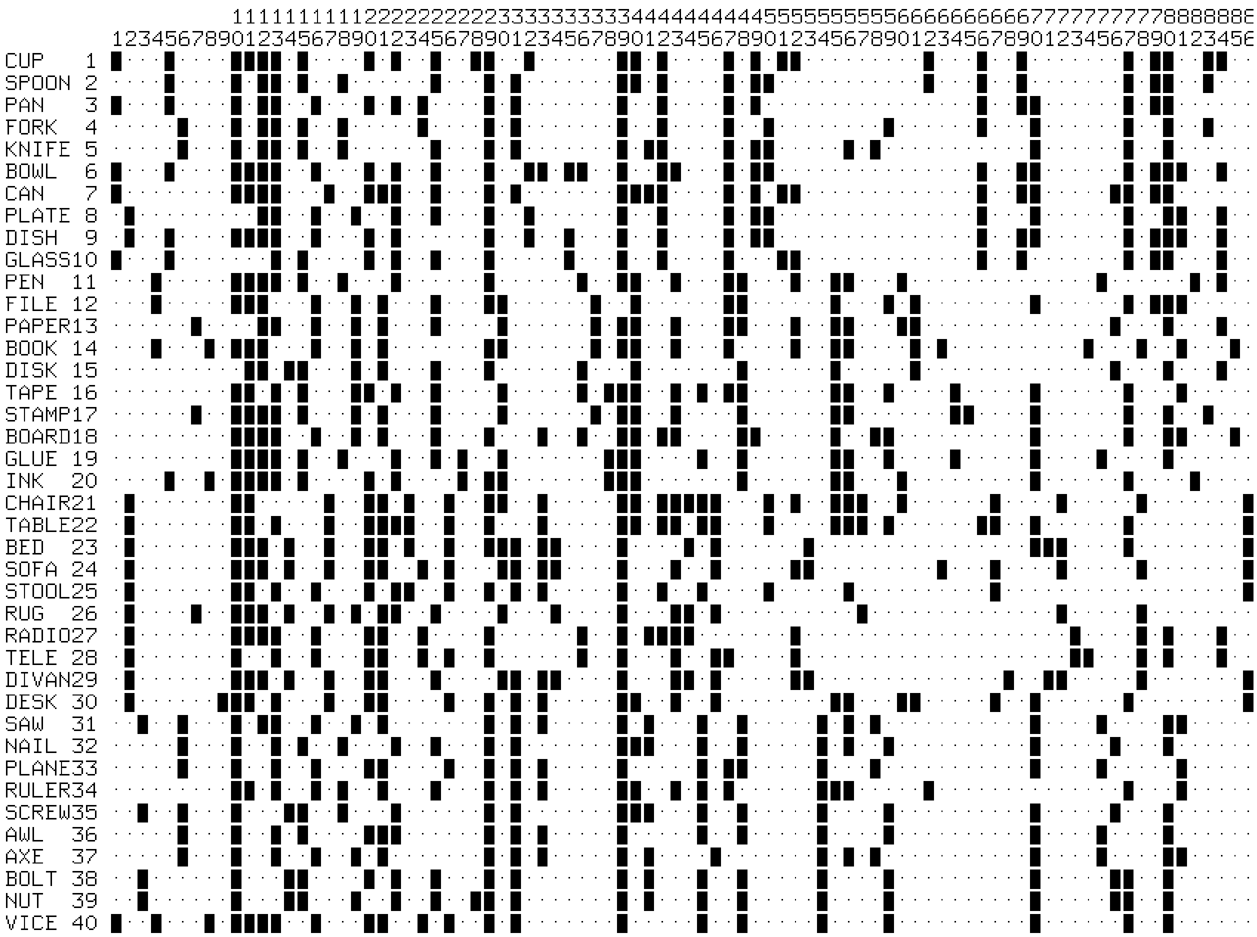

Figure 8.4: The assignment of semantic features to objects.

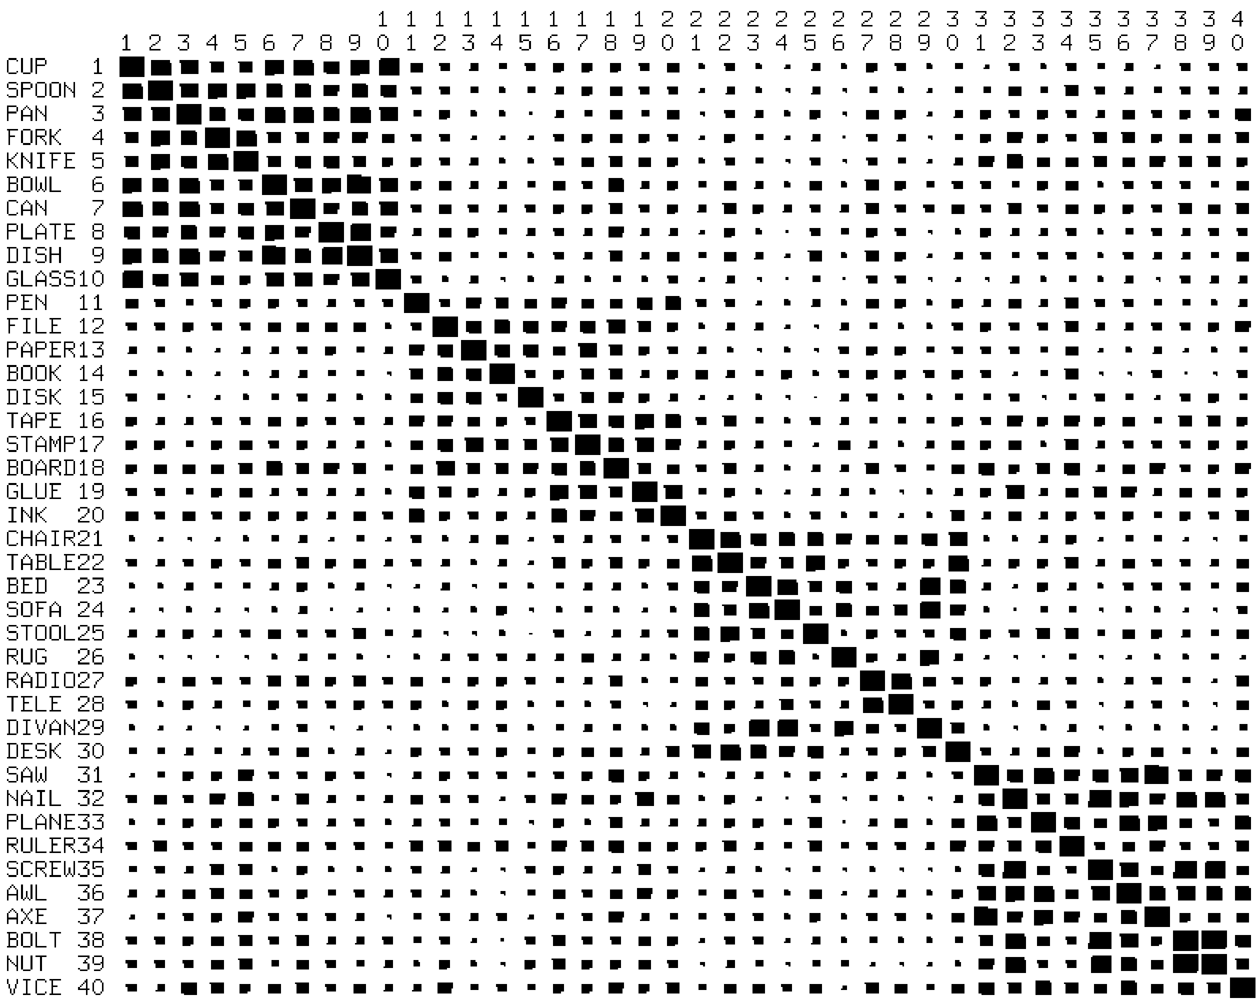

Figure 8.5: The similarity matrix for the semantic representations of the objects.

tightly clustered. In fact, *radio* and *tele* appear to be rather atypical items of furniture. “Tools” are relatively similar to each other, but also to some kitchen and office objects. In general, the semantic categories for objects are not as tightly clustered as those in the dyslexia simulations.

Despite these variations in the similarities within each category, we will use category membership as the basis for deciding if two objects (e.g. stimulus and response) are semantically related. The definition of visual relatedness of objects is somewhat more problematic as there is no simple analogue of letter overlap used with words. We can use the proximity of the visual representations of two objects to decide if they are visually related, but we then need a criterion. For semantic relatedness in the abstract/concrete word set, we adopted proximity criteria of approximately half-way between 1.0 and the mean proximity for all abstract or concrete word pairs (see Section 6.5). The mean proximity of all pairs of visual representations of objects is 0.465. Using a visual proximity criterion of half-way between 1.0 and this (0.73) yields the following chance rates of errors: visual $v = 0.0141$, semantic $s = 0.215$, mixed visual-and-semantic $m = 0.0154$, and other $o = 0.755$. Notice that $vs = 0.00304$ is much less than m. If visual and semantic similarity were unrelated, these two values should be about equal, as they are for the H&S word set (see Section 2.6.1). A more direct test of the relationship between visual and semantic relatedness is the correlation, over all pairs of objects, of visual and semantic proximity. In fact, there is a highly significant correlation between the visual and semantic similarity matrices for objects (0.52 ignoring diagonal terms, $t(1558) = 23.7$, $p < .001$). In contrast, the correlation for the H&S word set shows a slight *negative* trend (-0.04, $t(1558) = 1.56$, $p = .12$). Thus a major difference between our definition of object recognition as compared with word recognition is that there is significant structure in the mapping visual input to semantics for objects but not for words. This will prove important in explaining the rarity of visual errors in optic aphasic object naming compared with visual errors in deep dyslexic reading.

8.3.2 The network

The architecture of the network we will use to map visual representations onto semantic ones is identical to the 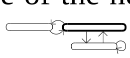 dyslexia network, except that the current network has 44 “visual” input units and 86 semantic output units (see Figure 8.6). Thus the network has a direct pathway through 40 intermediate units, a clean-up pathway via 40 additional units, and a feedback pathway from semantics to the intermediate units. There are no intra-sememe connections, and all pathways have 25% connectivity density, for a total of 4492 connections.

In addition to the conventional slowly learning weight, each connection has a short-term correlational weight that operates as described above.

