Computational Intelligence, Volume 18, Number 1, 2002

CONTROL AND CROSS-DOMAIN MENTAL COMPUTATION: EVIDENCE FROM LANGUAGE BREAKDOWN

William Frawley

Department of Linguistics and Cognitive Science, University of Delaware

This paper uses the notion of control from programming languages to look at the organization of mental code. Data for the analysis comes principally from language breakdown. The paper first outlines the well known distinction between logic and control in algorithms and argues that the same distinction holds in mental code. Discussion then focuses mainly on control—the management of data flow—and shows that a variety of language disorders affect either the logic component of the mental algorithms for language (e.g., Specific Language Impairment) or the control component (e.g., Williams syndrome and Turner syndrome). A comparative study of the loss of morphology in Williams syndrome and Specific Language Impairment reinforces the logic/control split as an accurate guide to the explanation of linguistic behavior in these disorders. The data, moreover, are not accountable to sheer performance factors, but to the way the disorders disrupt the structure of mental algorithms. The paper closes with a discussion of how control and the management of cross-domain computation fit into recent theories of modular mental architecture and proposals about the explicitness of representations and their availability to working memory.

Key words: cognitive science, philosophical implications, knowledge representation, natural language, control, mental architecture.

CAVEAT LECTOR

What follows is a speculative paper that seeks to bring together claims about mental architecture, information transfer and tracking in programming languages, modularity in linguistic theory, and neurodevelopmental disorders. This is admittedly a tall order, and by trying to speak in one voice to four or five quite different audiences—and stay within manageable page limits—I fear that I may have annoyed readers as much as pleased them. I have been forced both to egregiously understate some issues so as to promote readability and egregiously overstate others in order to nail down points. For example, readers from AI will correctly observe that there is much more to how programming languages manage control than what I have said, and an entire book can (and should!) be written on how linguistic interface rules ought to be sensitive to the formalisms of computable interface rules: in trying to suggest that control might simply provide a vocabulary for talking about cross-domain mental computation, I have admittedly understated the case. In contrast, I have gone into specifics on how one view of morphology works, and how it might account for a variety of technical details from experimental psychological work: here I have probably overstated the case, but I did so for the sake of illustration. While not trying to shirk my responsibility for balanced exposition, I must acknowledge the risk that this kind of unevenness is in some sense intrinsic to the territory. I would like the readers to see this paper as a thought-piece, certainly one in need of many corrections, but also one that is an invitation to expansion.

1. COMPUTABLE MENTAL CODE

In the computational theory of mind, domains of knowledge are collections of *computable* representations and *computable* operations over such representations. This

Address correspondence to William Frawley at the Department of Linguistics and Cognitive Science, University of Delaware, Newark, DE 19716. E-mail: billf@udel.edu.

claim is axiomatic to virtually all accounts of cognitive architecture. It is obviously part and parcel of symbolic, "sentence-crunching," but also essential to the alternatives—input-sensitive recurrent networks, vector coding, tensor products, and varieties of constraint satisfaction. The classic connectionist-symbolist debate, in point of fact, turns on *how abstract or interactive* computable mental code is—on what is *believably* computable—not on computability itself.

The commitment to computability seems to be consistent with positions that are otherwise skeptical about *mental* computation. Proposals like situated action (Suchman 1987), externally distributed representation (Zhang and Norman 1994), and existential cognition (McClamrock 1995) push mind outside of its architecture, into the environment, making the world its own best model and thus ridding the mental architecture of computable representations. In this view, intrinsically representational culture is the operating system of a multi-trillion gigabyte external hard drive.

But even if it is true that, as Dennett (1996, p. 124) says, humans have the "habit of *off-loading* as much as possible of our cognitive tasks into the environment itself—extruding our minds," this in no way precludes computable mental code. Unless we are empty zombies bouncing around at the mercy of richly symbolic affordances that we ourselves have made—as if off-loading were a kind of perpetual clean-install, which I doubt (though see Brooks 1991; Braitenberg 1984)—the symbolic environment must still be registered *in minds* for it to be effective for the agents who have off-loaded it into existence. The social world must have a landing site in computable mental code (see Frawley 1997 for additional arguments in this vein).

Even anti-representationalist dynamical systems theory (hereafter DST) is not antithetical to computability. In DST, mind is an instantiation of a geometric space of continuous interdependent variables characterizable by the changes across states, not the structure of the states themselves (Thelen and Smith 1994; Port and van Gelder 1995; van Gelder 1998). In rejecting mind as intrinsically representational states mediating input and output, DST commits to a real-time, situated, analogue system explained by differential equations, not code. But while DST is not obligated to be a computational theory of mind at all (Thagard 1996, pp. 169–281), it is consistent with the theory of computability. As van Gelder (1998) points out, a dynamical system is computable, just not a computer. So even if mind is a dynamical system, it could also be one that is effectively computable.

To say that the mind's code is computable is to say that mentalese, to use Fodor's (1983) term, is implementable—that there is not only a language of thought (however abstract) but a *programming* language of thought. An important dichotomy comes with this requirement of implementable mental code. Programming languages are made out of algorithms, and algorithms have two components, what Kowalski (1979a, 1979b) famously distinguished as *logic vs. control*, roughly "data structures" vs. "information flow." Much of the energy in cognitive science has been spent on the former—the logic of mental computation—because somehow data structures seem to be the mind's true content.

