

anatomical organization to the levels of cortical maps, columns, and synapses (Kosslyn & Koenig 1992; Kosslyn 1994; Posner & Rothbart 1994). My specific goals, however, are fairly modest, since I will restrict my discussion to the implementation of the syntactic comprehension system as a whole as well as the implementation of the major processing subsystems outlined in section 3.2 (i.e., parsing, interpretation, syntactic STM, and attentional control). There are several reasons for this conservative approach, one of which is that research in cognitive neuroscience typically focuses on rather high levels of functional-anatomical organization; the lower levels are the province of a closely related but different field of research called computational neuroscience (Churchland & Sejnowski 1992). The most prominent reason for concentrating on the higher levels, though, is simply that virtually nothing is known about the lower levels of implementation for the syntactic comprehension system. In fact, very little evidence is available regarding even the higher levels of implementation for this system, and the evidence that does exist is controversial. It will probably take many more decades of research before we develop a basic understanding of how syntactic comprehension is accomplished by the human brain, and it is likely that the most illuminating explanations will be pitched at the lower levels of implementation and will draw heavily on neural network computer modeling. Recent advances in research on the primate visual system have shown that distinct visual functions are carried out in segregated cortical maps consisting of strange anatomical configurations such as blobs and stripes (Zeki 1993), which suggests that the brain areas supporting linguistic functions may ultimately turn out to have any number of similarly odd designs. Only time will tell.

3.3.1 Hemispheric Asymmetry

In 1861 Paul Broca claimed on the basis of clinical data that "the faculty for articulate language" resides in the left hemisphere of the brain. In the 135 years since, this view has been corroborated by a tremendous range of additional evidence, and it is now widely

accepted that in a very large number of adults, most language functions, including syntactic comprehension, are lateralized to the left hemisphere.¹ Moreover, it is known that left-hemisphere dominance for language is determined by organic rather than external factors, since it is unaffected by literacy, the number of languages a person speaks, or the type of language a person speaks—even sign language is realized in the left hemisphere (Caplan 1987; Poizner et al. 1987). There is, however, some variability across the population with respect to hemispheric asymmetry for language. This variability depends on handedness, familial handedness, and sex.

Several studies in this century have shown that about 98% of right-handed individuals have strong left-hemisphere dominance for language. Russell and Espir (1961) investigated the incidence of aphasia following left-side or right-side head wounds in right-handed war veterans, and found that 213 of 348 veterans with left-side wounds suffered aphasia, whereas only 10 of 276 veterans with right-side wounds did. Kimura (1983) conducted a similar study with right-handed stroke and tumor victims, and reported that 95 of 216 people with lesions in the left hemisphere suffered aphasia, compared to only 3 of 169 people with right-hemisphere lesions. Another source of evidence for left-hemisphere dominance for language in right-handers is the Wada test, in which sodium amytal, a short-acting barbiturate, is injected into one carotid artery, thereby temporarily paralyzing one entire hemisphere (Wada 1949). Milner et al. (1964, 1966; Milner 1974) used this technique to investigate hemispheric asymmetry for language, and demonstrated that transient aphasia occurs with left-hemisphere deactivation, but not with right-hemisphere deactivation, in 98% of right-handed subjects. The discovery that only about 2% of right-handers experience transient aphasia following right-hemisphere but not left-hemisphere deactivation is consistent with H caen and

¹ Language functions that appear to be lateralized to the right hemisphere include emotional-attitudinal prosody, connotative semantics, and the kinds of reasoning necessary for establishing discourse cohesion and for understanding anomalies, as in jokes (Chiarello 1988; Joannette et al. 1990).

Albert's (1978) finding that crossed dextral aphasia occurs in only 0.4 to 2.0% of right-handers with right-hemisphere brain damage. More recently, a number of ERP and PET studies have confirmed predominantly left-hemispheric involvement in language functions, including syntactic comprehension, for right-handed individuals (see Garrett 1994 and Kutas & Van Petten 1994 for reviews of ERP studies, and Petersen & Fiez 1993, Frackowiak 1994, and Stowe et al. 1995 for reviews of PET studies).

In contrast to right-handers, left-handers have a more variable neural implementation of language functions. In a study of 123 left-handed individuals with brain damage, Goodglass and Quadfasel (1954) showed that while 53 of 65 patients with left-hemisphere lesions suffered aphasia, 50 of 58 patients with right-hemisphere lesions did too. This general pattern of results has been obtained in several similar studies (Bryden et al. 1983; H caen et al. 1981). Perhaps the most revealing source of evidence about the neural basis of language in the left-handed population comes from the Wada test. Milner (1974) administered this test to a group of 74 left-handers and found that 51 (69%) showed left-hemisphere dominance for language, 10 (13%) showed bilateral representation, and 13 (18%) showed right-hemisphere dominance.

Another factor that influences hemispheric asymmetries for language is familial handedness. Looking first at right-handed individuals, Luria (1947) showed that those with all right-handed family members (RHF) have more severe and longer-lasting aphasia following left-hemisphere brain damage than those with some left-handed family members (LHF). However, this should not be taken to mean that LHF right-handers have right-hemisphere dominance for language, because if that were the case -we would expect to find a higher incidence of crossed dextral aphasia in LHF than in RHF right-handers, but in fact the incidence is not significantly greater for the former group (Bryden et al. 1983). Rather, it is more likely that LHF right-handers have a more diffuse, bilateral implementation of language functions than RHF right-handers. This view has received support from a series of studies conducted by Bever and his col-

leagues (Bever et al. 1989). One of the most intriguing findings to come out of this research is that RHF right-handers tend to rely more on grammatical than semantic information during on-line sentence processing, whereas LHF right-handers tend to exhibit the opposite preferences.

A third factor that appears to affect hemispheric asymmetry for language is sex. Several studies have reported data which suggest that males are more left-lateralized for language than females (McGlone 1980; Bryden et al. 1983; Kimura 1983, 1992, 1993). For instance, it has been observed that males are more likely than females to experience aphasia following left-hemisphere lesions. However, there are some problems with drawing strong conclusions from these studies. Caplan (1987) points out that many of the studies that have reported sex-by-laterality interactions for language are not statistically reliable, and that several other studies have failed to find such an interaction altogether. For instance, aphasia does not occur more frequently in females with right-hemisphere lesions than in males with right-hemisphere lesions (Damasio et al. 1989; Kimura 1992). On the other hand, a few recent PET studies provide solid support for the view that while right-handed males have left-hemisphere dominance for language, right-handed females have more bilateral representation (Shaywitz et al. 1995; Lockwood et al. 1996). Unfortunately, none of these PET studies dealt with syntactic comprehension.

Attempts to relate left-hemisphere dominance for language to neurobiological differences between the cerebral hemispheres have generally focused on a cortical region known as the planum temporale, which, on the left side, is an extension of Wernicke's area. In a classic study, Geschwind and Levitsky (1968) examined the brains of 100 right-handers and found that the length of the lateral edge of the planum temporale was longer in the left hemisphere than in the right in 65% of the cases, equal in the two hemispheres in 24% of the cases, and longer in the right hemisphere in 11% of the cases. Furthermore, the asymmetry was sometimes very dramatic, with the left planum being ten times larger than the right in some brains. Several researchers have speculated that this

anatomical asymmetry is causally related to the fact that the left hemisphere is dominant for language in the vast majority of humans (Geschwind & Levitsky 1968; Galaburda 1984; Charles et al. 1994). However, there are a number of problems with this view. The finding that the left planum is bigger than the right in only 65% of the population does not dovetail with the well-established fact that the language-dominant hemisphere is the left in over 90% of the population. In addition, anatomical asymmetries of the same kind have been observed in some nonhuman primates as well as in hominids whose language capacities are controversial (Holloway 1995; Wilkins & Wakefield 1995). Finally, traditional methods of measuring the size of the planum temporale are fraught with complications (Witelson 1976). By using a new computer program that creates three-dimensional reconstructions of the cortical surface by unfolding it and laying it out like a sheet, Loftus et al. (1993) discovered that the anatomical asymmetry originally postulated by Geschwind and Levinsky is actually an illusion. In particular, they found that as many brains have a larger planum in the left hemisphere as have a larger planum in the right hemisphere.