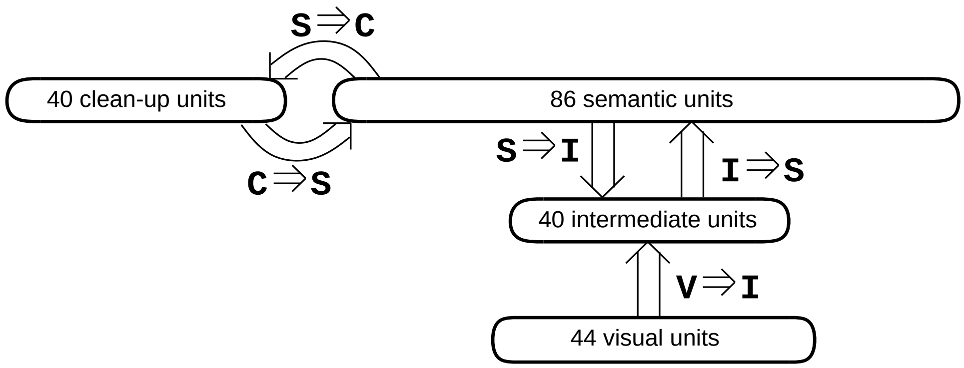

Figure 8.6: The architecture of the optic aphasia network. Notice that the set of connections from the visual (input) units to the intermediate units is labeled V⇒I rather than O⇒I.

8.3.3 The training procedure

The network was trained using back-propagation to activate each of the appropriate semantic units to within 0.2 of its correct value over the last three of eight iterations when presented with the visual representation of each object. At the end of processing each object, the short-term weights were modified according to Equation 8.1. This consisted of averaging (with $\lambda = 0.5$) the previous short-term weights with those specified by the current unit activities. In this way each object was presented in the context of the outcome of the presentation of the previous object. Objects were chosen randomly without replacement for presentation during a sweep to ensure that they were all presented equally often and in an unbiased order. The network must derive a set of weights that enables it to recognize each object when preceded by each other object. To the extent that the unit correlations for one object are unrelated to those for another, the short-term weights effectively act like noise in the weights, forcing the network to develop stronger semantic attractors with the long-term learning weights.

Although the operation of the network is deterministic, the random order of object presentations causes performance to vary somewhat over successive training sweeps. However, the network reliably satisfied the training criteria after about 10,000 sweeps through all 40 objects.

8.3.4 The lesioning procedure

After the network had learned to recognized each object, we subjected each of its major sets of connections to lesions of the standard range of severity. We have not implemented a phonological output system that would map the semantics of objects onto their pronunciations. Thus we must resort to applying distance and gap criteria directly to semantics. We will continue to use a distance criterion of 0.8 and a gap criterion of 0.05. The expected proximity of random vectors decreases with increasing dimensionality, so these criteria are somewhat more stringent in the current context (with 86 semantic features) than when we applied them to the semantics of words (with 68 features). While this will lower the overall rate of explicit responses (correct or error), it should not significantly bias the distributions of error types (see Hinton & Shallice, 1991).

We would like to measure the performance of the damaged network on each object as stimulus when preceded by every other object. We will refer to the preceding object as the "prime." One possible procedure for gathering data is to administer a particular lesion, and then measure performance on all of the objects after using each of them in turn as the prime (i.e. setting the short-term weights based on the unit correlations when the prime object is presented to the damaged network). However, this procedure has the drawback that the pattern of errors will be quite similar for each different prime—the tendency for a particular lesion to cause particular errors may dominate any perseverative effects. The alternative procedure that we adopt is to administer a different lesion for each prime object. This means that data is gathered over 800 instances of a particular type of lesion (40 primes $\times$ 20 lesions per prime) rather than just 20. In this way, the effects due to particular lesions are better sampled, while still enabling perseverative effects to emerge. Although the first procedure is more analogous to the testing situation for an individual patient, the latter should produce results that better reflect the extent to which lesions to the network *in general* produce behavior like that found in optic aphasia.

For each lesion, we presented the prime with the short-term weights to zero, and then set the short-term weights on the basis of the resulting unit activities. We then presented each object in turn (with the short-term weights fixed), and classified the response of the network as correct, an error, or an omission.

8.3.5 Correct performance

Figure 8.7 presents the correct performance of the network after lesions to each set of connections, as a function of lesion severity. By comparison with the 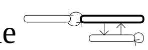 network for mapping orthography to semantics (see Figure 4.13, p. 96), the optic aphasia network is remarkably sensitive to lesions, particularly to the clean-up pathway. Even lesions to the feedback connections $\mathtt{S} \Rightarrow \mathtt{I}$ produce significant impairment in correct performance, whereas in the 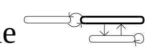 network they had negligible effect even at the highest severities. Clearly the network finds it quite difficult to correctly recognize objects when the short-term weights provide a bias towards previous objects. It relies heavily on the clean-up pathway(s) to overcome this bias, but under damage it often fails.

8.3.6 Horizontal errors

We are concerned with two types of effects in the errors produced by the network under damage, roughly corresponding to Lhermitte & Beauvois' (1973) distinction of "horizontal" vs. "vertical" errors. Horizontal errors consist of the standard comparisons of the relatedness of the stimulus and response. For these we will use the definitions of visual and semantic relatedness described above,

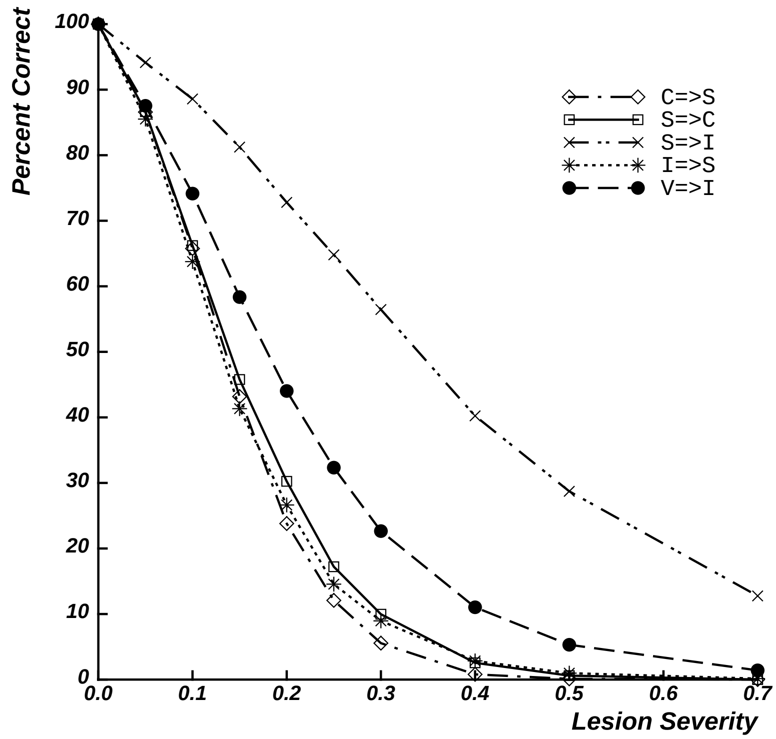

Figure 8.7: Overall correct performance of the optic aphasia network after lesions to each main set of connections, as a function of lesion severity.

and classify errors as visual, semantic, mixed visual-and-semantic, and other in the standard way. We address vertical (perseverative) errors in the following section.

Figure 8.8 presents the rates of each type of error after lesions resulting in correct performance between 20-80%, as well as the distribution of types for error responses chosen randomly from the set of objects. The total error rates are fairly low, ranging from 5.0% for V⇒I lesions to 0.5% for C⇒S lesions. For all lesion locations there is a strong bias towards semantic errors. This is made clear by comparing the ratio of semantic to "other" errors for each lesion location with the ratio for the "chance" error distribution. The observed ratios for V⇒I, I⇒S, and S⇒I are nearly 10 times larger than the chance ratio; for clean-up lesions they are nearly 20 times larger. The comparison for mixed visual-and-semantic errors is even more dramatic. The observed ratio for V⇒I lesions is over 40 times the chance value. For I⇒S and S⇒I it is over 100 times the chance value, and for clean-up lesions it is over 250 times the chance value. These factors are much larger than any observed in the dyslexia networks. Thus the network shows a remarkable bias towards semantic and mixed similarity in its errors relative to the chance error distribution. In contrast, the observed ratios for visual errors do not exceed the chance ratio to nearly the same degree, although they are still much larger (8 times for V⇒I, 6 times for I⇒S, 9 times for S⇒I and S⇒C, 4 times for C⇒S). Thus visual errors occur at an above-chance rate, but the network is much more strongly biased towards making semantic and mixed visual-and-semantic errors.

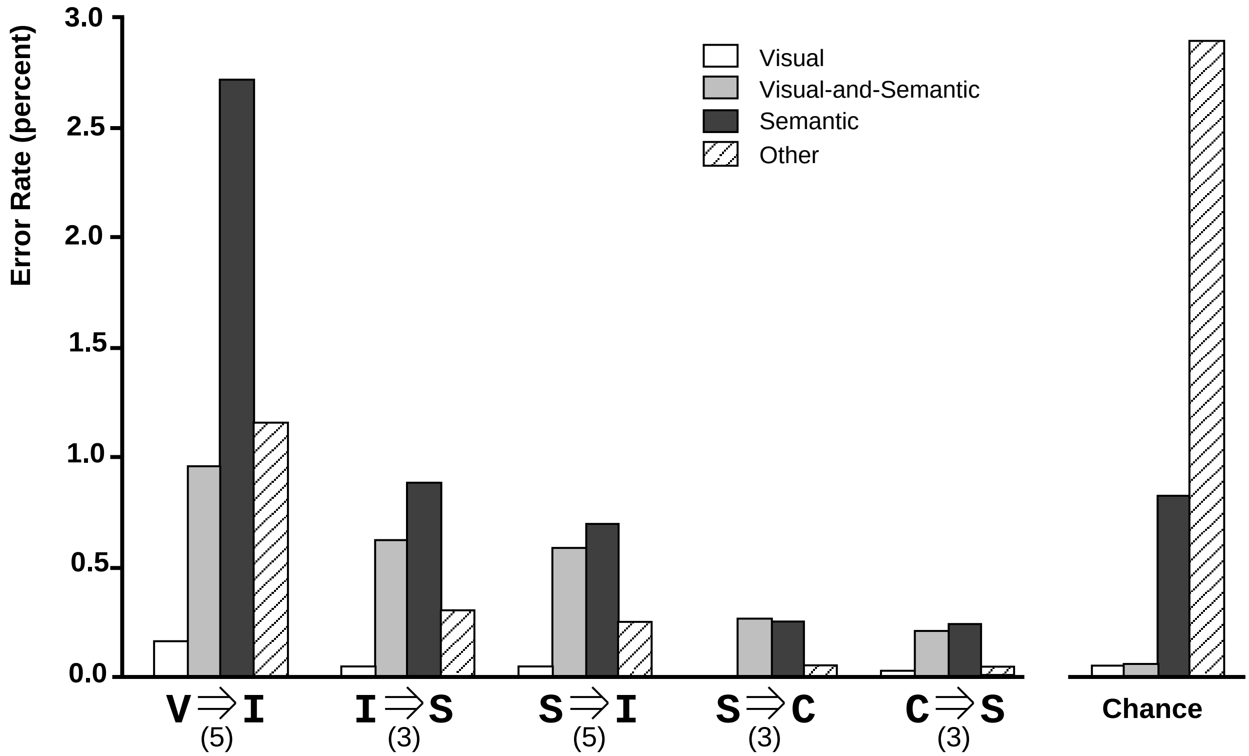

Figure 8.8: The distribution of error types for lesions to the optic aphasia network producing correct performance between 20-80%, as well as the distribution of types for error responses chosen randomly from the set of objects.