But early on in the development of the computational theory of mind, Pylyshyn (1984, pp. 78–86) argued for equal time for the study of control structures:

> The commitment to a model that actually generates token behaviors forces one to confront the problem of how, and under what conditions, internal representations and rules are pressed into service as actions are generated. These are questions that concern *control* of the process. Although a central topic in computer science, they have almost never been raised in cognitive psychology... [B]ecause control issues are a major area of study in computer science, progress in developing computer models of cognitive processes will very likely depend on technical ideas originating in that field. (pp. 78–79)

Indeed, in some proposals in artificial intelligence, control plays an increasingly important role. It is one of the real advances in blackboard architectures (Hayes-Roth 1985), and hybrid connectionist-symbolic models often rely on explicit control for beneficial effects (Schneider and Oliver 1991). Some views, in fact, have mind entirely as a control system (Sloman 1993).[1]

Even with these concerns for computational control, data flow and information management remain the poor relations in cognitive science, no doubt because they seem more an issue of performance than do data structures (a point I will dispute below). In contrast to this trend, I want to argue that control structures are an essential part of computable mental representations, are a competence issue, and deserve study in their own right (see Frawley 1997 for a variety of arguments). Crucial evidence for this comes from language breakdown, which, I hope to show, manifests a distinction between logic disruptions and control disruptions. The latter, moreover, are not traceable to general processing factors. Rather, these breakdowns suggest something fundamental about the organization of the code of mental computation.

2. WHAT IS CONTROL?

Control is thought of in two forms in computation. The first is control within the programming environment proper, which involves managing the flow of the execution of a program: sequencing information, handling interrupts and exceptions, and overseeing the tradeoff and coordination of data across chunks of program, such as subroutines and coroutines. The second is control of real-time computing. In this sense, control is the way a system continuously monitors input and output in order to respond appropriately under real-time pressures. I will focus principally on the former kind of control (which I will hereafter refer to simply as *control*), but the latter also deserves serious investigation—especially since Damasio's (1994) work suggests that the brain has an area dedicated to monitoring the real-time fit between its decisions and the environment and since theories like DST see mind as a real-time device.

2.1. Computational Control

Kowalski's (1979a, 1979b) work on the nature of programming languages underscores the importance of distinguishing the knowledge in a program from the efficient manipulation of that knowledge. For example, as Kowalski (1979b, p. 129) observes, *if X, then Y* statements can be understood as instructions to the computer on how to manipulate the declarative information in X and Y (logic) to solve the problem of the relation of X to Y in a top-down fashion (control).

The independence of logic and control means that the overall behavior of a computational system can be affected by a modification in either one: e.g., logic could break down separately from control (as we will see in humans). Moreover, control can be more or less explicit, depending on the programming language and the style of the programmer. Structured programming, for example, is an attempt to build control implicitly into the programming environment itself and has come about as a response to brute-force,

[1]All the more reason to see putative computational alternatives like DST in computational terms. If mind is a kind of dynamical system and hence is explained by measuring changes across collections of variables (mental states) rather than by characterizing the variables themselves, mind must have powerful control structures for managing the information changes across mental states as collections of variables.

all-purpose, explicit control mechanisms like GOTO, which can send the flow of computing anywhere.

Control comes in two forms: *machine-level control*, hardware constraints on data flow, and *high-level control*, control structures of a particular programming language (Fischer and Grodzinsky 1993; Ghezzi and Jazayeri 1987; Teufel 1991). Booting up, for example, is a machine-level control process directing the machine to load the operating system and transfer control to it. Sequencers (*and*) and conditional statements (*if/then/else*), which link chunks of program code, are examples of high-level control.

There is an important further distinction in high-level control that speaks directly to the nature of mentalese. High-level control structures can be *statement-level* or *unit-level*, differentiated by the range of code to which control applies. Statement-level control is local and affects individual statements and expressions in the program: typically this involves sequencing, selection, and repetition of information. Unit-level control is global and affects chunks of a program or collections of statements and expressions—*program units*. Unit-level control is involved when control must be passed from one program unit to another (e.g., in subroutines), when the system crashes, and when, for recovery, control must be passed to some program unit or to the user.

The effects (or even the existence) of these two types of control very much depend on the nature of the programming language in which they are implemented and the kinds of processing and memory demands that compiling the code requires. Some programming languages lack what are known as *statements*—forms that do not return a value when called. Statements contrast with *expressions*, which do return values. Those programming languages that lack statements (e.g., LISP) are known as *functional languages*. In contrast, *procedural* or *sequential languages* (e.g., C) have both statements and expressions. Both kinds of languages have control, but as a consequence of their constituent forms, they manage the flow of data quite differently. Functional languages are constituted by forms that deliver outputs, and so control is communicated and passed via the memory stack, where the outputs are recorded. But procedural languages do not have to manage control via the memory stack because they have forms that do not deliver outputs and so communicate via the program environment directly.