To summarize, in the vast majority of the population, most language functions are carried out in the left hemisphere. Some factors that influence the participation of the right hemisphere in language include left-handedness, familial left-handedness, and possibly being female. Some researchers have associated left-hemisphere dominance for language to the left planum temporale, but recent studies have rendered this association problematic.

3.3.2 Intrahemispheric Localization

Numerous studies have demonstrated that within the left hemisphere, language functions involved in both comprehension and production are carried out predominantly in the perisylvian cortex (see Caplan 1987, 1994 for reviews). As shown in Figure 22, this large region of the brain includes the posterior half of the third frontal convolution

Figure 22: Left perisylvian language areas (from Fuster 1995).

(BA 44, 45, often referred to as Broca's area), the pre- and post-central gyri (BA 6, 43), the insular cortex buried within the sylvian fissure, the supramarginal and angular gyri of the parietal lobe (BA 39, 40), the posterior superior temporal gyrus (BA 22, often referred to as Wernicke's area), and the anterior superior temporal gyrus and temporal pole (BA 22, 38). Using the computer program described in the previous section, Gazzaniga (1989) created a "brainprint" illustrating the spatial contiguity of the various cortical areas making up the perisylvian language region; this is shown in Figure 23.

Figure 23: View of the "unfolded" cortical surface of the left hemisphere (from Gazzaniga 1989).

Other regions within the left hemisphere that have recently been attributed language functions include areas of the prefrontal cortex (BA 47, 10, 46) (Greenfield 1991; Deacon

1992; Grossman et al. 1992; Posner & Raichle 1994; Naidoo et al. 1995; Jaeger et al., in press), the anterior cingulate cortex (BA 24, 32; Grossman et al. 1992; Posner & Raichle 1994), the middle temporal gyrus (BA 21; Damasio 1992; Mazoyer et al. 1993; Damasio et al. 1996; Jaeger et al., in press), the basal ganglia and thalamus (Crosson 1990; Damasio 1992; Grossman et al. 1992), and the hippocampus (Squire et al. 1992; Jaeger et al., in press); also, portions of the right and left cerebellum have been shown to contribute to language processing (Fiez et al. 1992; Leiner et al. 1993; Posner & Raichle 1994; Silveri et al. 1994; Jaeger et al., in press; Miller, in press). In the sections that follow, I will concentrate on evidence regarding the neural implementation of the components of the syntactic comprehension system.

3.3.2.1 Left Perisylvian Areas that have been Associated with Syntactic Comprehension

Since the mid-1970s, syntactic comprehension has been associated most strongly with the cortex in the vicinity of Broca's area. Prior to this time, individuals with damage to this region of the brain—i.e., Broca's aphasics—were thought to have just a disorder of language production, since their most prominent symptoms (often but not always occurring together) are apraxia of speech and agrammatism, the first of which involves an impairment of articulation and the second of which involves the omission of function words, inflections, and sometimes verbs, and the avoidance of complex grammatical constructions (Kean 1985; Menn & Obler 1990; Caplan 1991). However, in 1976 Caramazza and Zurif (1976) showed that although Broca's aphasics have good single-word comprehension and are able to understand complex sentences when semantic or pragmatic cues are available, their comprehension of many types of complex sentences drops to chance when they are forced to rely solely on grammatical information. Since then, a great deal of effort has been devoted to exploring the nature of the syntactic comprehension deficits of Broca's aphasics, or agrammatic aphasics, as they are sometimes called. For present purposes, the important point is simply that these findings

have led many people to infer that the cortex in and around Broca's area is necessary for syntactic comprehension (Mesulam 1990; Damasio 1992).

While this view has become quite popular, other studies indicate that it is not completely adequate. First of all, a number of cases have been reported of Broca's aphasics who have impaired language production but normal syntactic comprehension (Miceli et al. 1983; Nespoulos et al. 1984; Kolk & van Grunsven 1985). In addition, Caplan et al. (1985; see also Caplan & Hildebrandt 1988) conducted an extensive investigation of the syntactic comprehension abilities of three groups of 58, 37, and 49 aphasic patients who were unselected, i.e., who belonged to a variety of diagnostic categories and had lesions in a variety of brain regions. The researchers tested the patients' performance on a wide range of grammatical constructions, and then carried out clustering analyses to identify subgroups of patients that differed in the overall severity of their syntactic comprehension deficits. Next, the researchers attempted to correlate the subgroups with the patients' lesion sites, which were classified as either purely frontal, purely parietal, or purely temporal. They found that patients with lesions confined to any of the three lobes in the perisylvian cortex were equally likely to fall into any of the subgroups identified by the clustering procedure. In other words, the results showed that damage restricted to any one of these lobes can cause a severe disruption of syntactic comprehension, no disruption whatsoever, or any degree of disruption. The inference that Caplan et al. draw from these findings is that the neural implementation of the major components of the syntactic comprehension system is not universal but rather appears to vary across individuals. Caplan (1994) suggests that the specific localization of these components is largely determined by genetic factors. This view is plausible simply because there must be genetic variation across individuals with regard to linguistic capacity; otherwise language could never have evolved. Further support for Caplan's view comes from a recent study demonstrating that while individual brains typically exhibit considerable variation in the gyral/sulcal pattern and in the relative size of different

cortical areas, the brains of monozygotic twins show very little variation along these parameters, especially in the left hemisphere, which is dominant for language (Tramo et al. 1995; see also Whitaker & Selnes 1976).

Although Caplan et al.'s study is clearly very important, it has the limitation of not providing very narrow lesion localization data for the patients. Lesions were classified as either purely frontal, purely parietal, or purely temporal, but each of these lobes contains a very large amount of cortex. Thus, it is possible that two patients who are both treated as having frontal lesions actually have lesions affecting nonoverlapping areas of the frontal lobe. In such an event, if one patient suffers syntactic comprehension deficits whereas the other does not, the most appropriate inference would be that the former patient's lesion affected a subarea of the frontal lobe which is necessary for syntactic comprehension whereas the latter patient's lesion did not. Moreover, it could be that the subarea that is damaged in the former patient is one that typically plays an important role in syntactic comprehension. The upshot is that we should not be too quick to accept Caplan et al.'s view that the localization of the major components of the syntactic comprehension system is variable across the population.

Several other studies provide data that give some support to Caplan et al.'s view. Vignolo et al. (1986) examined the CT scans of 37 global aphasics with severe production and comprehension deficits, and found that 22 had lesions including both Broca's and Wernicke's areas, eight had lesions extending from Broca's area to the anterior part of the temporal lobe, three had lesions in the parietal-occipital region, and four had deep lesions affecting the insular cortex. Following an investigation of the syntactic comprehension abilities of nine groups of aphasic patients, Naeser et al. (1987) concluded that although damage to the posterior two thirds of the superior temporal gyrus seems to cause the most severe deficits, damage to the surrounding frontal, parietal, and temporal areas can also produce deficits in some individuals. Finally, Ojemann (1983; Ojemann et al. 1989) reported that during intra-operative electrocortical stimulation studies, syn-tactic

comprehension can be disrupted by stimulation in fairly restricted brain regions, but these critical regions vary considerably from patient to patient.

Although most of the evidence discussed so far goes against the idea that there is a systematic, universal implementation of the major components of the syntactic comprehension system, two important new studies claim to have found reliable deficit-lesion correlations. Remarkably enough, however, these studies isolate different brain regions as being crucial for syntactic comprehension. In the first study, Kempler et al. (1991) examined the relationship between syntactic comprehension ability and two measures of brain damage—one structural (CT) and the other metabolic (PET)—in 43 aphasic patients. The major discovery was that across the entire group there was a strong correlation between syntactic comprehension impairment and glucose hypometabolism in the temporoparietal cortex. Furthermore, the severity of syntactic comprehension impairment was positively related to the degree of hypometabolism in this brain region. These findings are quite striking, but they must be interpreted with caution. A reduced metabolic rate in a specific brain area does not necessarily reflect less information processing in that area (Kosslyn 1994; Sergent 1994). And just because this group of patients has a common area of brain dysfunction does not mean that this particular area is absolutely necessary for syntactic comprehension in all of the patients. It may be the case that some of the patients' deficits are due to structural or metabolic lesions that they *don't* have in common. Such a situation is not only possible in principle, but is empirically supported by the studies reviewed earlier that suggest variable localization of the components of the syntactic comprehension system.