It is important to emphasize that these measures are relative to the chance distributions of the error types, and so the results are not due to the relative rarity of purely visual errors among all potential errors. In fact, potential mixed visual-and-semantic errors are just as rare, and yet they occur over 5 times more frequently than visual errors after `V`⇒`I` lesions, over 10 times more frequently after `I`⇒`S` and `S`⇒`I` lesions, over 30 times more frequently after `S`⇒`C` lesions, and 75 times more frequently after `C`⇒`S` lesions.

Why does the network show such a strong bias towards semantic and mixed visual-and-semantic errors relative to visual errors? Another way to phrase this question is, why is the bias towards semantic vs. visual similarity in errors so much stronger in the optic aphasia network than in the dyslexia networks? It is unlikely that it relates directly to the presense of short-term weights in the optic aphasia network. A bias in errors towards the previous object would match the chance distribution of error types, since all combinations of primes and stimuli were tested. Rather, it more likely relates to a difference in the nature of the tasks of object recognition and word recognition. Earlier we claimed that the relationship between visual and semantic representations is more structured for objects than it is for words, both in general and in our versions of the tasks. That is, two objects with similar visual forms are more likely to have similar meanings and functions than are words that share letters. In a sense this follows the distinction made by Gibson (1979) that shapes have particular "affordances"—they allow for certain types of manipulations and actions independent of specific knowledge of their identity. This distinction was echoed by Riddoch & Humphreys (1987) to explain the relative preservation of gesturing after impaired semantic access (under their interpretation) in optic aphasia.

On this basis, one might be tempted to explain the predominance of semantic errors in optic aphasia on the basis that visually similar objects are also often semantically similar, and so would often be classified as mixed visual-and-semantic errors. The classic example of a semantic category with high visual similarity among its members is "animals." In essence, the "chance" rate of visual similarity without semantic similarity is much lower for object naming than for word reading. In fact, we have attempted to reflect this in our criteria for visual relatedness in the two tasks. However, this explanation does not account for why associative agnosics produce predominantly *visual* errors in object naming under the same criteria for visual relatedness. And in fact, our optic aphasia network shows a bias towards semantic and mixed errors *relative* to their chance distributions. The explanation for the predominance of semantic errors in optic aphasia, both in patients and in the network, must relate to the nature of the representations and processes that are developed in learning a more structured task.

We explain word reading errors in deep dyslexia, and object naming errors in optic aphasia, in terms of the same computational principles: damage to an attractor network causes the initial pattern of activity generated by a stimulus to be "captured" by the attractor for a related stimulus. Although we demonstrated in Chapter 4 that these attractors need not be *semantic*, they are in the optic aphasia network and most of the dyslexia networks, and it will clarify the following description to refer to them as such. In mapping orthography to semantics, there is strong pressure to position and shape the attractors so as to separate the initial semantic activity for visually similar stimuli into their quite distinct, final semantics (see Figure 2.10, p. 43). In fact, for many words (e.g. BUG), visual similarity with words in other categories (e.g. MUG, BOG) is often much greater than with any of the words in the same category (e.g. DOG, PIG). The result is that there are large areas within semantic space where the attractor basins for purely visually related words adjoin, providing ample opportunity for visual errors. In contrast, in object naming it is less common that visually similar objects must be separated into completely different semantics. Furthermore, even for visually similar, semantically distinct pairs of objects (e.g. *fork* and *awl*), typically there are other objects within each category that are just as visually similar (e.g. *spoon* and *screw*, respectively). When the initial semantics for *fork* is corrupted by damage, the additional bias of semantic similarity makes the mixed error *fork* ⇒ "spoon" much more likely than the visual error *fork* ⇒ "awl". Thus potential visual errors are often preempted by semantic or mixed visual-and-semantic errors.

In summary, similar inputs are free to have similar effects on the output in a structured task, and so there is much less direct influence of input similarity as distinct from output similarity in the layout of attractor basins developed by a network. Semantic attractors must still provide clean-up to compensate for variations in the initial semantic activity, due either to noise or the short-term weights from the previous object. Since this clean-up is semantically structured, the naming errors produced after damage show predominantly semantic similarity.

8.3.7 Vertical errors

One of the more interesting aspects of the naming errors of optic aphasics is that they are biased by the responses given to previously presented objects. These perseverative errors, termed "vertical" by Lhermitte & Beauvois (1973), are most frequently identical to previous responses, but can also be semantically related. Mixed visual-and-semantic perseverative errors occur as well, but purely visual perseverations have not been documented.

In the preceding section we analyzed error responses based only on the relationship between stimulus and response—so-called "horizontal" errors. Classifying vertical errors is more complicated as they involve the relationship of the stimulus and response with previous objects. For simplicity we will confine ourselves to considering the effects of only the immediately preceding object, which we call the "prime." This simplification also applies to most of the errors produced by patients. We will also consider only *semantic* relatedness between the prime and the stimulus and/or response. Finally, if the damaged network misnames the previous object, we will consider the *named* object to be the "prime" for the purposes of comparison with the stimulus and response.

Figure 8.9 presents the possible relationships between stimulus, response, and prime, and their classification into error types. Each type is labeled in two parts. The first part (before the "+")

	Relationship			Example			
Type	Stim-Resp	Prime-Stim	Prime-Resp	Prime	Stim	$\Rightarrow$	Resp
S+P	semantic		identical	"spoon"	*fork*	$\Rightarrow$	"spoon"
S+$\overline{\text{P}}$	semantic	identical	semantic	"spoon"	*spoon*	$\Rightarrow$	"fork"
S+C	semantic		semantic	"spoon"	*fork*	$\Rightarrow$	"knife"
S+U	semantic		none	"spoon"	*table*	$\Rightarrow$	"chair"
O+P	none		identical	"spoon"	*table*	$\Rightarrow$	"spoon"
O+$\overline{\text{P}}$	none	identical	none	"spoon"	*spoon*	$\Rightarrow$	"table"
O+C	none		semantic	"spoon"	*table*	$\Rightarrow$	"fork"
O+$\overline{\text{C}}$	none	semantic	none	"spoon"	*fork*	$\Rightarrow$	"table"
O+U	none		none	"spoon"	*table*	$\Rightarrow$	"nail"

Figure 8.9: The possible types of errors based on semantic relatedness between stimulus, response and prime.

refers to the nature of the stimulus-response error:

- S Stimulus and response are semantically related (includes mixed visual-and-semantic errors).
- O Stimulus and response are not semantically related (includes visual errors).

The second part refers to the nature of the perseveration, depending on the relationship between the prime and the stimulus and/or response:

- P The response is identical to the prime ("perseveration").
- $\overline{\text{P}}$ The stimulus is identical to the prime but the response is not. In this case the prime-response relationship is *opposite* to an item perseveration.
- C The response is semantically related ("coordinate") but not identical to the prime.
- $\overline{\text{C}}$ The stimulus is semantically related to the prime but the response is not. Thus the response goes against a semantic perseveration.
- U The response and prime are unrelated.

The easiest way to understand these types is in terms of the different influences that contribute to the error. An S+U error (e.g. ("spoon") *table* $\Rightarrow$ "chair") is a standard semantic error with no perseverative influence, while an O+P error (e.g. ("spoon") *table* $\Rightarrow$ "spoon") is a repeated response unrelated to the stimulus. In an S+C error (e.g. ("spoon") *fork* $\Rightarrow$ "knife"), both horizontal (stimulus-response) and vertical (prime-response) semantic similarity contribute to the error. In an

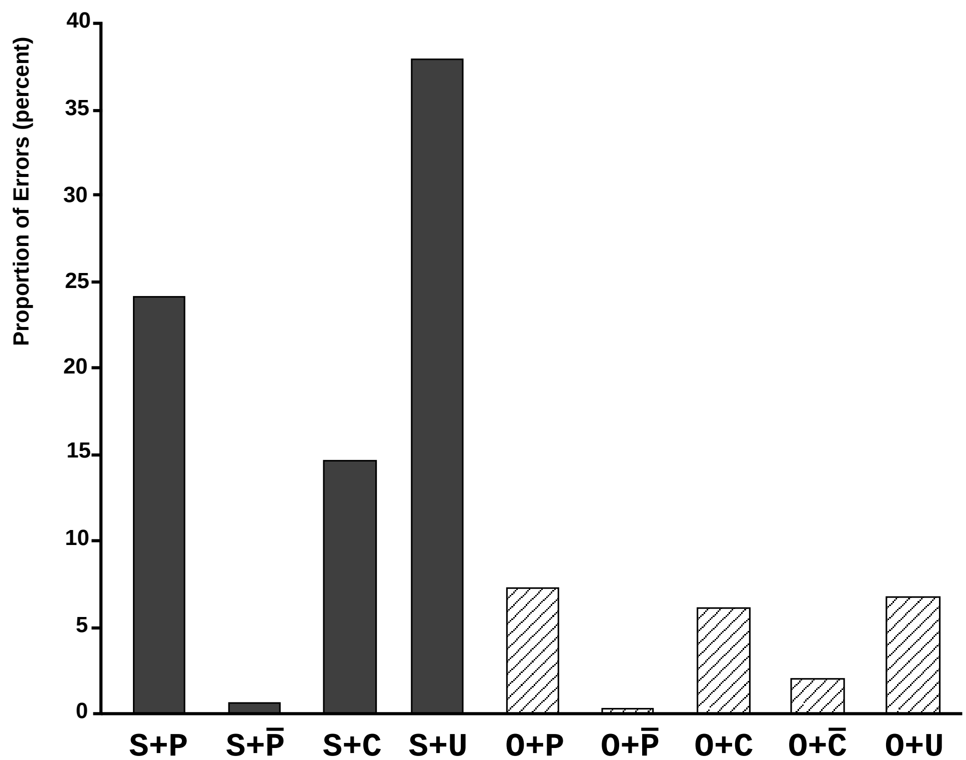

Figure 8.10: Distribution of perseverative error types, averaged over all lesion locations and severities producing correct performance between 20-80%.

error involving the perseverative relation $\overline{\mathrm{P}}$ or $\overline{\mathrm{C}}$, the stimulus is consistent with the perseverative effect but the response is not—thus the error is contrary to the perseverative influence.