Consider an example of variation in the execution of statement-level control. In LISP, a functional language, conditional control is an expression which can be nested in other code, allowing a hierarchically structured, embedded conditional sequence. But in languages that have conditional statements, not expressions, conditional control is managed in a sequential way. Pascal, for example, uses a conditional statement, and so communicates with the rest of the code via the program environment, not the memory stack.

Similar variations can be found in unit-level control, where, for example, the effects of control under breakdown also depend on the programming language (Teufel 1991). Most languages have an *exception handler*, a piece of code designed to respond to specific interrupts, but different languages handle crashes differently: in some, the exception handler returns computing to the point where the interruption occurred while in others, computing is terminated and does not return to the point of the interrupt. These two strategies have different effects and processing demands.

2.2. Control in Mental Code

If the structure of the code in programming languages affects the way computation is managed, do these effects likewise transfer to control in mentalese? Does the nature of the programming language of thought also carry with it particular cognitive demands on information management? More specifically, is mentalese a functional or procedural language? Is mentalese LISP or C? (Again, see Pylyshyn 1984, pp. 78–86, for suggestive observations.)

One way to begin to get an answer is to see how mental domains communicate with each other, a classic unit-level control problem.[2] A first guess would be that intramodular control—say, the communication between phonology and syntax in language, or between low-level and high-level properties of objects in spatial knowledge—would look to be procedural. There would seem to be no need to report and record output via an independent memory stack across related domains because the information communicated is relatively close. On the other hand, intermodular control might look more functional. The coordination of, say, language with motor programs for speech would seem to place greater demand on the mind's resources and might require explicitly recorded outputs.

The analogy also goes through in considerations of interrupts. When processing crashes, how and where does mentalese return the mind to mental computing? There are generally two computational strategies for resetting the system: (1) at the point of the crash or (2) at some earlier point (often the initial state) to clear the system completely. Remarkably, there is some evidence for both of these options in crashes of human mentalese (Kaczmarek 1987).

While these points are admittedly speculative, they do raise empirical questions. Are there types of control in mentalese, just as there are for programming languages? Does intramodular control break down differently from intermodular control? Importantly, these questions would not arise without considering mentalese as the *programming* language of thought. Now we turn our attention to trying to answer some of them by looking at how the mental algorithms for language fail.

3. CONTROL DISORDERS (VS. SPECIFIC LANGUAGE IMPAIRMENT)

Over the past fifteen years or so, a number of congenital disorders of global knowledge have been the focus of intense investigation: Williams syndrome, Turner syndrome, spina bifida with hydrocephalus, autism, and a variety of conditions either unlabeled or vaguely characterized. Interest in these conditions has grown despite their lack of common—or, sometimes, even known—etiology. Williams syndrome (1 per 25,000 births) appears to result from abnormalities in elastin production and the neuropeptide CGRP, which affects calcitonin, and consequent hypercalcemia (Udwin 1990; Udwin and Yule 1991; Bellugi, Wang, and Jernigan 1994); very recent work appears to have identified the genetic locale of the disorder (Frangiskakis et al. 1996; Lehnoff et al. 1997). Turner syndrome (1 per 2,500 females only) is caused by abnormalities in the X chromosome (White 1994; McCauley et al. 1987). Spina bifida with hydrocephalus (1 per 1,000) results from a neural tube defect during the first month of intrauterine life; most cases also involve hydrocephalus—cerebrospinal fluid in the ventricles of the brain (Reigel 1993). Autism (1.5 per 1,000) has essentially unknown etiology, though recent studies implicate hydrocephalus, tuberous sclerosis, Fragile-X phenomenon, rubella, metabolic disorders and, increasingly, the cerebellum (Courchesne et al. 1994). The other conditions are a kind of grabbag: Yamada's

[2]This presumes, of course, an architecture with dedicated processing areas, however abstract and circumscribed one's theory might have them be. Obviously this rings true to modularists (Fodor 1983), but even devout connectionists have to admit dedicated processing, if only in weak form (see Elman et al. 1996). Karmiloff-Smith (1998) has a nice argument running between the two schools of thought—dedicated processing areas emerge developmentally, as domain-relevant processing becomes domain-specific with the functional reorganization of representations.

(1990) description of what might be called *Laura's syndrome*, Blank et al.'s (1979) *language without communication*, other conditions labeled *nonverbal learning disabilities* (Rourke 1988) and *right hemisphere dysfunctions* (Kaplan et al. 1990). These syndromes are frequently of unknown etiology and have varying incidence.

What makes these conditions of interest to linguists and cognitive scientists is that they ostensibly reveal a dissociation between nonlinguistic world-knowledge and domain-specific linguistic knowledge, frequently leaving a deficit in the former. Linguistic development is often slowed, but mostly has a normal outcome—even superior in some cases. The conditions are thus crucial to claims about domain-specificity, encapsulation, and modular mental architecture since they show that normal linguistic development can proceed independently of world knowledge. (Karmiloff-Smith's 1998 arguments that the linguistic domain is not fully preserved and that there is no hard dissociation of cognition and language are discussed in more detail below.)