The second study, which has not been published yet, was conducted by Dronkers et al. (submitted). These researchers assessed the syntactic comprehension abilities of 26 unselected aphasic patients, and subsequently carried out a cluster analysis to form three subgroups. Nine patients had consistently poor performance, 12 had consistently good performance, and five had mixed performance. Then they determined the common areas

of brain injury for the different subgroups by applying a computer program that superimposes multiple CT scans. The results showed that while all but one of the patients in the impaired subgroup had lesions that included Broca's area, two of the patients in the normal-like subgroup did too. Moreover, three patients from each subgroup had lesions affecting Wernicke's area. When the researchers focused on a different cortical area, however—specifically, the anterior third of the superior temporal gyrus—they discovered that it was damaged in all of the patients in the impaired subgroup, but was spared in all of the patients making up the normal-like subgroup. The authors also review a variety of other sources of evidence that are consistent with the notion that this brain region plays a critical role in syntactic comprehension (from PET: Mazoyer et al. 1993, Stromswold et al. 1996; from patients with temporal lobe epilepsy and anterior temporal lobectomies: Milner 1958; Rochetta 1986; Frisk & Milner 1990; Shih & Peng 1992; from ERPs: Kluender & Kutas 1993; Hagoort & Kutas 1995). Dronkers et al. conclude by stating that although it is likely that many different areas within the left hemisphere contribute to syntactic comprehension, the anterior sector of superior temporal cortex may be an especially important area.

Like Kempler et al.'s study, this localization study is quite impressive, but it is not problem-free. In fact, it has the same basic weakness as Kempler et al.'s study. One cannot safely infer from superimposing lesion data from several patients that the overlapping sites are (part of) what cause the deficits, since it is always possible, on strictly logical grounds, that the deficits are due to nonoverlapping areas of damage in at least some of the cases. For instance, if three patients have lesions that include the anterior superior temporal gyrus, but one of the lesions also includes Broca's area, another also includes Wernicke's area, and the last also includes the middle temporal gyrus, it is logically possible that these patients' syntactic comprehension deficits are due to the nonoverlapping areas of damage. Thus, I believe that although the results of Dronkers et al.'s study are clearly very important, they are not entirely sufficient to warrant the

conclusion that the anterior superior temporal gyrus is always necessary for syntactic comprehension.

Given that it is so hard to specify precisely a region within the left perisylvian cortex that is reliably associated with syntactic comprehension in general, it is no doubt even riskier to make any strong claims about the neural substrates of the various components of the syntactic comprehension system. Nonetheless, there is some evidence suggesting that each of the major components *tends* to be implemented in a particular region of the left hemisphere. Most of this evidence comes from studies using the PET and ERP techniques, but I should acknowledge at the outset that these studies are by no means easy to understand. Different studies sometimes produce conflicting results, and the results of virtually all of these studies contradict the results of some of the clinical studies described earlier: while the PET and ERP data suggest a fairly regular implementation of processing subsystems, some of the clinical data suggest variable implementation. Problems like these are vexing, but they do not require that we abandon any hope of ever making sense of the data. For instance, different findings in different PET studies may be due to different methodologies, and PET and ERP data may mask variability across the subjects. Also, the subjects in PET and ERP studies are usually selected very carefully so that they are maximally similar with respect to sex, age, handedness, education, socioeconomic status, race, and so forth. Thus, it is possible that the subjects in PET and ERP studies have fairly similar neural substrates for syntactic comprehension, but that the patients in clinical studies have more variable implementation, since they are not forced to meet such stringent inclusion criteria. Having quickly mentioned these difficult issues, I now proceed, albeit with some trepidation, to discuss the neuro-biology of the major components involved in syntactic comprehension. I will focus first on parsing and interpretation, and will then shift to the two processing resources of syntactic STM and attentional control.

3.3.2.2 Parsing

Grossman et al. (1992a) conducted a PET study in which one of the subjects' tasks was to determine whether or not each of a sequence of sentences contained an adjective. The results showed significant levels of activation in the left middle and inferior frontal lobe—a region encompassing Broca's area—but not in the left temporal or parietal lobes.² Since this receptive language processing task involves monitoring the categories of words, one might think that the brain areas that are activated are related to parsing. More specifically, one might think that the areas of activation reflect the process of assembling syntactic constituent structures and "scanning" them for nodes with the adjective category. This is not the only possible task analysis, however. An alternative is that the subjects were simply monitoring for *lexical* items of the adjective category, a task which does not in itself require parsing at all. Thus, although the brain areas that Grossman et al. found to be activated in this task may be related to parsing, they may instead be related to making detection judgements about the presence or absence of particular kinds of lexical items.

In another PET study, Mazoyer et al. (1993) presented 16 subjects with a variety of different types of linguistic stimuli. In two of the conditions, the stimuli were designed "to disrupt semantic integration while preserving syntax and prosody" (p. 468). Specifically, the stimuli in these conditions consisted of, first, sentences in which the content words were replaced with pseudowords and, second, sentences in which the content words were replaced with semantically unrelated words of the same grammatical category, frequency, length, and imageability. The researchers found that while neither of these conditions elicited activation in Broca's area, both of them elicited activation in the entire extent of the left and right superior temporal gyri. The posterior sectors of the left

² There was also significant activation in several brain areas that are involved in the visual processing of words as well as in several areas that contribute to attention; I will discuss these latter areas later in the text.

and right superior temporal gyri were also activated in a separate condition in which the subjects simply listened to a list of words. The authors suggest that in the first two conditions the right anterior temporal activation may underlie the perception of prosody (see my footnote 11, p. 110), whereas the left anterior temporal activation may be related to parsing or at least attempted parsing, given the distorted nature of the stimuli (recall that the stimuli involved sentences with pseudowords and semantic anomalies). On the other hand, the researchers point out that because the stimuli were so unusual, it may be that the left anterior temporal activation merely reflects verbal memory search. Also, in another condition in which the subjects listened to a story in French (their native language), activation was found not only in the left and right superior temporal gyri, but also in the left middle temporal gyrus and Broca's area. Thus, it is possible that either of these additional areas of activation is related to parsing.

In summary, although the two PET studies by Grossman et al. and Mazoyer et al. do not produce convergent results and conflict with some of the clinical data presented earlier, they do suggest that there is a tendency for parsing operations to be implemented in the anterior portion of the perisylvian cortex.

This view receives further support from several recent ERP studies of sentence processing. In an RSVP (i.e., rapid serial visual presentation) experiment, Neville et al. (1991) showed that phrase structure violations provoke a left anterior negativity that peaks between 200 and 300 msec post-stimulus. Using an auditory mode of presentation, Friederici et al. (1993) observed a similar early left anterior negativity in response to phrase structure violations. Additional replications based on violations of word category, argument structure, and inflectional agreement have been reported by Rösler et al. (1993), Münte et al. (1993), Münte and Heinz (1994), and Friederici et al. (1995). The ERP effects produced by syntactic violations contrast in latency and neurotopography with other ERP effects produced by semantic violations. Most notably, the former effects contrast with the classic N400 effect, which results from semantic incongruity and has a

bilateral temporoparietal distribution (although research with split-brain patients suggests that it is generated only by the left hemisphere—Kutas et al. 1988; Kutas & Van Petten 1994). One nontrivial problem with all of these ERP studies is that they are based on linguistic violations, and so there is no guarantee that we "know exactly what is 'expected' and therefore what is 'violated'" (Kutas & Kluender 1994:185). Another problem is that the ERP method does not allow completely reliable localization, since electrical currents not only travel across the brain and across the scalp, but are also distorted when they pass through the skull. Nonetheless, the studies mentioned above are still valuable insofar as they are consistent with the previously described PET data suggesting that parsing operations tend to be implemented in the anterior portion of the left perisylvian cortex.

3.3.2.3 Interpretation

I turn now to the neurobiology of the interpretive component of the syntactic comprehension system. There are two main issues here: first, the implementation of the LSs and macroroles of verbs; and second, the implementation of the linking operations that establish correspondences between NPs, macroroles, arguments in the LSs of verbs, and concepts expressed by nouns. I will address these issues in order.