Figure 8.10 shows the distribution of these perseverative error types produced by lesions to the optic aphasia network, averaged over all lesion locations and severities producing correct performance between 20-80%. The predominance of semantic errors is clear in the figure, with 77.1% of all error response being semantically related to the stimulus. Among these, the most common is the conventional "horizontal" semantic error with no perseverative relationship (S+U). However, since our current concern is with perseverative influences, we will first consider the remaining ("O") errors. These errors provide the clearest picture of perseverative influences because there is no confounding bias from semantic relatedness of the stimulus and response.

7.4% of the network's errors are pure perseverations (O+P), in which the previous response is repeated even though it bears no relation to either the current stimulus or response. If responses were generated randomly with no perseverative influence, they would have a 1 in 40 chance (2.5%) of being identical to the prime. Thus the observed rate of response perseverations is about three times the chance rate. Another indication of the strong perseverative influence is that it is extremely rare (0.15%) for the prime itself to produce an unrelated response when presented as the stimulus (O+$\overline{\mathrm{P}}$), compared with the rate of unrelated errors when there is no perseverative relationship (6.8%, O+U). Thus the prime is exerting a strong bias on the nature of the response independent of any relationship with the stimulus.

The network also produces many “semantic perseverations” (O+C), in which, rather than the prime itself, an object that is semantically related to prime is given as a response unrelated to the stimulus (e.g. (“spoon”) *table* ⇒ “fork”). In contrast, the network is much less likely to produce an unrelated response when the *stimulus* is semantically related to the prime (O+$\overline{\mathrm{C}}$, e.g. (“spoon”) *fork* ⇒ “table”). The rate of these errors (2.5%) is also much less than the completely unrelated (O+U) errors. Thus the prime biases the network towards responses that are semantically related to it—this increases O+C errors and decreases O+$\overline{\mathrm{C}}$ errors.

Now consider the errors in which the stimulus and response are semantically related. Among those for which there is a perseverative influence, the most common (24.1% of all errors) are response perseverations (S+P, e.g. (“spoon”) *fork* ⇒ “spoon”). Semantic perseverations (S+C) are somewhat less common (14.6%). Thus the prime induces a strong bias towards an identical response to the next object rather than simply one in the same category. Also notice that it is very rare (0.59%) for the prime to produce another object in the category rather than itself (S+$\overline{\mathrm{P}}$) when presented as the stimulus.[4] Thus even within a category, the prime biases responses towards itself compared with other objects in the category.

8.3.8 Effects of type of response to the preceding object

For the purposes of categorizing perseverative errors, we have define the “prime” to be the response given to the preceding object (correct or an error), or the object whose semantics are nearest those generated by the network in the case where presentation of the preceding object resulted in an omission. While we have grouped these conditions together in the analysis presented above, it would seem likely that the influence that the prime has on the naming of subsequent objects would vary considerably with how well the network responded to the prime itself. To investigate this possibility, we separated errors based on how the damaged network responded to the prime.

Figure 8.11 presents the same data on the distributions of perseverative error types, now separated by whether the network named the preceding object correctly, made an error, or failed to respond. Consider the balance of responses that are identical to the prime (P) vs. those that are unrelated (U), both for semantic errors (S) and other errors (O). This provides a rough measure of the “strength” of perseverative influences. First notice that the proportions of all errors that are response perseverations (S+P and O+P) are much lower when the prime is an omission than when it is an explicit response. Conversely, the proportion of errors showing no perseverative influence (S+U and O+U) are much higher when the prime produces no response as a stimulus. In fact,

[4]The error type S+$\overline{\mathrm{C}}$ cannot occur because semantic relatedness—defined by category membership—is transitive. Thus it is impossible for the stimulus and response to be semantically related (S) and the stimulus and prime to be semantically related ($\overline{\mathrm{C}}$) without the response and prime being semantically related (which would result in an S+C error).

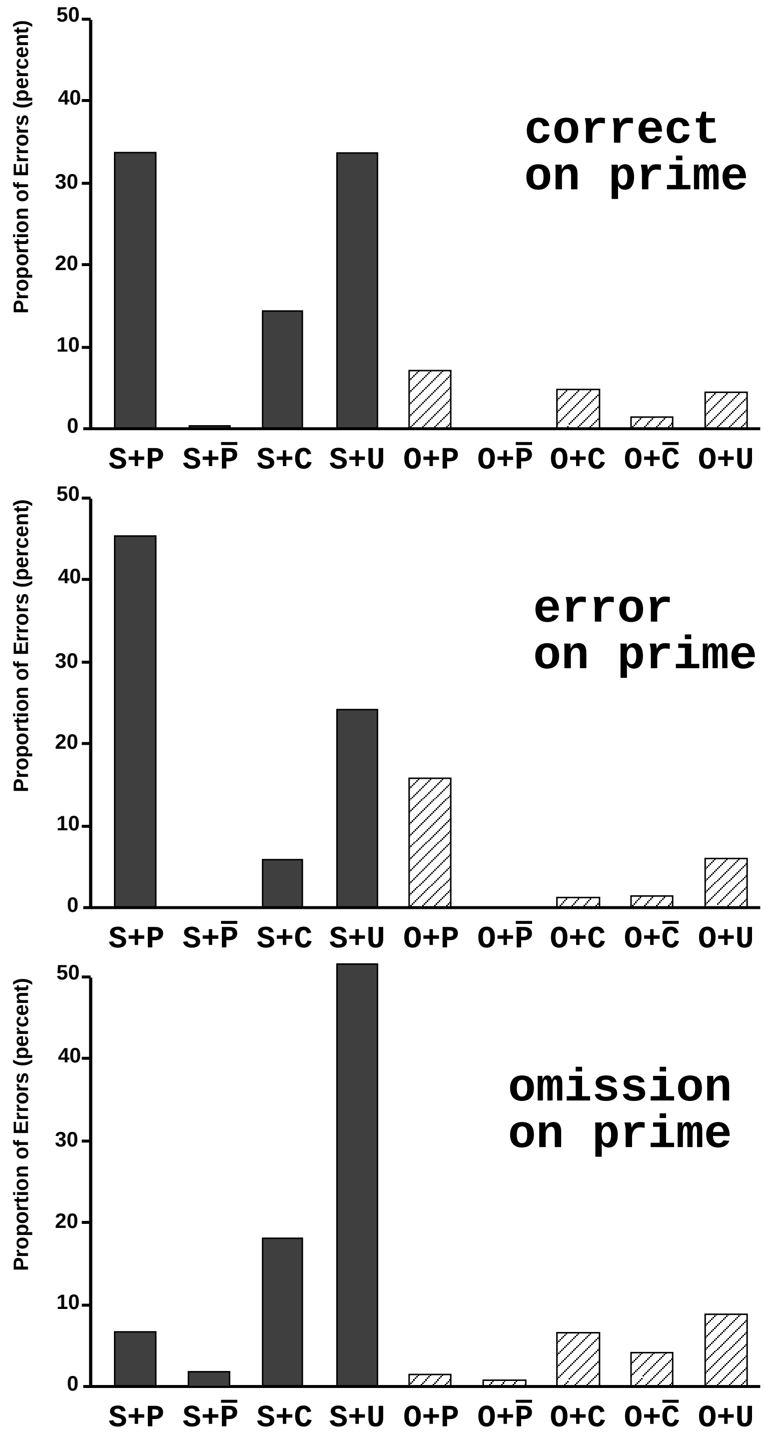

Figure 8.11: Distributions of perseverative error types following prime presentations resulting in correct responses, errors, and omissions.

the proportion of all error response that are identical to the prime is 41.6% for primes producing explicit response vs. only 8.4% for primes producing omissions ($F(1, 6100) = 894.9$, $p < .001$). Thus when an object generates no response, it also has less influence on the naming of subsequent objects than when it generates a correct or error response. This makes sense given that omissions occur when the semantics generated by the network do not match the nearest object very well. When the short-term weights are set on the basis of poorer semantics, they provide a weaker bias towards the prime than when set by more accurate semantics. Thus the combined data reported in Figure 8.10 significantly underestimates the perseverative influence from previous objects that evoke an explicit response.

Even among explicit responses, there are some interesting differences between correctly vs. incorrectly named primes. When the preceding object is named correctly, the proportion of errors involving a response perseveration is about equal to that involving no perseverative relation (both for semantic and other errors). In contrast, when the preceding object is named incorrectly, the next object is much more likely to elicit the same response. In fact, the proportions of semantic perseverations not involving the same exact response (S+C and O+C) are lower when the preceding object is misnamed. The proportion of errors identical to the prime is 61.2% for primes producing errors vs. 40.9% for primes named correctly ($F(1, 4352) = 31.0$, $p < .001$). In essence, the attractor for the incorrect response has become abnormally strong as a result of damage, producing the error the the preceding object. When the short-term weights are set on the basis of this objects semantics, there is even stronger pressure for other objects to succumb to the same attractor.

The perseverative responses of optic aphasics can come from correctly named objects, incorrectly named objects, or even objects generating no response (Gil et al., 1985; Lhermitte & Beauvois, 1973). However, the proportions of error types that follow each of these conditions has not been analyzed in detail, so it is difficult to compare the network's behavior with that of patients in more than a qualitative manner.