Still, a close look at the cognitive and behavioral manifestations of these disorders does not reveal Nature cutting the mind-brain so neatly at its joints. There are a number of unexplained (or incorrectly explained) linguistic disruptions in these syndromes. My claim will be that these syndromes can affect the computational mechanisms involved in the *coordination of linguistic domains* rather than within-language representations and thus are a particular type of computational disruption. In effect, they preserve computation but have defective report. They are thus breakdowns of the *control component of mental code*, very much like the well known disorders of consciousness, such as blindsight. As such, they contrast with specific language impairment (SLI), which preserves report, but report of defective representations, and hence is a disorder of the logic component of mentalese. Unfortunately, these two kinds of disruptions often look alike because they affect opposite sides of mental algorithms. Control disorders involve interface-management breakdown and so can look like a logic deficit because of failure in information coordination. SLI involves defective computation within a knowledge domain itself—a logic breakdown—but this can surface as an apparent failure in information management because the representations themselves are affected. We can see this distinction by examining what is preserved or lost in the control disorders and comparing those findings to the effects of SLI.[3]

3.1. What Exactly Is Preserved or Lost in Control Disorders?

Control disruptions can be found throughout language. We will restrict our consideration to four standard components: phonology, morphology, syntax, and semantics.

3.1.1. Phonology. Phonological performance in these syndromes is generally excellent, even superior to normals in some cases. Individuals are notoriously fluent (Cromer 1994; Yamada 1990, p. 59–61)—even auditorily hypersensitive (Neville, Mills, and Bellugi 1994; Schopler 1994, p. 91)—and often better at verbal imitation than normals (Morrow and Wachs 1992; Karmiloff-Smith et al. 1995; Karmiloff-Smith et al. 1997). Remarkable testimony to this phonological capability is that Williams children, who have a

[3]After this paper was initially written (in 1998), work by Clahsen and Almazan (1998) came to my attention, a paper that contrasts Williams and SLI. I will show in section 3.1.2 that their arguments, while correct in characterizing the syndromes as opposites, are narrowly linguistic and in the end need a more comprehensive framework for explanation. Their findings about Williams syndrome turn out to be a particular case of computational control, which, in turn, brings their paper into contact with the larger issues of mental computation (which they hint at but do not explicate) and the variety of language deficits that characterize Williams syndrome.

marked deficit in reading because of disruptions in spatial integration, can sometimes be taught to read through *phonetics*: whole-language, global knowledge approaches fail (Udwin, Yule, and Martin 1987; MacDonald and Roy 1988). Moreover, these individuals can repeat words as forms but are unable to explain their meanings, suggesting preservation of the phonological aspects of lexical form (Morrow and Wachs 1992; Dennis et al. 1987). This would also mean that segmental phonological representation is unaffected. Indeed, when Laura is given phonological parts of words, syllables for example, she is able to complete the word normally on the basis of segmental phonological information alone (Yamada 1990, p. 61).

However, more detailed examination of the clinical and experimental reports shows some notable failures. Laura (Yamada 1990, p. 61) appears to have marked prosodic problems. The same sorts of prosodic difficulties have been observed in Turner girls. Silbert et al. (1977, p. 18) report significantly lower scores than normals on tests of rhythm, tonal memory, and auditory figure-ground structure.

By most accounts, prosody is either a template separate from segmental information—since it operates over groups of segments—or found in the late (postlexical) rules of lexical phonology. Moreover, prosody is the one feature of phonological knowledge that interfaces with other mental modules, arguably with what Fodor (1983) calls the central system since prosody is known to be a reliable signal of planning (Garman 1990, pp. 121–133). It is here that individuals with control disorders—and ostensibly preserved linguistic information—falter in performance.

One might then argue that what is really going on in these patterns of preservation and loss is an intact ability to manage the flow of phonological data across representational levels in phonology, but significant problems where this flow of information has to be managed across modules. In effect, they have phonological computation but defective report. This is a classic issue in *unit-level control* issue, a description further supported by the patterns of preservation and loss in other kinds of linguistic knowledge.

3.1.2. Morphology. It has often been argued that morphology is essentially unaffected by these disorders. Tager-Flusberg (1994, p. 190) observes this about autistics, and the same claims have been made for Williams syndrome (Thal, Bates, and Bellugi 1989), Turner syndrome (Yamada and Curtiss 1981), and spina bifida with hydrocephalus (Cromer 1994). The unnamed syndromes also manifest remarkably good morphology (Yamada 1990; Curtiss 1981). Perhaps most remarkable in this respect is the individual in Smith and Tsimpli's (1995) study, a savant who can communicate in some fifteen languages but apparently learns them principally by morphological analysis.

The major exception to these accounts of preserved morphology is Karmiloff-Smith et al.'s (1997) study of gender knowledge in French-speaking Williams children, whose morphological performance is well below normal. (Yamada 1990, p. 38, also notes that Laura has some performance difficulties in morphology also, although her difficulties seem restricted to comprehension only, and do not affect all morphology.) Karmiloff-Smith et al. (1997) take this to indicate that Williams syndrome can have *within-domain deficits* (within morphology), which, to them, calls into question the whole idea of a modular architecture because it would suggest that the language module, otherwise dissociable as a block from general cognition, is not uniformly preserved. But there are many other possible explanations, as Clahsen and Almazan (1998) also observe—accounts, in fact, that piggyback on the control claim above with respect to the preservation of segmental phonology but disruption of prosody.