While a fair amount of work has been done during the past few years on how the meanings of concrete nouns are represented in the brain (for reviews see Caramazza et al. 1994, Damasio & Damasio 1994, and Gainotti et al. 1995), much less attention has been paid to the neural underpinnings of verb meanings, and the evidence that has been gathered does not fit together very well. Nonetheless, I will summarize the major findings and attempt to identify the most plausible possibilities regarding localization. It is known that the "dorsal" processing stream leading from the occipital lobe through the parietal lobe to the premotor and supplementary motor cortices represents the visual motion patterns of entities, the spatial and temporal relations that obtain among entities,

and schemas for executing different bodily actions (Kosslyn & Koenig 1992; Kosslyn 1994). Since this is the kind of information that verbs typically encode, one might expect that the LSs and macroroles of verbs would be represented in cell assemblies distributed throughout these regions of the brain (Damasio & Tranel 1993; Damasio & Damasio 1994). So far, however, I have not encountered any direct evidence for this hypothesis. Instead, it appears that two distinct anatomical areas in the left hemisphere, one in the frontal lobe and the other in the temporal lobe, play especially important roles in accessing and representing verb meanings.

Petersen et al. (1988, 1989; see also Posner & Raichle 1994) conducted a PET study in which one of the conditions required subjects to view a series of nouns and generate, for each one, a semantically related verb (e.g. *cake* --> *eat*). They found significant levels of activation in the ventrolateral prefrontal cortex, especially in area 47, as well as in Wernicke's area.³ Subsequent studies led to two further discoveries about the interaction of these brain areas in the verb generation task. First, when subjects were allowed to practice the verb generation task before being scanned, all of the areas of activation seen in the first study disappeared and the only area that showed any significant activation was a completely different one—namely, the insular cortex (Raichle et al. 1994; see also Posner & Raichle 1994). The researchers concluded that the network of areas observed in the first study subserves nonautomatic processing for the task, whereas the single area observed in the second study subserves overlearned, automatic processing. Second, when naive subjects (i.e., subjects not allowed to practice) were imaged with both PET and ERP techniques simultaneously, it was shown that the ventrolateral prefrontal cortex is activated very quickly, about 200 msec after the noun is presented,

³ In addition, activation was found in the anterior cingulate cortex and in the right cerebellum. I will consider the cingulate activation later on when I discuss the neurobiology of attentional control. It is worth noting here that the overall pattern of results has been replicated in similar studies using PET (Wise et al. 1991), fMRI (McCarthy et al. 1993), and electrical stimulation (Ojemann et al. 1993).

and Wernicke's area is activated much later, about 700 msec post-stimulus (Snyder et al. 1995). Thus, it is possible that many semantic tasks could be carried out in the anterior region well before the posterior region becomes activated. Still, the entire set of results suggests that both regions of the brain contribute to the processing of verb meanings.

Support for the importance of the left frontal lobe in processing verb meanings comes from the clinical literature. Numerous studies have reported that lesions in this brain region typically disrupt the production and comprehension of verbs more than nouns (Miceli et al. 1984; McCarthy & Warrington 1985; Bates et al. 1991; Daniele et al. 1994; Kellogg 1995; Breedin & Martin 1996). None of these studies, however, distinguishes between two possible functions of this frontal region: (1) accessing the LSs and macroroles of verbs, and (2) actually representing these properties of verbs. Kosslyn and Koenig (1992) suggest that the frontal activation observed in the PET studies of verb generation reflects the former kind of operation rather than the latter. They base this speculation on the fact that other research—e.g., in the domain of visual object recognition—has demonstrated that the left ventrolateral prefrontal cortex plays a role in the process of "looking up" specific information stored in memory (see also Kosslyn 1994). On their view, the actual representations of verb meanings are most likely implemented in the posterior superior temporal cortex (i.e., Wernicke's area) and perhaps also in the temporal-parietal-occipital junction. Kosslyn and Koenig's ideas about the lateral prefrontal cortex are consistent with Petrides's (1995) theory that one of the general executive functions of the ventrolateral prefrontal cortex is to actively retrieve information stored in short-term or long-term memory (see 2.1.1.4, pp. 26-7).

This approach accords nicely with the finding mentioned earlier that, in the verb generation task, the anterior region is activated much sooner than the posterior region. In addition, there is some clinical and PET data that dovetails with this approach. First, Damasio & Tranel (1993) discovered that four patients with lesions in the left frontal lobe performed normally on tests requiring the manipulation of noun concepts, but exhibited

an impaired ability to retrieve verb concepts in both production and comprehension modalities. The patients' knowledge of verb concepts was intact, however, as shown by the finding that some of the items that were failed on one experimental epoch were passed on the next, and vice versa. This suggests that a mechanism for just accessing the LSs and macroroles of verbs is implemented in the frontal region. Second, in an on-line sentence processing study, Shapiro et al. (1993) discovered that while normal subjects as well as Broca's aphasics momentarily activate all of the possible argument structures of a verb when they encounter it, Wernicke's aphasics do not. This suggests that the LSs and macroroles of verbs may actually be stored in the vicinity of the posterior superior temporal cortex. There is, however, a problem with Shapiro et al.'s results—in particular, if a verb-retrieval mechanism is implemented in the ventro-lateral prefrontal cortex, wouldn't we expect it to be impaired in Broca's aphasics, and hence wouldn't we expect these patients to be just as deficient as Wernicke's aphasics at activating all of the possible argument structures of verbs? Although this is a difficult question, one possible answer is that the retrieval mechanism in the ventrolateral pre-frontal cortex may not be used every time a person accesses a verb meaning; rather, it may only be necessary in tasks that require the subject to carry out a guided search through semantic memory. While the single-word tasks administered in the PET studies and in the clinical studies cited previously (including the one by Damasio and Tranel 1993) seem to tap this kind of operation, the on-line sentence processing study conducted by Shapiro et al. may not have.

A final piece of evidence supporting the view that the LSs and macroroles of verbs are implemented in the posterior superior temporal cortex comes from the PET study by Grossman et al. (1992a) that was mentioned in the discussion of parsing. In one of the conditions in this study, the subjects viewed various types of sentences—actives, passives, and relatives—and were required to determine, for each one, whether or not the actor was female. Since this task involves monitoring the syntactically determined semantic

relationships between noun concepts and verb concepts, it seems reasonable to assume that the brain areas that are activated contribute to representing the LSs and macroroles of verbs. The results showed that the most significant amount of activation elicited by this task was in the left posterior superior temporal cortex.

The evidence that I have presented concerning the localization of the LSs and macroroles of verbs is far from straightforward. Although the evidence suggests that two distinct regions of the left hemisphere—the ventrolateral prefrontal cortex and the posterior superior temporal cortex—are both relevant, the precise function of each area is still a matter of controversy and remains to be elucidated (Posner & Raichle 1995). Another problem is that, as I have already mentioned, the findings summarized here are inconsistent with the studies on deficit-lesion correlations which indicate that damage to either the anterior or the posterior region does not always cause syntactic comprehension deficits. Setting these problems aside, however, the balance of evidence seems to favor the hypothesis that the semantic properties of verbs tend to be stored in the posterior superior temporal cortex, and I will assume that this view is the most plausible one in the rest of the thesis.

With respect to the second aspect of the interpretive component of the syntactic comprehension system—namely, the linking operations that establish correspondences between NPs, macroroles, the arguments of predicates, and the concepts expressed by nouns—there is no solid evidence concerning implementation, and hence I have very little to say. In general, one can conceive of the neural substrates of linking operations as consisting of chains of connections between cell assemblies representing the various types of information that must be related by chains of correspondences—specifically, lexical units and noun and verb concepts, lexical units and NPs, NPs and macroroles, macroroles and the arguments of predicates, and, ultimately, noun concepts and the arguments of predicates. It is likely that these connections are strongly influenced by other cell assemblies that keep track of various types of cues for linking, such as word order and

inflectional morphology (probably in the anterior perisylvian cortex) as well as syntactically relevant semantic properties of verbs (probably in the posterior peri-sylvian cortex).