8.3.9 Effects of lesion location

The data from different lesion locations are averaged together in the results presented above. However, the distribution of perseverative error types differs significantly as a function of the location of damage in the network. Figure 8.3.9 presents the distributions of these error types separately for each lesion location. The pattern for `V`⇒`I` lesions is most similar to that for the entire network because the largest proportion of errors (52.9%) occur after these lesions. There is an interesting progression in the error pattern as the lesion location moves closer to semantics. For `V`⇒`I` lesions, most semantic errors show no perseverative relationship, while other errors are about balance between O+P and O+U. For these lesions, the proportion of error responses that are identical to the prime is 22.5%. For lesions to connections between the intermediate units and semantic (`I`⇒`S` and `S`⇒`I`), S+P and S+U errors are more balanced, and a higher proportion of errors

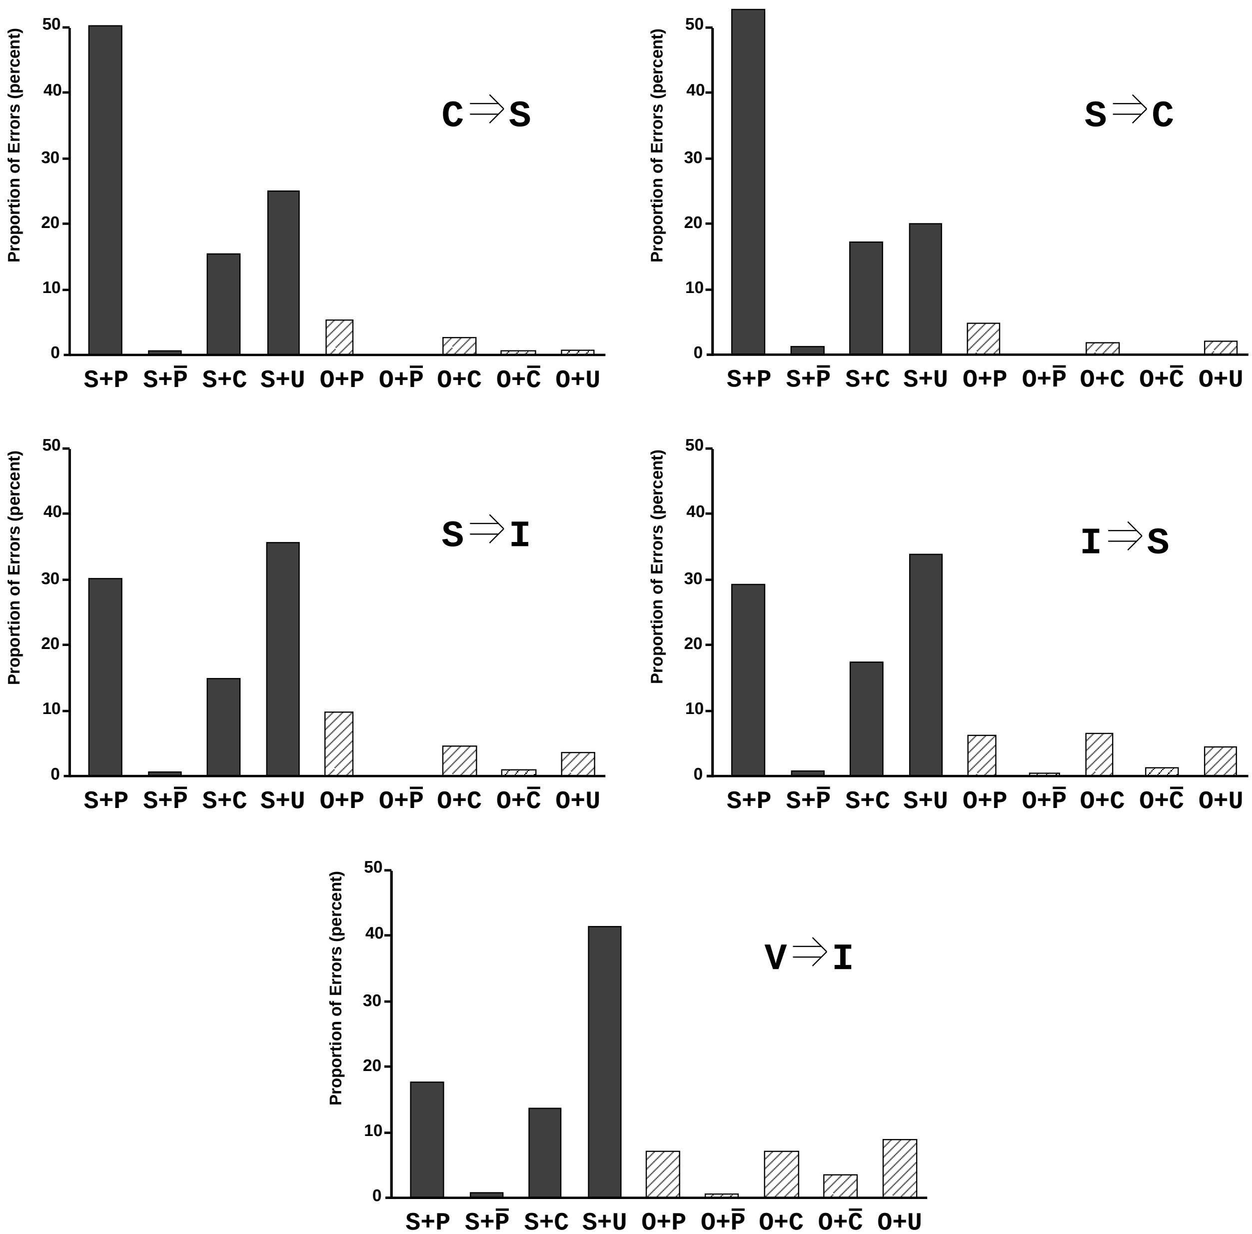

Figure 8.12: Distributions of perseverative error types for each lesion location.

are identical to the prime (37.6% for I⇒S, $F(1, 3815) = 96.1$, $p < .001$ vs. V⇒I; 37.4% for S⇒I, $F(1, 4804) = 161.1$, $p < .001$ vs. V⇒I; $F < 1$ for I⇒S vs. S⇒I). For clean-up lesions, many more semantic errors are perseverations than are unrelated, with the same holding true (although to a lesser extent) for other errors. The proportion of error responses that are identical to the prime is 62.7% for S⇒C ($F(1, 972) = 59.5$, $p < .001$ vs. I⇒S) and 58.7% for C⇒S ($F(1, 1024) = 46.6$, $p < .001$ vs. I⇒S; $F(1, 590) = 1.05$, $p = .31$ for S⇒C vs. C⇒S). Thus the strength of the perseverative influence increases as lesions move closer to semantics. However, the proportion of semantic perseverations (S+C and O+C) are much less affected by lesion location.

Why should lesions near or within semantics produce a stronger bias towards response perseverations than lesions closer to the input? Clean-up lesions corrupt the semantic attractors for objects, resulting in far fewer overall naming errors than do lesions to the direct pathway (see Figure 8.8, p. 252). After clean-up lesions, the preceding object is named correctly on about two-thirds of the trials.[5] When this occurs, the short-term weights within the clean-up pathway are set in a way that magnifies the clean-up influences which generated the semantics of that particular object. This additional bias has more influence after clean-up lesions compared with direct-pathway lesions because the normal clean-up influences are diminished after the former but not the latter. The bias towards the semantics of the preceding object can dominate the weakened clean-up for the correct semantics of the stimulus, causing the network to more frequently produce a response perseveration. Even when the semantics generated by the network in response to the preceding object do not satisfy the response criteria, they still evoke short-term weights that bias the network on the next trial towards the semantics of that object compared with others.

8.3.10 Item effects

How do the semantic and perseverative effects vary across objects? Are only a few objects responsible for the majority of perseverations? To investigate this, we re-analyzed the data for individual objects when presented as stimuli, when given as responses, and when serving as the prime.

Figure 8.13 presents the correct performance of each object as stimulus, averaged over all lesion locations and severities producing overall correct performance between 20-80%. Performance on individual objects varies quite considerably, from only 15.4% correct on *disk* up to 77.8% correct on *pan*. Nonetheless, correct responses are fairly evenly spread throughout the set of objects. Although there is the suggestion of differences in correct performance between categories, with "tools" named more poorly than other objects, this difference falls just short of significance when analyzed with objects as the random variable ($F(3, 36) = 2.31$, $p = .09$).

Now consider the errors made by the network. For each object, we can determine how often

[5] This proportion is much higher than the average correct performance for subsequently presented objects (45.8%) because the preceding object is presented with all of the short-term weights set to zero.

Figure 8.13: Average correct performance for each object. The bars represent the average performance for each category.

it produces an error when presented as stimulus, as well as how often it is given as the incorrect response in an error. Figure 8.14 shows the proportion of semantically related errors for each object as stimulus and as response. The proportions of semantic errors produced by most objects as stimuli are reasonably near 2.5%, which is the expected value for 40 objects if these errors were uniformly distributed across objects. However, there are a few objects that produce a disproportionate number of semantic errors—the top six objects (15%) account for 36.3% of all semantically related errors. Although there appear to be overall differences between categories, the variability prevents these differences from reaching significance ($F(3, 36) < 1$). Similarly, the responses in semantic errors are distributed across most of the objects, although there are a few error responses (e.g. *disk*, *ink*, *ruler*, *vice*) that are rarely if ever occur. Conversely, *saw* is given as the response in over 13% of all semantic errors; *divan* and *desk* are also common responses. Thus it appears that semantic errors are not confined to particular objects as stimuli or responses, although damage can cause the attractors for some objects to dominate those of many others.