Using a picture identification task, Karmiloff-Smith et al. (1997) tested normal and Williams children on their use of real and nonce French words with concord and discord between article and ending, plus nonce words with a suffix clue only. For example, subjects were given an imaginary object and told *ça, c'est un matton* ("that's a *matton*") or *ça, c'est deux bicrons* ("those are two *bicrons*"), both phonotactically possible but nonexistent French words. They were later shown pairs of identical pictures of the object, each of a different color; another object was hidden under one of the pictures. Subjects were then asked to identify the locale of the hidden object by completing a phrase: *J'ai caché ma bague*. . . ("I hid my ring. . ."). A response with the correctly inflected article and agreement on the color adjective revealed their intuition of the gender of the nonce or real form: e.g., *sous le matton vert*, ("under the green *matton*"), *sous la plichon grise* ("under the grey *plichon*"), etc.

Unsurprisingly, for both groups, real words are easier than nonce words. The Williams children are superior to normals on repeating nonce words, but since Williams children are known to be hypersensitive to phonology, their superior performance in repetition of nonce words is not unexpected. However, despite their greater chronological age and higher scores on language measures, the Williams children are much worse than normals on the agreement task. They perform much worse than normals on using article cues for adjective agreement (saying *la fourni vert* instead of using the article *la* as a cue to the correctly agreeing *verte*) and on using nominal endings to infer gender when the articles give no clue (saying *le faldine vert* instead of *la faldine verte* when given the cue *faldines*, which has a feminine nominal ending).

Karmiloff-Smith et al. (1997) explain these failures as a *morphological representation deficit*. They argue that arbitrary generalizations, like gender, have to be learned early and must be semantically bootstrapped; since language development in Williams children is delayed, they have missed the crucial time in which to acquire proper morphology. French gender is tied to input frequency, which, by some measures, Williams children are not sensitive to. More tellingly, French gender is multi-cue based—a variety of lexical, phonological cues signal gender—and so the system is intrinsically irregular. All this suggests to Karmiloff-Smith et al. (1997) that Williams children have an alternate course of morphological development, more like second-language learning, thus undermining the fashionable use of Williams children as exemplars of the independence of language and world knowledge.

But there are alternatives to this account of Williams syndrome as a disruption of the logic of morphological algorithms. For one thing, there is no compelling reason why arbitrary features of the grammar have to be learned early. For another, the multiply-interactive status of French gender has long been observed as a central factor in the explanation of native language learning (Tucker et al. 1968), second language learning (Vuchic 1993), and computer modeling (Sokolik and Smith, 1992), a point Karmiloff-Smith et al. (1997) themselves acknowledge in observing that French gender depends on the interplay of a number of features not tied to morphology proper. In point of fact, this morphological deficit surfaces *only when the subjects have to manage multiple sources of information* (see Section 4.1 for more detailed argument). In this respect it is important to consider Clahsen and Almazan's (1998) counters to Karmiloff-Smith et al. (1997), which, while linguistically compelling, can ultimately be cast in a broader, more computationally sensitive, and hence more explanatory light.

Clahsen and Almazan (1998) show that morphological failures in Williams children are restricted to problems with irregulars. In their account, irregularly inflected forms are stored as lexical templates with internal structure (subnodes) that cues their dependence on phonology and other out-of-domain information. As such, irregulars are unlike regularly inflected forms, which are stored as stems and affixes and adhere to rule-based morphological combinatorics. Their experiments show that Williams children are like normals on the latter but not the former, and they trace the failures to lack of access to the content of the lexical templates. For Clahsen and Almazan (1998), then, Williams children have normal computational competence (regulars) but disordered associative competence (irregulars).

Three points are worth noting here, all of which return us to computational control. First, to say that the lexical representation of irregulars contains cues to out-of-domain information, like phonology, is to say that irregulars pose a problem of computational control. The subnodes in the lexical templates are a signal to the processor to run a subroutine or co-routine into phonology while the main program (lexical representation) is being run: this is a classic problem in dataflow management. Williams children have the representations, but miss these control cues. Second, to make a distinction between computational and associative competence in an otherwise wholly computational mind is not to push the explanation very far. *Associative mechanisms are also computational*. But how so? What Clahsen and Almazan (1998) call *computational* is really within-domain combinatorics, and what they call associative is cross-domain tracking: respectively, logic and control. But both are computational, just different sides of the algorithms. Third, as we have seen with phonology and will see with syntax and semantics, Williams children have a variety of linguistic deficits, not just problems with morphological irregularity. A problem with subnode retrieval in lexical representations is but one manifestation of a broader computational deficit, namely the management of cross-domain computation. Thus, all these problems and their explanations come together in a single account via the control component of computable mental code.

The ostensible within-domain morphological deficit in Williams children surfaces within a domain whose rules require the management of several domains at once. The failures, then, on French gender are much like failures on prosody.