Establishing all of the necessary correspondences between lexical, syntactic, and semantic information during the course of on-line sentence processing is one manifestation of a much larger neurocomputational problem that has come to be known as the binding problem. This problem has been studied most intensively in the domain of low-level visual perception, where it takes the following form: features of shape, color, motion, and so forth are all represented in anatomically distinct areas of the occipital lobe, so when our visual field contains several objects each of which bears these kinds of features, how does the brain manage to temporarily attribute the right features to the right objects? During the past few years, a great deal of excitement has built up over one possible solution to this problem. Basically, the idea is that transient binding is achieved by the synchronization of neural firing rates or oscillations in cell assemblies distributed throughout the relevant anatomical structures. For instance, cell assemblies representing the shape, color, and motion features of one object may all fire together in phase-locked fashion every 40 Hz, whereas cell assemblies representing the appropriate features of another object may all fire together in another phase-locked cycle that is out of synchrony with the first. This solution to the binding problem was originally presented in purely theoretical terms by von der Malsberg and Schneider (1986), but a number of researchers have recently accumulated a substantial amount of empirical evidence supporting it (for reviews see Singer 1993, 1994); moreover, a few researchers have begun to explore ways in which temporal binding can be applied to problems in domains other than vision (Hummel & Holyoak 1992; Shastri & Ajjanagadde 1993; Hummel et al. 1994; Desmedt & Tomberg 1994; Vaddia et al. 1995).

In principle, it seems that a similar binding mechanism would be useful for establishing correspondences during on-line sentence processing. To pick a highly simplified

example, consider how correspondences between NPs and macroroles might be established when processing a passive sentence like *Sally was seen by Harry*. A cell assembly representing the first NP initially starts to fire in synchrony with a cell assembly representing the actor macrorole, since there is always a high probability that the first NP of a sentence will be an actor. As the multiple cues for the passive construction are encountered, however—specifically, the auxiliary, the perfect participial form of the verb, and the preposition *by*—cell assemblies that represent these cues serve to break up the original synchronized oscillatory pattern and replace it with one that synchronizes the firing rate of the cell assembly for the first NP with the firing rate of the cell assembly for the undergoer macrorole. When the oblique NP is encountered, a cell assembly representing it quickly develops a synchronized oscillatory pattern with the cell assembly representing the actor macrorole, except this pattern is in a different phase than the other one.

Presumably, the sentence is ultimately fully interpreted by extending the first oscillatory pattern to include cell assemblies for the lexical unit *Sally*, the concept **Sally**, and the first argument of **see**, and extending the second oscillatory pattern to include cell assemblies for the lexical unit *Harry*, the concept **Harry**, and the second argument of **see**. In the end, the cell assemblies for the lexical items and the syntactic template of the passive construction can be allowed to revert to chaotic firing rates (Skarda & Freeman 1987; Freeman 1991). However, because correspondences between the noun concepts and the arguments of the predicate have been successfully established via grammatical mediation, the cell assemblies for these structures can continue to fire in distinct synchronized oscillatory patterns in semantic STM. This embodies the surviving propositional representation of the basic meaning of the sentence.

These ideas about how linking operations are carried out dynamically during on-line sentence processing are, of course, completely speculative. The functional-anatomical details of the matrix of cell assemblies that support linking can only be specified by future

research involving neural network computer modeling together with empirical investigations of the microcircuitry of the left perisylvian cortex.

3.3.2.4 *Syntactic STM*

I shift now to the neurobiology of the first major processing resource of the syntactic comprehension system—namely, syntactic STM. Stromswold et al. (1996) conducted a very tightly controlled PET study which managed to isolate the probable neural substrates of this particular processing resource. In the first condition, subjects viewed a series of sentences which all contained subject-object relative clauses (e.g., *The limerick that the boy recited appalled the priest*), and in the second condition, they viewed a series of sentences that all contained object-subject relative clauses (e.g., *The biographer omitted the story that insulted the director*). In each condition, the subjects' task was to determine whether or not the sentences were semantically anomalous (e.g., *The teenager that the miniskirt wore horrified the mother* or *The woman tipped the hairdresser that pleased the haircut*). Since this task forced the subjects to parse and interpret the sentences, the researchers could be confident that they were attending closely to the stimuli. Of special interest is the well-established fact that the two types of relative clause impose different demands on syntactic STM. As I pointed out in section 3.2.2.1 (pp. 94-5), subject-object relatives place a heavy load on syntactic STM, since the pivot NP must be retained until both the embedded predicate and the matrix predicate are encountered; by contrast, object-subject relatives do not tax this processing resource, since all of the NPs can be interpreted quite rapidly. Thus, in the experiment, subtraction of the blood flow map for the second condition from the blood flow map for the first condition should reveal the location of the brain area(s) that implement syntactic STM. What the researchers discovered when they carried out this subtraction was that a structure within Broca's area—specifically, the pars opercularis—was significantly activated. They suggest that this brain region may subserve the memory resources

necessary for comprehending sentences containing subject-object relative clauses. They acknowledge, however, that although the activation observed in this region may reflect purely syntactic STM, it could also reflect verbal STM, since the subjects could have been rehearsing the sentences to ensure correct template selection and linking (see §3.2.2.2, p. 100). Consistent with this possibility is the fact that both neuroimaging and clinical studies indicate that Broca's area contributes to the articulatory aspect of verbal STM (Vallar & Shallice 1990; Paulesu et al. 1993; Awh et al. 1995). The authors also acknowledge another, more serious problem with their results—specifically, that it is difficult to reconcile the idea that Broca's area is necessary for syntactic comprehension with the many studies which show that lesions in Broca's area do not always cause syntactic comprehension deficits.

Another interesting finding emerged from Stromswold et al.'s PET study and deserves to be mentioned here. This finding is related to the differences between, on the one hand, the blood flow maps for the first two conditions and, on the other hand, the blood flow map for a third condition in which the subjects viewed a mixture of sentences with subject-object and object-subject relatives and were required to detect nonwords (e.g., *The sculpture that the artist exhibited shocked the findle; The economist predicted the recession that chorried the man*). When the blood flow map for the third condition was subtracted from the one for the first condition (i.e., the subject-object relative condition), significant activation was found in both Wernicke's area and the anterior sector of the left superior temporal gyrus, this latter area being the one that Dronkers et al. (submitted) claim to be crucially involved in syntactic comprehension. However, when the blood flow map for the third condition was subtracted from the one for the second condition (i.e., the object-subject relative condition), no significant activation was found in either of these areas; moreover, similar findings resulted from subtracting the blood flow map for the second condition from the one for the first. It is possible to make sense of this pattern of data if we assume that both Wernicke's area and the anterior sector of

the left superior temporal cortex were activated rather strongly in the first condition, less strongly in the second condition, and still less strongly in the third condition. Such an interpretation can explain why the differences between the first and second condition, and between the second and third condition, were *not* significant, but the differences between the first and third condition *were* significant. Now the question arises as to what this pattern of data implies about the neurobiology of syntactic comprehension. One reasonable hypothesis which is consistent with some of the other studies that I have discussed in this chapter is the following: the pars opercularis of Broca's area is involved in syntactic (and verbal) STM, the anterior sector of the left superior temporal cortex is involved in parsing, and Wernicke's area is involved in interpretation.

In this context, I should point out that there is nothing unnatural about a mechanism for short-term memory being implemented in a different brain area than the one that implements the actual representations that get held on-line. In fact, Goldman-Rakic (1987, 1995) and others (e.g., Fuster 1989; Petrides 1995) have demonstrated that a major function of the prefrontal cortex is to maintain in an activated state cell assemblies located in the posterior association cortices, especially when the environmental stimuli necessary to activate the cell assemblies in a bottom-up fashion are no longer present. Such short-term memories may be achieved through reverberatory circuits that include not just specific prefrontal and posterior cortical sites, but also the basal ganglia-thalamocortical pathways that I discussed in section 2.1.1.4 (pp. 17-27) (Goldman-Rakic 1994). Recall that in that discussion I suggested that there may be a basal ganglia-thalamocortical pathway involving the ventrolateral prefrontal cortex, which includes Broca's area (pp. 26-7). If so, this would enable the basal ganglia to use information received from the left superior temporal cortex to signal to Broca's area when it is appropriate to hold certain syntactic templates on-line and when it is appropriate to allow them to decay (this hypothesis is based on the treatment of the basal ganglia as

a "cortical biasing system" in °2.2.4.1, pp. 39-42). The linguistic short-term memory functions of Broca's area may also be facilitated by dopaminergic innervation from the substantia nigra (via the mesocortical projection system); however, this is by no means certain. As I mentioned in Chapter 2 (°2.1.1.3, p. 12-14), studies with macaques show that the dopamine supply to the prefrontal cortex is strongest dorsally and weakest laterally and mesially (Williams & Goldman-Rakic 1993). Still, it is possible that evolutionary changes led to a richer dopamine supply to the lateral region of the prefrontal cortex in humans; future research is needed to determine whether this is the case.⁴