Figure 8.15 presents the confusion matrix for all responses produced by the network after damage. The size each square in the matrix represents the frequency which which the stimulus listed on the left produces the response numbered across the top. On-diagonal values represent correct responses; off-diagonal values represent errors. Semantic errors are represented in each of the four 10-by-10 blocks along the main diagonal—the predominance of these errors is reflected by the fact that most errors are within one of these blocks. The two most common errors are *table*

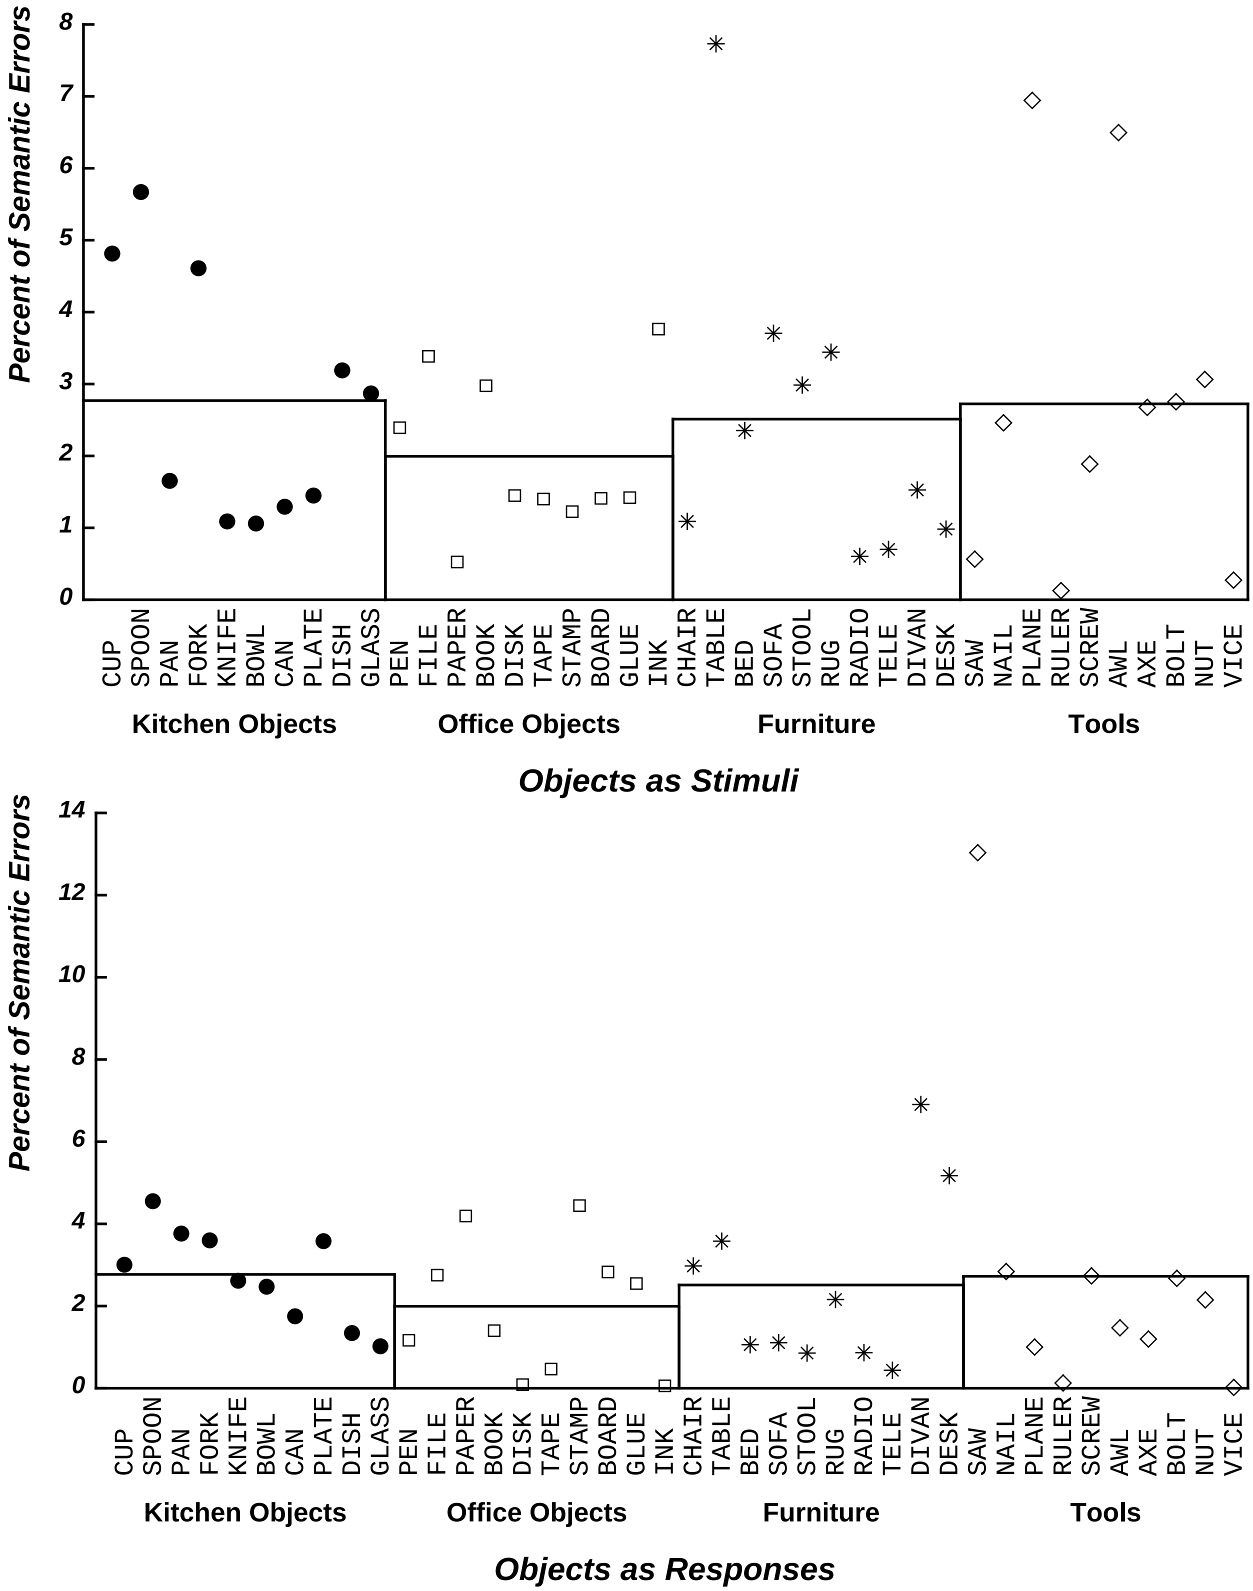

Figure 8.14: The distribution of semantically related errors for each object as stimulus (top) and response. Notice that plots are scaled differently.

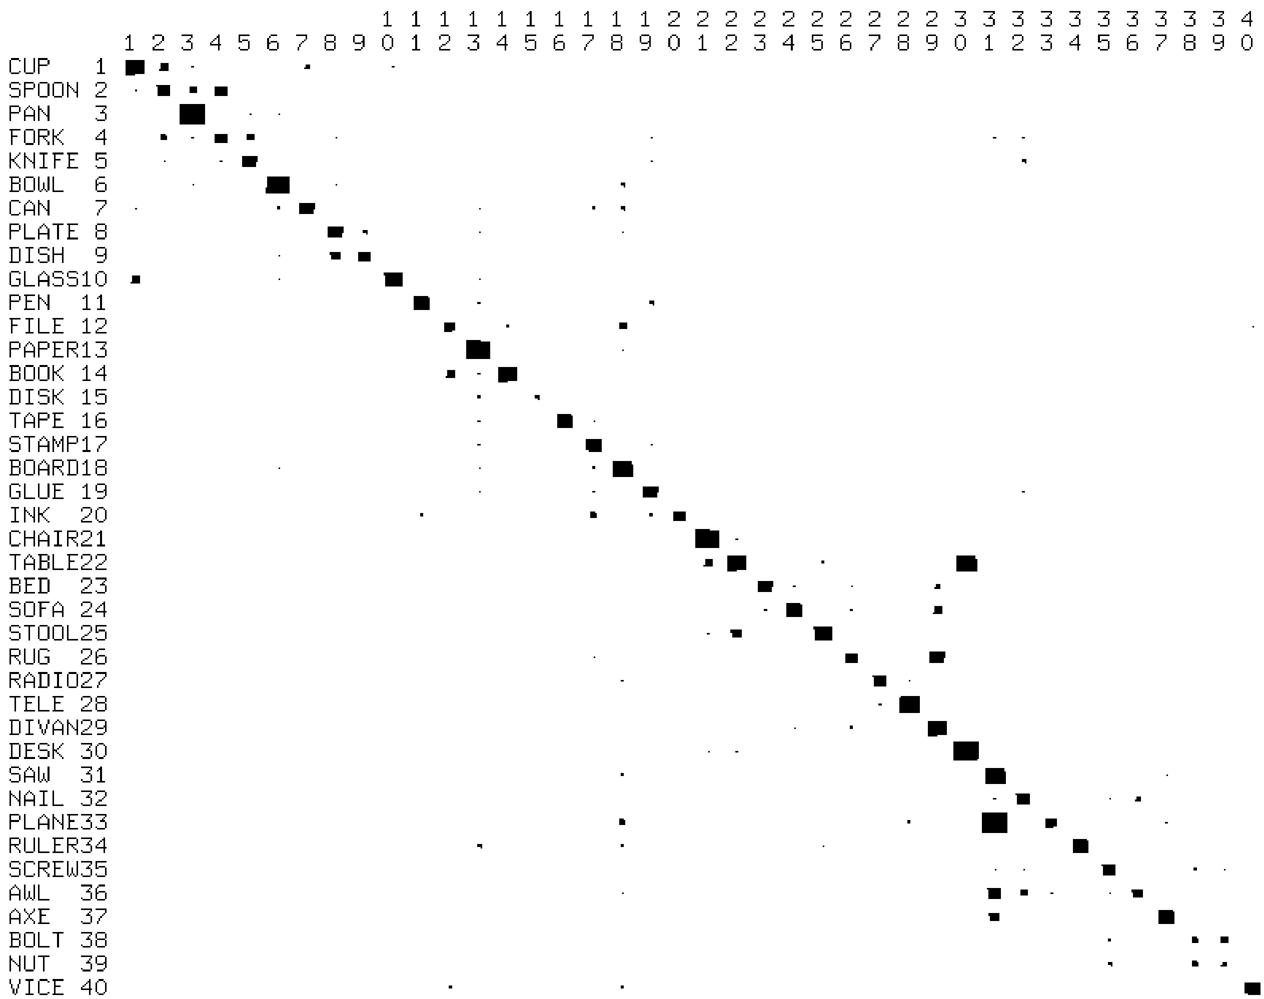

Figure 8.15: The confusion matrix for responses produced by the network after damage. These two sets of values are scaled separately relative to the highest frequency stimulus-response pair among each.

Figure 8.16: The distribution of perseverative errors (S+P and O+P) for objects as primes.

⇒ "desk" and *plane* ⇒ "saw". Many other errors also occur, but the frequency of some of them relative to the most common errors is too low to be strongly indicated in the display. The magnitude of the squares vertically within a column reflects the overall "strength" of the corresponding object as a response independent of the particular stimulus presented. Thus *saw* (column 31) is a strong response for many semantic errors within "tools," while *board* (column 18) is given in response to many other objects in other categories. Overall, the matrix indicates that a few object pairs account for a high proportion of the errors produced by the network, but that many other pairs also contribute to a lesser degree.