3.1.3. Syntax. A similar story can be told for syntactic performance. Williams syndrome (Thal, Bates, and Bellugi 1989; Mervis et al. 1999, who argue for delayed development with essentially good outcome), Turner syndrome (Yamada and Curtiss 1981), spina bifida with hydrocephalus (Cromer 1994), and the various unnamed syndromes also show good syntax (Yamada 1990; Curtiss 1981), as does autism, at least in some studies (Tager-Flusberg 1994, p. 190). Laura, for instance, has full phrase structure, including operators and empty categories, and understands transformational relations across structures and the difference between well-formed and ill-formed constituency. Christopher, the savant in Smith and Tsimpli's (1995) study, likewise has intact syntax despite serious deficits in inferences with world knowledge.

Perhaps the most striking indication of the preservation of syntactic knowledge in these syndromes is that individuals have good metagrammatical performance. Cromer's (1994, p. 150) hydrocephalic subject has excellent metalinguistic judgments of grammaticality and ungrammaticality, as does Laura (Yamada 1990, p. 36). Turner girls are fine at verbal completions and so must know lexical and syntactic structure. Williams children are quite competent at correcting grammatical anomalies (Bellugi, Wang, and Jernigan 1994, p. 28), as is Christopher (Smith and Tsimpli 1995, pp. 44–60).

Still, if one looks closely at the reported data, there are some errors in syntactic performance. Notable here is the work by Karmiloff-Smith et al. (1998), who have studied the implicit and explicit syntactic processing of Williams individuals vs. normal controls. In the implicit processing task, subjects had to monitor a word (in CAPS below) in both grammatical and ungrammatical constructions in which the ungrammaticality resulted from a violation in phrase structure (*I expect special the PILLS...*), auxiliary structure (*he might expecting*

SPEECHES at the. . .), and subcategorization (*Maria always needed for PARTNERS*). Assuming that word monitoring taps on-line processing and hence implicit syntactic knowledge, they predicted that individuals with intact syntactic representations ought to devote more monitoring time to the word in the ungrammatical constructions than in the grammatical ones, given that the word appears at the point of ungrammaticality, thus forcing re-processing. Indeed, normals and Williams showed comparable differences between grammatical and ungrammatical constructions for word monitoring time for both the phrase structure and auxiliary structure, where they devoted more monitoring time to the word in the ungrammatical constructions. But for subcategorization constructions, while the normals had the same monitoring differences between grammatical and ungrammatical constructions, with more time devoted to the word in the ungrammatical items, Williams individuals showed no time difference between grammatical and ungrammatical constructions.

In the explicit task, subjects had to match pictures to sentences. Assuming that such a task requires individuals to use the syntactic representation of an expression and to explicitly match the propositional content, those individuals with intact syntactic knowledge ought to be better at the task than those with syntactic deficits. Given the expression, *The clown photographed the policeman*, subjects had to pick the correct picture from a set of three: one that correctly depicted the proposition, one that reversed the roles, and one distractor. Normals did this task well, as might be expected, but Williams individuals performed very poorly, most often choosing the reversed role picture.

Karmiloff-Smith et al. (1998) take these findings as evidence that Williams syndrome does selectively affect syntactic representation. Interestingly enough, the explanations they give suggest that the deficits are computational-control disorders, and not disruptions in the syntax per se. As for the failures on the sentence-picture matching task, Karmiloff-Smith et al. acknowledge that the sentence-picture task does "not map into language processing directly, but also involve[s] a number of additional cognitive processes" (p. 347). That is, successful performance on the explicit task requires the management of information across at least two domains—syntax and semantics—a typical instance of computational control (see also Section 5.2, where the issue of explicitness is explored more fully).[4]

Their explanation for the failure on subcategorization is even more telling: Williams individuals "are able to access subcategory information associated with a verb, but they are slow to integrate this information with the incoming input" (p. 348). The reason for this is that Williams individuals are slow overall on subcategory processing, whether grammatical or ungrammatical, and so this suggests that these kinds of constructions as a whole pose a unique processing burden. Indeed, in the architecture of language, subcategorization is located *at the interface of the lexicon and syntax*; it is not a phenomenon wholly within the syntactic module. Jackendoff (1997, p. 102) says:

> lexical constraints in syntax (e.g., subcategorization. . .) are imposed. . . at S-Structure, where the lexicon interfaces with syntax. . . [T]here is no need to preserve throughout the derivation all the lexically imposed structural relations. They need to be involved only at the single level of S-Structure.

Thus, the crucial evidence Karmiloff-Smith et al. (1998) adduce for a within-domain deficit in syntax, failure on subcategorization, is in fact a cross-domain computational loss: a

[4]Cromer (1994, p. 150) remarks that the only case in which his subject fails to perform normally in syntax is on the judgment of double-object constructions (nonprepositional datives), which are also classic cases of syntax-semantic interactions.

problem in computational control, not in logic. Once again, the ostensible exceptions to the rule prove the larger point about control at the interface of components.

3.1.4. Semantics. In semantics, a similar pattern emerges, with these individuals exhibiting strong performance on the logic or core of the interpretation module, but failures where this information must be reported out or interface with other domains. Admittedly, the organization of semantic interpretation is currently a matter of debate: compare Jackendoff's (1990) argument that interpretation involves a set of conceptual primitives recruited specifically for grammatical meaning with Heim and Kratzer's (1998) argument that interpretation consists of relatively content-free logical mechanisms for grammatical interpretation. For consistency with the assumptions about the other modules of language discussed above, I will take semantic interpretation as logical and minimal and guided by rudimentary formal mechanisms (e.g., functions and set operations) that establish denotation and point to conceptual content stored outside of the interpretation module.