Although the ERP method does not provide spatial resolution fine enough to isolate Broca's area, a recent ERP study conducted by Kluender and Kutas (1993) does provide support for the more general idea that the anterior portion of the left perisylvian cortex plays an important role in syntactic STM. These researchers monitored the electro-physiological activity of subjects' brains from thirteen recording sites while they were processing visually presented sentences that require filler-gap integration. A wide range of constructions were used as stimuli, including yes/no-questions and WH-questions that

⁴ Neuroscientists often take macaques to be "model" or "representative" primates and assume that their findings can safely be generalized to the human brain, but there are some serious problems with this kind of reasoning (Tooby & Cosmides 1989; Preuss 1995). Over 200 different species of primates have been identified, and although all of them have certain traits in common, each has a variety of unique evolutionary specializations as well. Macaques and humans share a long history of common ancestry, and this is reflected in a relatively large number of shared features, such as eyes set together in the front of the face rather than on the sides, dextrous hands with opposable thumbs and nails instead of claws, and a complex social organization requiring the cognitive ability to keep track of a large number of ever-changing dominance and mating relationships. On the other hand, the macaque and human lineages diverged about 25 million years ago, and since then each one has adapted to different ecological and social conditions. For instance, macaques have evolved cheek pouches for hiding food from higher-ranking individuals, as well as a social organization based on a stable core of closely related females. By contrast, humans have evolved such features as bipedalism, language, concealed female ovulation, and male parental investment. Unique evolutionary specializations like these make extrapolations from macaque brains to human brains inherently suspicious, because when we consider a particular aspect of macaque neurobiology, there is no direct evidence indicating whether or not it is also part of human neurobiology. As Preuss (1995: 1229) puts it, "How can we tell whether we are studying the neural analogues of the opposable thumb and frontated orbits, rather than something akin to a cheek pouch?" Consideration of the mesocortical dopaminergic projection system is but one of an unlimited number of specific instances of this problem.

varied in terms of gap site. There were two primary points of interest in each sentence: first, immediately after the filler NP; and second, immediately after the gap. Simplified examples of the yes/no-questions, along with controls, are presented below (fillers are indicated by italics, gap sites by blank lines, and points of measurement by capital letters):

ERPs to Function Words Immediately Following Fillers

- a. Target Items:
 - i. subject: Has she forgotten *who* ____ IS . . . ?
 - ii. object: Has she forgotten *what* THEY . . . ?
- b. Control Items:
 - i. that: Has she forgotten that THEY . . . ?
 - ii. if: Has she forgotten if THEY . . . ?

ERPs to Function Words Immediately Following Gaps:

- c. Target Items:
 - i. subject: Can't you remember *who* ____ tried to scare him INTO . . . ?⁵
 - ii. object: Did he wonder *who* he could coerce ____ INTO . . . ?
- d. Control Items:
 - i. that: Can you believe that he was able to lure them INTO . . . ?
 - ii. if: Did he wonder if he could coerce her INTO . . . ?

With regard to the post-filler measurements, Kluender and Kutas observed that between 300 and 500 msec after the function words in the object condition (a-ii), a distinct effect of enhanced left anterior negativity (LAN) occurred. This effect was not seen, however, at the corresponding positions in either the subject condition (a-i) or the two control conditions (b-i,ii). The researchers suggest that the LAN effect indexes the need to suspend interpreting *what* and instead retain it in syntactic STM. In addition, they note that the absence of a LAN effect in the subject condition is not really surpris-

⁵ The actual gap in this sentence is after *who*, but ERPS were measured at *into* because this is the point that corresponds to the point immediately following the gap in the subsequent example. Hence, it may be best to think of this sentence as another control for the critical object-gap condition.

ing, since the post-filler function word is a verbal element that unambiguously indicates that the filler serves as its subject and thus does not need to be held in syntactic STM until a gap is encountered. With regard to the post-gap measurements, the very same pattern of results was obtained. A LAN effect appeared between 300 and 500 msec after the function words in the object condition (c-ii), but not at the corresponding positions in either the subject condition (c-i) or the two control conditions (d-i,ii). Kluender and Kutas interpret this occurrence of the effect as a reflection of the retrieval or reactivation of the semantic properties of the filler for purposes of being associated with the predicate. Although other grammatical constructions that require syntactic STM (see Table 3) were not used in the study, there is a clear prediction that they should elicit LAN effects at the post-filler and post-gap positions as well. In summary, the results of this ERP study are consistent with the results of Stromwold et al.'s (1996) PET study in relating syntactic STM to the cortex in the vicinity of Broca's area.

3.3.2.5 Attentional Control

The last component of the syntactic comprehension system whose neural substrates I will consider is the processing resource of attentional control. I will argue that one brain area—the anterior cingulate cortex—definitely contributes to attentional control, and that two others—the ventrolateral prefrontal cortex and the basal ganglia—are likely to contribute as well.

In the discussion of basal ganglia-thalamocortical circuits in Chapter 2, I mentioned several studies which suggest that the anterior cingulate cortex plays an important role in monitoring the activities within particular mental domains and, when necessary, influencing response selection in a top-down fashion by suppressing inappropriate responses and promoting appropriate ones (see §2.1.1.4, pp. 23-5). And in the discussion of attentional control earlier in this chapter, I argued that similar operations may be needed in order to regulate template selection and linking during the on-line processing of cer-

tain types of challenging grammatical constructions, such as object-gap relative clauses (see §3.2.2.2, pp. 98-107). Thus, a clear prediction is that the left anterior cingulate cortex should be activated when subjects process these types of constructions, but not when they process simpler types of constructions, such as subject-gap relative clauses. This prediction has received support from two recent PET studies.

The first study is the one by Stromswold et al. (1996) that I described in the previous section. As I said before, the researchers found significant activation in Broca's area for subject-object relatives (e.g., *The limerick that the boy recited appalled the priest*), but not for object-subject relatives (e.g., *The biographer omitted the story that insulted the director*), which suggests that Broca's area is involved in syntactic STM. In addition, the researchers reported significant activation in the left anterior cingulate cortex during the processing of subject-object relatives but not during the processing of object-subject relatives. It is reasonable to suppose that this activation reflects the intervention of attentional control to regulate template selection and linking. More precisely, the activation may reflect the attentional processes involved in, first, detecting an impasse in on-line sentence processing and, second, responding to this impasse by suppressing the incorrect subject-gap template and its associated linking pattern and promoting the correct object-gap template and its associated linking pattern.

It is noteworthy that there is a connection here, albeit a tenuous one, with my earlier speculations about how linking operations may be mediated neurophysiologically by synchronized oscillatory patterns among participating cell assemblies (see §3.3.2.3, pp. 128-31). Recent theoretical and experimental work on how such patterns may provide a solution to the more general "binding problem" has led to the hypothesis that attention is one of the chief mechanisms influencing the buildup of patterns and the competition between rival patterns. Thus, it may be the case that the left anterior cingulate cortex regulates linking during sentence processing by manipulating the synchronized oscillatory patterns between the relevant cell assemblies.

The second PET study is the one by Grossman et al. (1992a) that I mentioned earlier in the sections on the neurobiology of parsing and interpretation (see §3.3.2.2, p. 122, and §3.3.2.3, p. 127-8). Recall that in one condition subjects were presented with a series of sentences and had to determine whether each one contained an adjective, and in another condition subjects were again presented with a series of sentences and this time had to determine whether the actor in each one was female. In both of these conditions, significant activation was observed in the left anterior cingulate cortex. A justified speculation is that in the first condition this brain area serves to amplify the processing efficiency of parsing operations and detect the presence of adjectives, while in the second condition it serves to amplify the processing efficiency of linking operations and detect the presence of female agents. A problem with Grossman et al.'s study, though, is that several different types of constructions—actives, passives, and both subject-gap and object-gap relatives—were included in each condition; hence, it is impossible to tell if one type of construction elicited greater anterior cingulate activation than the others. As a consequence, Grossman et al.'s PET study does not provide as strong support as Stromswold et al.'s for the hypothesis that attentional control is needed more for the processing of object-gap relatives than for the processing of subject-gap relatives.