In addition to the relationship between stimulus and response, we can also consider the relationship between the prime and response—in particular, the occurrence of perseverative errors (S+P and O+P), in which the prime and response are identical. Figure 8.16 presents the distribution of perseverative errors for objects as primes. A few objects (e.g. *disk*, *ink*) rarely produce perseverative errors as primes—these are also the objects that are rarely given as the response in semantic errors. In contrast, a few other objects (e.g. *paper*, *board*, *saw*) often lead to perseverations. Interestingly, while *saw* is also a common response in semantic errors, many of the other frequent perseverations are not particularly frequent semantic error responses. It seems some objects benefit from perseverative influences somewhat independently of their strength as semantic responses. However, the distributions across objects of purely horizontal semantic influences (S+U) and purely vertical

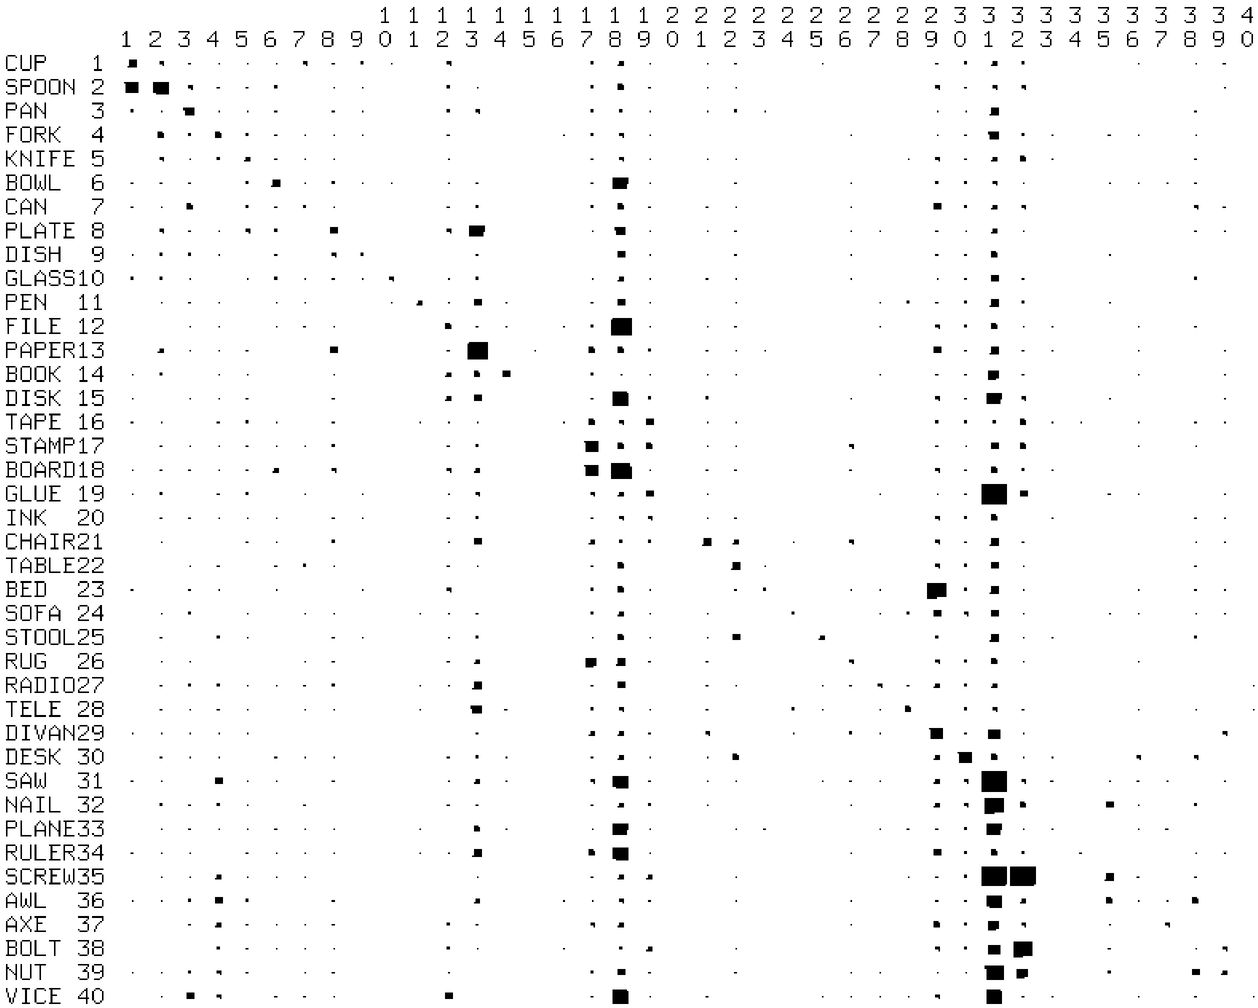

Figure 8.17: The frequency with which each object as prime (listed on the left) results in each object as response (numbered along the top).

perseverative influences (O+P) are significantly correlated (0.51, paired $t(38) = 3.62$, $p < .001$). Thus damage changes the relative strengths of the attractors for objects, and these differences are reflected in the proportion of errors in which each object is the response, somewhat independently of the nature of the influence that leads to the error.

Figure 8.17 plots how often each object as the prime leads to each object as the response. This is not actually a “confusion matrix” because, in generating the response, the network is presented with the stimulus rather than the prime. However, it provides an indication of the strength and distribution of the perseverative influence of each object. The diagonal values reflect the relative frequency across objects of perseverations in which the prime and response are identical—these are the same values as those plotted in Figure 8.16. Because all combinations of primes and stimuli are presented equally often, the overall strength of the vertical column for an object indicates the frequency of that object as a response independent of perseverative influences. The influence of

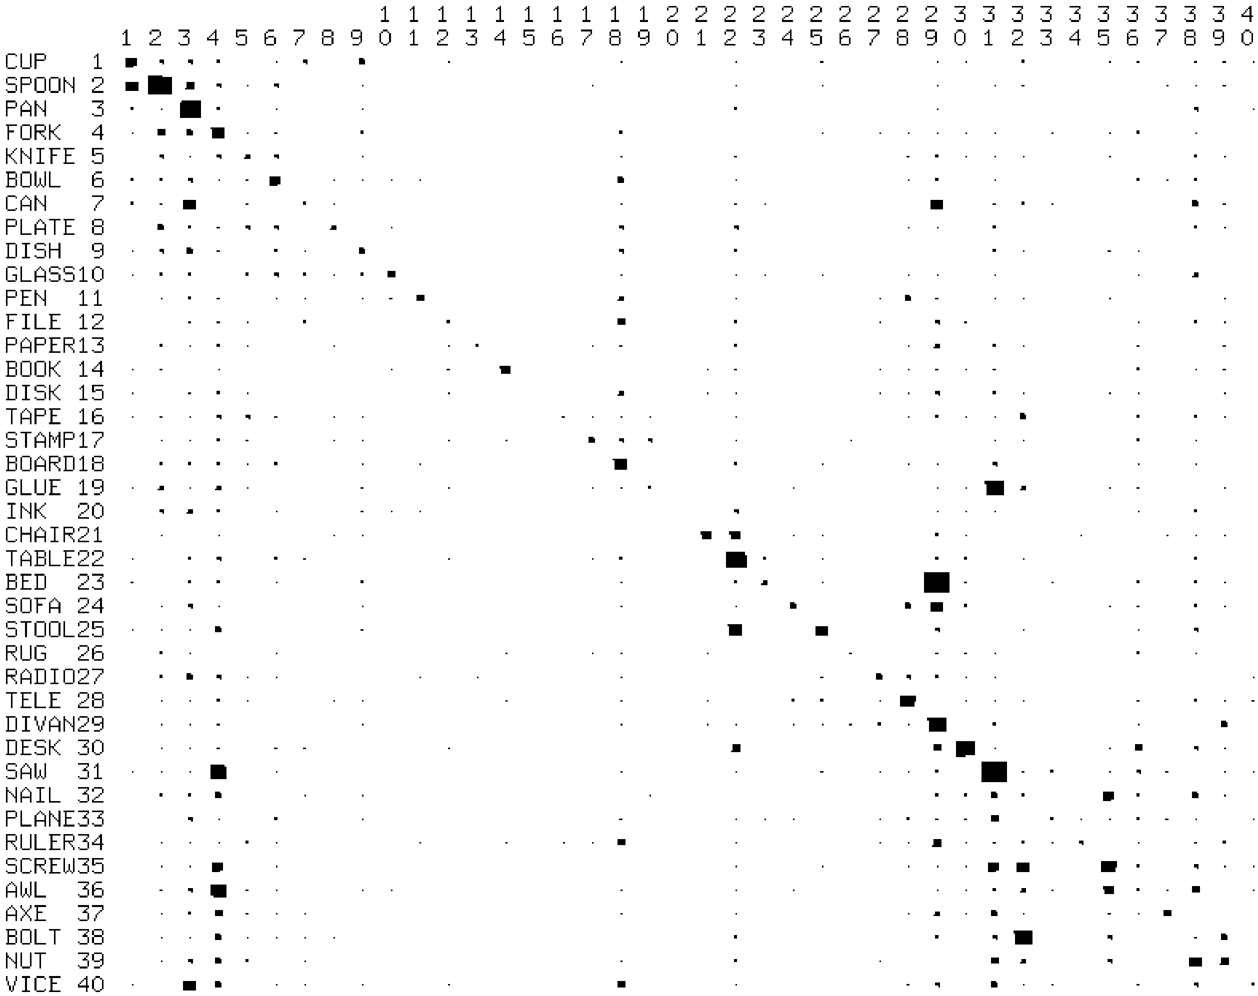

Figure 8.18: The frequency with which each object as prime (listed on the left) results in each object as response (numbered along the top), for errors produced by lesions other than to `V`⇒`I` connections.

the prime is reflected in the *distribution* of values along each vertical column. Thus the tendency to produce semantic perseverations is indicated by the larger values within an object's column for other objects in the same category.