Clinical and experimental results suggest that all the syndromes preserve the formal mechanisms of interpretation.[5] There are no indications of problems in the denotational function, the elemental computation of formal semantic competence (Bach 1989). Cromer (1994) reports good denotation by individuals with spina bifida, Tager-Flusberg (1991) the same for autistics, and Bellugi et al. (1990) for Williams children. The disorders also appear to preserve the operations of logical form. Laura (Yamada 1990, p. 49), for instance, knows how to interpret variables, though what she gets wrong are the specifics of the world knowledge that fill the variable—just what we might expect if the syndromes leave *semantic form* unaffected. Furthermore, all the syndromes leave intact the form and relational organization of the mental lexical network (Yamada 1990, p. 43; Bellugi, Wang, and Jernigan 1994, p. 32; Temple and Carney 1993, p. 696; Mervis et al. 1999). Tyler et al. (1997) survey challenges to the claim for normal semantics in Williams children and find that semantic representation and access to lexical meaning are in fact preserved.

But there are some notable semantic deficits, and they fall into a by-now-familiar computational pattern. Williams and spina bifida children exhibit hypersemantic excitation: the entire lexical network is activated by a prime (Bellugi, Wang, and Jernigan 1994, p. 49). This problem appears to be a deficit in the very sort of computational function that control structures perform: inhibition of computation and restriction of computation to a certain range of representations.

These syndromes also cause problems in tasks requiring semantic/pragmatic interactions (Dennis et al. 1999), like presupposition and deixis (Tager-Flusberg 1991; Cromer 1981; Bellugi, Wang, and Jernigan 1994, p. 44), implicature, mental search for appropriate lexical items (Tager-Flusberg 1991), explanation and judgment of the meanings of words (Dennis et al. 1987), and the management of word lists as lists (Cull and Wyke 1984, p. 181), even though these individuals otherwise exhibit high verbal IQ and quite normal lexical behavior (Bender, Linden, and Robinson 1989). Significantly, these tasks require cross-domain computation—either the management of semantic and pragmatic knowledge or the coordination of semantic knowledge with behavioral control. That is, while they might look like within-domain semantic deficits, they are really

[5]Although I argue that the logical mechanisms are preserved and hence come down on the side of the formalists, there is also evidence that grammatically relevant conceptual categories, such as boundedness, specificity, animacy, and shape, are also normal in these syndromes (Sigman 1994, p. 142; Tager-Flusberg 1994, p. 179, footnote 4; Yamada 1990, p. 43; McGlone 1985; Morrow and Wachs 1992; Mervis et al. 1999), and so a theory like Jackendoff's (1990) is not necessarily ruled out.

specific kinds of cross-domain breakdowns. Moreover, they are not simple performance deficits since, again, the tasks vary widely in their performance pressures and yet have the same results.

However, perhaps the cleanest evidence for splitting logic from control in semantic disruptions comes from Stevens and Karmiloff-Smith's (1997) study of lexical acquisition, otherwise designed to show that Williams children have within-domain semantic deficits. Stevens and Karmiloff-Smith (1997) examine the deployment of four lexical principles by normals and Williams children in learning word meanings: fast mapping (new names map to new objects), mutual exclusivity (each object has only one name), whole-object scope (words are taken to initially denote entire objects, not their parts), and the taxonomic constraint (additional referents for a word will be categorial, not thematic: e.g., if *blick* refers to a robin, another *blick* with no robins around will evoke an animal (categorial), and not, say, a nest or branch (thematic)). Stevens and Karmiloff-Smith (1997) find that Williams and normals alike use fast mapping and mutual exclusivity, but diverge on whole-object scope and the taxonomic constraint. For this reason, they argue that Williams children have an alternate course of lexical-semantic development (but see Mervis et al. 1999, who argue for delayed rather than deviant lexical development).

Here again the exceptions prove the rule. Williams and normals are identical on the most basic and content-free lexical principles: fast mapping is the equivalent of the denotational function—i.e., that forms map to the world—and mutual exclusivity is a logical restriction on fast mapping—i.e., that initial mapping must be one-to-one: no many-to-one or one-to-many functions allowed. These two principles are just the sort of abstract formal mechanisms that constitute the within-domain logic of the semantic component. If control disorders preserve logic at the expense of computational control, then we ought to see fast mapping and mutual exclusivity operative in their lexical learning.