Further support for the view that the left anterior cingulate cortex contributes to the on-line processing of object-gap relative clauses comes from the ERP study by King and Kutas (1995) that I mentioned in section 3.2.2.2 (p. 102). Recall that these researchers monitored the electrophysiological activity of subjects' brains while they viewed subject-subject relatives and subject-object relatives (e.g., *The reporter that attacked the senator admitted the error* vs. *The reporter that the senator attacked admitted the error*). The results showed that the two consecutive verbs in the subject-object relatives elicited a long, sustained negativity over the left lateral frontal and left central frontal recording sites; by contrast, the two corresponding positions in the subject-subject relatives did not elicit this kind of waveform. In line with Kluender and Kutas's (1993) ERP study of

syntactic STM, it is likely that the negativity observed over the left lateral frontal site indexes the reactivation of the semantic properties of the filler NP for linking purposes (note that this NP must be linked not only to an argument of the embedded predicate but also to an argument of the matrix predicate). As for the negativity observed over the left central frontal site, King and Kutas point out that slow waves such as this "are generally taken to reflect additional processing instigated by perceptually or conceptually difficult operations (Ruchkin et al. 1988)." For this reason, they suggest that the "standing" negativity seen in their own study reflects attentional processes implemented in the anterior cingulate cortex. It is remarkable that King and Kutas's ERP data regarding the processing of subject-object relatives dovetails perfectly with Stromswold et al.'s PET data regarding the processing of the same kind of sentences: the ERP finding of a distinctive left lateral frontal waveform corresponds to the PET finding of activation in Broca's area, and the ERP finding of a distinctive left central frontal waveform corresponds to the PET finding of activation in the anterior cingulate cortex.

Although the evidence that I have presented so far strongly suggests that attentional control is implemented in the anterior cingulate cortex, the real functional-anatomical situation may be more complicated than this. Recall that Broca's area is part of the ventrolateral prefrontal cortex (BA 45, 47, inferior 46) and that, according to Petrides (1995), this large region of the frontal lobe is involved not only in holding information from different mental domains (including the linguistic domain) active in STM, but also in making judgements about this information, i.e., operating on it in a top-down fashion (see §2.1.1.4, pp. 26-7). This opens up the possibility that Broca's area contributes to both syntactic (and verbal) STM and the decision-making or executive aspect of attentional control, the aspect that is relevant to regulating template selection and linking during the on-line processing of grammatically challenging sentences, such as subject-object relatives (see also Goldman-Rakic 1995). If this is the case, then some sort of division of

attentional labor exists between the anterior cingulate cortex and Broca's area. The nature of such a division, however, is not clear.

Furthermore, it is likely that the basal ganglia contribute to attentional control for syntactic comprehension in several different ways. First of all, it is known that the basal ganglia influence the anterior cingulate cortex through a specialized circuit. In addition, if it is true that another circuit exists between the basal ganglia and the ventrolateral prefrontal cortex, and if it is true that Broca's area plays a role in the decision-making aspect of attentional control, then this provides another route by which the basal ganglia could influence attentional processes. Based on the discussion of the basal ganglia in Chapter 2, it is reasonable to suppose that these two circuits may operate in the following manner. During on-line sentence processing, the basal ganglia receive continuous input from the left temporal cortex and recognize in this input morphosyntactic, lexical, and semantic cues that are relevant to the parsing and interpretation of complex constructions, including ones with noncanonical linking. The basal ganglia then translate these cues into a recommendation for template selection and linking and relay this information up to the anterior cingulate and ventrolateral prefrontal cortices, where it serves to bias decision-making. Yet another way in which the basal ganglia might contribute to attentional control is through the mesocortical dopaminergic projection system (although this innervation may be weak in the anterior cingulate and ventro-lateral prefrontal cortices, but see footnote 14, p. 134). Dopamine may function in the cortex by reinforcing or boosting certain attentional processes that have been successful in similar contexts in the past—e.g., processes that enable the proper detection of, and response to, impasses that occur when the syntactic comprehension system is confronted with complex noncanonical constructions.

Empirical support for the view that the basal ganglia contribute to attentional control comes from the PET study conducted by Grossman et al. (1992a). In both of the critical linguistic conditions in this study—first, viewing sentences and determining whether each

one contains an adjective, and second, viewing sentences and determining whether the actor in each one is female. Significant activation was found not only in the cortical areas described previously, but also in the left caudate nucleus and the left thalamus.⁶ As I mentioned above, however, the design of this study makes it impossible to tell if these activations were stronger for the sentence types that putatively require attentional control (e.g., object-gap relative clauses) than for the sentence types that do not (e.g., simple transitive actives). The PET study conducted by Stromswold et al. (1996) should also provide information about the role of the basal ganglia in sentence processing; however, the basal ganglia were not among the authors' designated ROIs (i.e., regions of interest), and as a result they do not mention the basal ganglia anywhere in their paper. This is unfortunate, since it would be interesting to know if the basal ganglia were activated more during the processing of subject-object relatives than during the processing of object-subject relatives, as the ideas presented above would predict.

3.3.2.6 *Summary*

I have reviewed a wide range of evidence concerning the neural implementation of the syntactic comprehension system. Taken together, the evidence leads to the following hypotheses. In the vast majority of the population, the syntactic comprehension system is realized in the left hemisphere, and the primary factors influencing variable lateralization are handedness, familial handedness, and sex. Within the left hemisphere, the major components of this system are realized primarily in the perisylvian cortex.

A few studies have claimed that certain areas of the perisylvian cortex are especially important for syntactic comprehension—e.g., Broca's area, the temporoparietal region, and

⁶ Grossman et al. do not indicate whether the ventral striatum was also activated. This is important because it is the ventral striatum, not the caudate, that is involved in a circuit with the anterior cingulate cortex. The caudate might be involved, however, in a circuit with the ventrolateral prefrontal cortex, but nothing is known about this for sure. It is possible that the ventral striatum was also activated in the study, but that this activation could not be distinguished from the activation of the caudate.

the anterior superior temporal cortex—but other studies have shown that damage restricted to any of the three lobes in the perisylvian cortex—frontal, parietal, and temporal—may or may not affect syntactic comprehension abilities. Given these latter findings, it is possible that the neural substrates of the major components of the syntactic comprehension system vary across individuals. Nonetheless, additional evidence suggests that there are "central tendencies" for the localization of these components. There seems to be a tendency for parsing to be implemented in the anterior sector of the perisylvian cortex (perhaps especially in the anterior superior temporal cortex) and for interpretation to be implemented in the posterior sector (perhaps especially in Wernicke's area). In addition, it appears that syntactic STM tends to be implemented in the anterior perisylvian cortex (primarily Broca's area) and that this processing resource may receive support from the basal ganglia. Finally, there are good reasons to believe that several brain areas contribute to attentional control: the anterior cingulate cortex, Broca's area, and the basal ganglia.

3.4 Predictions about Syntactic Comprehension in Parkinson's Disease

Now that a rough multilevel model of the normal syntactic comprehension system is in place, it can be used as a frame of reference for formulating and testing predictions about what types of English constructions we would expect to be easy or difficult for early-stage, nondemented PD patients to comprehend, given the background information about the neuropathology and neuropsychology of PD provided in Chapter 2. This final section of Chapter 3 is devoted to formulating predictions about the syntactic comprehension abilities of PD patients, and Chapters 4 and 5 are devoted to describing previous studies as well as new studies that address these predictions.

I will begin by briefly recapitulating the most distinctive features of PD. Neuro-pathologically, PD causes degeneration of the dopaminergic projection systems of the basal ganglia. The nigrostriatal system is affected in all patients, leading to dysfunction in the putamen and, to a lesser degree and in only about 50% of patients, in the caudate. The dopamine reduction in the putamen and caudate prevents these structures from efficiently processing their cortical and thalamic input, and this in turn leads to the "demodulation" of several areas in the frontal lobe that receive the output of the basal ganglia. Since the putamen is involved in the motor circuit, all patients develop motor disorders, and since the caudate is involved in the dorsolateral and orbitofrontal circuits (and perhaps also a ventrolateral circuit), around 50% of patients develop cognitive disorders. The mesocortical dopaminergic projection system is also compromised in PD, but not as severely as the nigrostriatal system. This leads to moderate dopamine depletion not only in the ventral striatum, which is involved in the anterior cingulate circuit, but also in a number of limbic sites (amygdala, hippocampus) as well as in the frontal cortex (these projections may be more dense in the dorsal than the lateral or mesial regions). This aspect of the disease contributes to the cognitive deficits exhibited by patients.