Earlier we noted that `V`⇒`I` lesions produce the majority of errors, and far more semantic errors without perseverative influence (S+U) than other lesion locations. The remaining errors showed a stronger tendency towards perseverative influences. This can be seen in the distribution of the prime and response for these errors, shown in Figure 8.18. Comparing this "confusion matrix" with the previous one, the reduced relative influence of horizontal semantic similarity is reflected in the reduction of the large horizontal stripes for most objects. Interestingly, the few that still show this effect (e.g. *fork*, *table*, *divan*, *bolt*) are different than the objects showing the strongest horizontal semantic influences after `V`⇒`I` lesions (*board* and *saw*). Lesions near or within semantics bias response towards different objects than lesions nearer to the input. The general "strength" of the perseverative influence of an object as prime is reflected in the size of the squares along the horizontal row for the object. No object stands out as particularly stronger than the others, although in general, objects producing frequent response perseverations (on-diagonal value) also tend to have stronger perseverative influence toward semantically related responses.

8.4 Relation to deep dyslexia simulations

The current simulations demonstrate that the computational principles that account for the reading behavior of deep dyslexics can be extended to reproduce many of the characteristics of visual object naming in optic aphasia. The main similarity in the patterns of performance of the two classes of patients—mixtures of visual and semantic influences in errors—is naturally accounted for by the effects of damage in networks that build attractors. However, there are significant differences between deep dyslexics and optic aphasics—how well have the current simulations accounted for these differences? Some comparisons are made difficult by the differences in the nature of words vs objects as stimuli. Thus the advantage of deep dyslexics in reading concrete vs. abstract words would be hard to investigate in optic aphasics because object are concrete almost by definition. However, two main differences in the pattern of performance of deep dyslexics and optic aphasics that are appropriate to compare relate to the relative frequency of purely visual errors and perseverations.

Errors that are visually but not semantically related can constitute a fairly high proportion of all errors in some deep dyslexics (e.g. 51% for P.S., Shallice & Coughlan, 1980) and yet are quite rare in optic aphasics (Coslett & Saffran, 1989b; Gil et al., 1985; Lhermitte & Beauvois, 1973). One possible contribution to this difference is that the definition of visual similarity may be more stringent for objects than for words. However, with presumably the same criteria for visual similarity, associative agnosics make a high proportion of visual errors in naming objects (e.g. 46% for F.Z.; Levine, 1978; also see Larrabee et al., 1985). Thus the rarity of visual errors by optic aphasics cannot be completely explained by a criteria difference. Our explanation is that the relative difference in visual errors in reading vs. object naming is due to the different amount of structure in the two tasks. As demonstrated above, the greater structure between the visual and semantic representations of objects diminishes the influence of purely visual similarity on the layout of attractor basins within semantics. However, is it legitimate to assume there is greater structure in mapping objects vs. words to semantics? We have assumed that the mapping from orthography to semantics is completely unstructured, which is only approximately true (see Lakoff, 1987). However, to the extent that visual characteristics are *included* in the semantic representation of objects, there will be more structure between the visual form of objects and semantics than for words. This would seems to particularly apply to biological semantic categories. In addition, members of functionally defined categories (e.g. tools) often share visual characteristics because similar shapes are appropriate for similar actions (e.g. elongated shapes for pounding, horizontal shapes for sitting/lying, Gibson, 1979). Thus the assumption that object naming is more structured than word naming seems justified—the network demonstrates that this difference can account for the relative rarity of purely visual errors in optic aphasia.

The second main difference between reading in deep dyslexia and object naming in optic aphasia relates to perseveration. Although perseveration is common after many types of language-related impairments (Albert & Sandson, 1986), is it not particularly prevalent in the reading errors of deep dyslexics. In contrast, a relatively high proportion of the naming errors of optic aphasics are related to previously presented objects (e.g. 28% for J.F., Lhermitte & Beauvois, 1973). We introduced short-term correlational weights to provide a means by which object naming could be influenced by the responses given to previous objects. While some amount of independent justification can be given for such weights, we were motivated to include them in the current simulations directly by the observation of perseverative effects in patients, and so they must be viewed as somewhat *ad hoc*. However, the fact that they lead to interesting interactions with other aspects of the operation of the network, such as semantic influences in errors, suggests that their introduction contributes in a significant way to understanding the nature of perseverative influences in optic aphasia.

However, a question remains: if there is independent motivation for including short-term weights in mapping between visual and semantic representations in object naming, why did we not include them in networks for mapping orthography to semantics? The simple answer is that the behavior of the patients with impairments in this mapping is better explained without them. However, this answer is unsatisfying without an independent explanation for why a temporary bias towards previous patterns of activity is *computationally* appropriate in object recognition but not reading. One possible explanation is that, unlike in object recognition, in reading for meaning there is great pressure to recognize successive words as quickly as possible. As long as the meaning of each individual word is unrelated to the next, any bias towards the semantics of previous words would induce a kind of "sluggishness" that would impede the network in deriving the correct semantics of the current word. In general, short-term weights are not appropriate in a network for a task in which the speed of separate successive interpretations is critical. In contrast, successively recognized objects—those found together—may tend more to be related than successive words, so a temporary bias in object recognition would be beneficial. In addition, the use of short-term interactions for feature binding and segmentation is more critical for natural objects than for words. However, it should be kept in mind that this argument for the appropriateness of fast weights in object recognition but not reading is rather speculative.

8.5 Recognition in optic aphasia

A major issue remains to be addressed, regarding the relationship between the current simulations and the *preserved* abilities of optic aphasics. The impaired visual object naming of optic aphasics is perplexing because their visual *recognition* of objects, as indicated by gesturing or categorization tasks, as well as their naming of objects presented in other modalities, appears relatively intact. As pointed out at the beginning of this chapter, no single lesion within the conventional framework depicted in Figure 8.1 seems able to account for such a pattern of impaired and preserved abilities.

The current simulations reproduce the error pattern of optic aphasics in visual object naming by introducing an impairment in deriving semantics from visual input. In this sense it follows Riddoch & Humphreys' (1987) claim that optic aphasia is more appropriately considered a "semantic access agnosia." However, the current research simulates neither intact visual recognition nor intact non-visual naming in the context of impaired visual naming. In what sense then is it a simulation of optic aphasia? The honest answer is that it isn't one—it is only a simulation of the error pattern of optic aphasics. However, such a simulation is only interesting as an explanation of patient behavior if it can plausibly be extended to encorporate the remaining characteristics of the syndrome.

In attempting to reconciling the current simulations with the preserved abilities of optic aphasics, we must re-emphasize that the visual recognition capabilities of optic aphasics may not be as intact as generally thought. This possibility was discussed in detail earlier in the chapter when describing Riddoch & Humphreys' theoretical account of the syndrome. There seems to be considerable evidence that visual recognition in optic aphasics is at least not *normal*. However, the question is whether an impairment sufficiently severe to produce the poor level of performance in naming can also be sufficiently mild to support whatever level of intact gesturing and categorization is observed in these patients. In the context of our simulations, how might the impaired semantics that produce poor naming still support relatively intact gesturing and categorization?

Regarding categorization, Hinton & Shallice (1991) demonstrated that, even when damage to their network impaired explicit naming performance quite significantly, often the network could still perform both within- and between-category forced choice discriminations quite well (see Section 2.6.4). They modeled these discrimination tasks by applying less stringent criteria to semantics, so this explanation amounts to a claim that the categorization tasks at which optic aphasics succeed require less precise semantics than naming. Riddoch & Humphreys provide some evidence for this by showing that their patient J.B. was significantly impaired at a categorization task that required distinctions *within* a category (as naming must).

However, the same argument is unlikely to account for the relatively preserved gesturing of optic aphasics. Lhermitte & Beauvois' (1973) patient J.F. *never* made an incorrect gesture to 100 visually presented objects, of which 31 were misnamed. If gesturing were based entirely on the same impaired semantics that underlies poor naming, occasional gesturing errors would be predicted. Riddoch & Humphreys propose that correct gesturing may often be based directly on

Figure 8.19: Allport's (1985) depiction of spoken and written word comprehension.

the (non-semantic) structural description derived by vision. However, as Farah (1990) points out, often quite different gestures are appropriate for visually similar objects (e.g. a sewing needle and a toothpick), making it unlikely that visual representations alone can support gesturing.

A possible resolution is to suggest that gesturing in optic aphasics may be based on a combination of intact visual structural descriptions and *residual* semantics. The residual semantics could narrow the range of gestures that are appropriate for the shape of the object to those that are consistent with the general semantics of the object, preventing inappropriate gestures.[6] Support for this type of explanation comes from recent work by Farah & McClelland (1991) in modeling category-specific semantic memory impairments. They argue that the apparent deficit in some patients in naming living vs. non-living things (Warrington & Shallice, 1984) is better interpreted as a modality-specific deficit in visual vs. functional semantics. Living things are selectively impaired after visual damage because a greater proportion of their semantics involves visual information, whereas the semantics of non-living things emphasizes functional information. Allport (1985) made a related proposal in the context of accessing semantics from orthographic and phonological representations (see Figure 8.19). Perhaps gesturing in optic aphasia can be based on the intact generation of functional/motoric portions of an object's semantics, even though the generation of other portions of semantics that normally support naming is impaired. This proposal amounts to a hybrid of Riddoch & Humphreys' account with Beauvois' (1982) distinction of "visual" vs. "verbal" semantics, except that, following Farah & McClelland, the division would be based on modality rather than type of information. Although a simulation based on this proposal, involving naming and gesturing to visual and non-visual stimuli, has yet to be developed, it seems plausible and would constitute a complete simulation of optic aphasia that is consistent with the current simulations.

[6]This would be analogous to the claim that residual operation of the phonological route in phonological alexics (see Section 2.1.2) "edits out" any potential semantic errors that might arise from the impaired operation of the semantic route (Newcombe & Marshall, 1980b).

8.6 Summary

The simulations presented in this chapter demonstrate that the computational principles that have proven useful in accounting for the reading behavior of deep dyslexics can be extended to reproduce the pattern of errors produced by optic aphasics in naming visual objects: (a) horizontal semantic errors, (b) horizontal visual-and-semantic errors, (c) rare horizontal visual errors, (d) vertical response perseverations, (e) vertical semantic perseverations, and (f) interactions between horizontal semantic errors and perseverations. These effects vary depending on how well the preceding object is named. Perseverative influence are minimal when the preceding object generates no responses, stronger towards a correct response, and strongest towards when the preceding object generates an incorrect response. Perseverative influences also vary with lesion location, with lesions near or within semantics producing more perseveration than lesions near the input. The proportion of semantic and perseverative errors vary considerably across objects, but in general the effects are reasonably spread throughout the object set. Although these simulations do not simulate the relative preservation of visual recognition and non-visual naming in optic aphasics, it seems plausible that the current framework can be extended into a complete simulation of the syndrome that would include these characteristics.