Neuropsychologically, around 50% of early-stage, nondemented PD patients display a kind of "environmental dependency syndrome" which is similar to that seen in patients with lesions to the prefrontal cortex. In a variety of cognitive domains, they perform well on tasks that provide clear, explicit guidelines for behavior, regardless of whether the tasks are simple or complex, routine or nonroutine; however, they perform poorly on tasks that require self-regulated problem-solving or response formation, or internal attentional control for shifting from one mental set to another or for maintaining a given mental set in the face of interference from competing ones. More generally, the patients suffer a decline in their ability to concentrate, to flexibly alternate among different trains of thought, and to construct imaginary scenarios without environmental support. It is likely

that this constellation of "executive" or "frontal type" impairments arises because cell assemblies in the prefrontal cortex are no longer receiving appropriate "boosting" or reinforcement from either the relevant basal ganglia-thalamocortical circuits or the mesocortical dopaminergic innervation.

Given this neuropathological and neuropsychological profile, how would we expect PD patients to perform with the kinds of English constructions described in this chapter? The most straightforward prediction is that around 50% of PD patients should have trouble understanding those constructions that depend on attentional control. As shown in Table 4 (p. 107), these constructions are as follows:

- a. relative clause:
 - i. subject-object relative: *The man that Sally saw knows me.*
 - ii. object-object relative: *I know the man that Sally saw.*

- b. cleft:
 - i. object cleft: *It was the man that Sally saw.*

- c. raising-to-subject:
 - i. subject-to-subject raising:
 - a. noncanonical: *Sally seems to Harry to be tall.*
 - ii. object-to-subject raising:
 - a. noncanonical: *Sally is easy for Harry to see.*

I argued in section 3.2.2.2 that all of these constructions are likely to require attentional control in order to suppress a high-frequency, canonical template and linking pattern and promote a low-frequency, noncanonical template and linking pattern. Although the foregrounding and backgrounding passive constructions (e.g., *Harry was seen/Harry was seen by Sally*) also involve low-frequency, noncanonical templates and linking patterns, they should not require attentional control because, unlike the constructions listed above, they contain multiple overt morphosyntactic cues that signal their atypical status;

moreover, they are structurally simpler than the relative, cleft, raising, and control constructions, since they only involve a single core. With regard to the active and passive undergoer-control constructions (e.g., *Harry persuaded Sally to be nice/Sally was persuaded by Harry to be nice*), I argued that even though the determination of the "controller NP" depends on implicit semantic properties of the matrix verb, attentional control should not be required because these semantic properties are easily accessible; in addition, for the passive version of the construction, the noncanonical status of the matrix core is clearly marked. As for the undergoer-intransitive construction, it involves implicitly signaled noncanonical linking, but special attention should not be needed to regulate the linking process because, first, the constituent structure is extremely simple, and second, as with the undergoer-control constructions, the proper linking strategy is determined by readily available semantic properties of the verb.

With regard to the underlying neurobiology, I argued in section 3.3.2.5 that the decision-making component of attentional control which is responsible for regulating template selection and linking in a top-down manner has a distributed implementation in the brain. Several sources of evidence indicate that the anterior cingulate cortex is crucially involved. In addition, the proposals advanced by Petrides (1995) lead to the possibility that the ventrolateral prefrontal cortex (more narrowly, Broca's area) contributes as well. Finally, there are both theoretical and empirical reasons for believing that the basal ganglia contribute in two ways—first, by means of circuits with the anterior cingulate cortex and, in theory, with the ventrolateral prefrontal cortex; and second, by means of meso-cortical dopaminergic innervation of these two cortical sites. Since it is known that the anterior cingulate cortex is "demodulated" in PD because of disruption of both the associated basal ganglia-thalamocortical circuit and the meso-cortical dopaminergic projection system, this provides a solid neural foundation for the prediction that PD patients should have difficulty understanding the constructions listed above. Furthermore, if it is really the case that Broca's area contributes to attentional control for syntactic compre-

hension, and if this area really is influenced by the basal ganglia through both a circuit and direct dopaminergic innervation, then this provides another avenue by which the processing of the constructions listed above could be impaired in PD.

What about the other subsystems that are necessary for syntactic comprehension? I argued that parsing—i.e., assembling constituent structures and assigning syntactic relations—tends to be implemented in the anterior portion of the left perisylvian cortex. It is difficult to localize parsing operations more precisely than this, but, taken together, the studies reviewed in section 3.3.2 suggest that it is more likely that these operations are carried out in the anterior superior temporal cortex than in Broca's area. If this hypothesis is correct, then we should not expect PD patients to have a basic parsing deficiency, since, although the anterior superior temporal cortex sends input to the basal ganglia, it does not receive output from the basal ganglia, nor does it receive a significant dopaminergic innervation from the mesocortical projection system. On the other hand, it is important to note that other modulatory neurotransmitter systems are mildly affected in early PD. In particular, the ascending cholinergic system is affected in many patients (DuBois et al. 1983, 1991), and this system projects rather densely to the superior temporal cortex, especially in the left hemisphere (Amaducci et al. 1981). Thus, it is possible that acetylcholine plays a role in parsing (Mimura et al. 1995), and that the reduction of this chemical in PD degrades the efficiency of parsing to a slight degree.

Next, consider interpretation—i.e., the processes of accessing the LSs and macro-roles of predicates and establishing correspondences between NPs, macroroles, the arguments of predicates, and the concepts encoded by nouns. I argued in section 3.3.2.3 that the LSs and macroroles of predicates may be stored in the posterior region of the left perisylvian cortex—i.e., in Wernicke's area and perhaps also the supramarginal gyrus (BA 40). If this localization hypothesis is correct, we would not expect the LSs and macroroles of predicates to be disturbed in PD, since this cortical region is not involved in a reciprocal circuit with the basal ganglia, nor does it receive a significant

dopaminergic innervation. Nonetheless, there is still the possibility that the mild cholinergic depletion mentioned above adversely affects semantic processing in the posterior perisylvian cortex to a slight degree. Another point that I made in section 3.3.2.3 was that correspondences between lexical, syntactic, and semantic structures are dynamically formed during on-line sentence processing via synchronization of the firing rates of the cell assemblies that represent the relevant structures. Thus, the establishment of correspondences could be disrupted if brain damage causes the "time windows" of activation for cell assemblies representing different types of structures to become discordant.

Accounts of this sort have been offered for developmental dysphasia, dyslexia, and the syntactic comprehension deficits in agrammatic Broca's aphasia (Llinás 1993; Merzenich 1993; Friederici 1995; for a theoretical discussion of lesion-induced "heterochrony" i.e., slowing of some neural processes vis-à-vis others see Brown 1988, 1994). With regard to PD, some researchers have claimed that patients exhibit bradyphrenia, i.e., a general slowing of thought processes, a kind of "psychic akinesia" (Rogers et al. 1987; Morris et al. 1988; Pillon et al. 1989). However, in a thorough review of the literature on this topic, DuBois et al. (1991) show that significant cognitive slowing only occurs during the performance of complex tasks that draw heavily on executive functions subserved by the prefrontal cortex. For this reason, it does not seem likely that the basic ability to establish correspondences during syntactic comprehension is impaired in PD patients.

The last component of the syntactic comprehension system to be considered is syntactic STM. I argued in section 3.3.2.4 that this processing resource tends to be implemented in the anterior region of the left perisylvian cortex, probably in Broca's area. Hence, whether we should expect PD patients to have intact or impaired syntactic STM depends on whether this cortical area participates in a reciprocal circuit with the basal ganglia and/or receives a significant dopaminergic innervation from the meso-cortical projection system. In both this chapter and the previous one, I have made it clear that although there is a good chance that these anatomical pathways exist, their reality has not

been confirmed. The status of syntactic STM in PD is therefore an open issue to be resolved through psycholinguistic experimentation. The studies that I will describe in the next two chapters will provide evidence bearing on this issue.