However, whole-object scope and the taxonomic constraint require reference to conceptual content: both state restrictions on the value for the range of the mapping function. That is, if semantic interpretation is constituted by functions that take well-formed expressions as their domain and map them into ranges of conceptual content, then whole-object scope restricts the range to individuals (roughly indivisible wholes) and the taxonomic constraint restricts the range to properties (boundedness, animacy, etc.: roughly the conceptual basis of intensions). Unlike fast mapping and mutual exclusivity, whole-object scope and the taxonomic constraint require cross-domain operations by linking well-formed expressions with conceptual domains.[6]

3.2. SLI: A Nutshell Contrast

The control disruptions outlined above contrast markedly with behavior in specific language impairments (SLI, also known as *developmental dysphasia*). Control disorders and SLI both dissociate language and world knowledge, but have, I would argue, algorithmically complementary effects on the language component (in contrast to the kind of explanation given by Clahsen and Almazan 1998). SLI children show competence in

[6]These results might allow us to say something about the brain's semantic interpretation module. Whole-object scope looks remarkably like the interpretation that gives rise to proper nouns and the taxonomic constraint like that for common nouns: the contents of proper names and common names are stored *outside of Wernicke's area*, the brain's semantic module, which looks to be basically just a set of pointers. What better testimony to a split between logic and control: Wernicke's area is constituted by interpretation functions (logic), but the range of these functions—how these map across domains—is a control issue.

world knowledge but clear deficits in linguistic representation. Control disorders leave their sufferers with relatively preserved linguistic representation, but the inability to coordinate that knowledge across domains. In effect SLI and control disruptions cut the computational mind at different algorithmic joints: SLI disrupts the logic of mental language-algorithms, while control disorders affect the control component of the algorithms and hence surface as impairments in the way the linguistic knowledge base interfaces with other domains.

Consider one of the standard findings from SLI—that children with SLI have marked morphological deficits. Some theorists argue that morphological impairment is the essential characteristic of the disorder (Crago and Allen 1996).[7] All of the explanations for this morphological vulnerability trace the impairment to a deficit in linguistic representation—even the counter-proposals, I would argue.

Among the accounts of morphological deficit in SLI are Gopnik's studies (Gopnik and Crago 1991, e.g.) of the familial aggregation of the inability to analyze lexical forms into their morphological features. In Gopnik's view, morphological SLI is caused by feature-blindness preventing the representation of abstract morphological paradigms and a consequent inability to construct morphological rules (Gopnik 1990; Gopnik and Crago 1991). Other, related theories with different slants also trace the problem to a deficit in the logic of morphological algorithms (Crago and Allen 1996).

Even direct counters to these representational theories ultimately support the hypothesis of a deficit in the logic component of the mental algorithms that constitute morphological knowledge. Leonard (1989, 1998) has proposed, for example, that SLI individuals lack the ability to analyze certain kinds of phonetic input, and, failing to construct the proper phonological representation of word forms, also fail to make the proper morphological analysis using these phonologically misanalyzed forms. But the phonetic properties he proposes as deficient are very much those that constitute the *logic of the algorithms* used for analyzing phonetic input. He finds that SLI children have "segmental inaccuracies" likely traceable to underlying "phonological representations [that]... are also relatively unstable" (Botolino and Leonard 2000, p. 144). His other work shows that SLI children have a deficit in the detection and representation of prosody and segmental duration, and this in turn translates into an inability to analyze phonetically reduced morphological forms; failures here in turn ramify to syntax, where complementizers and other phonetically reduced morphemes are also affected by SLI (Botolino and Leonard 2000 is a nice cross-language study of this phenomenon). But one might then argue that what Leonard has shown is not that SLI children have a general processing deficit, but a phonological impairment that then percolates into morphology as a consequence of the report of defective representations (see Gopnik and Crago 1991, p. 4, footnote 2—although they also report, p. 36, that the subjects had no auditory impairment). Indeed, he notes that phonological and morphological impairment can be independent (Botolino and Leonard 2000, p. 145), just as one might expect if the disruption is one of logic, not control.

It thus appears that whatever the proper analysis of morphological SLI (and each side concedes some ground to the other), the deficit is characterizable as an impairment *in information to be reported* across domains (SLI) rather than a deficit in *the mechanisms that pass* the information across domains (control disorder). Both types of impairment produce problems as a consequence of how domain-specific information fails to ramify throughout

[7] If morphological impairment is the core deficit in SLI, then this would make an interesting contrast with the behavior of Smith and Tsimpli's (1995) subject, Christopher, whose superior language abilities appear to be driven by his *preserved morphological component*. That is, Christopher and SLI children would seem to be in complementary distribution.

the architecture, although traceable to different aspects of the computational structure of mental algorithms. (Again, note the contrast of this computational explanation with that given by Clahsen and Almazan 1998: see also Karmiloff-Smith 1998, who points out how small neurocognitive deficits can have large-scale effects.) As Gopnik and Crago (1991, p. 14) say with respect to the organization of morphological knowledge: "A given underlying grammar constrains not only individual forms [*logic*, WF], but also the way in which these forms interact with other parts of the grammar [*control*, WF]."

A similar logic deficit that ramifies throughout the architecture can be found in syntactic problems in SLI. Van der Lely (1994) reports on SLI children with a deficit in the linking of syntactic and semantic roles. More particularly, on tasks requiring the children to judge the semantic roles of NPs in sentences with nonce verbs, SLI children and normals are no different on tests that require forward linking—using the semantic roles as clues to the syntactic roles. But on reverse linking, where syntactic clues alone have to be used to determine semantic roles, SLI children, unlike normals, are significantly impaired.

Van der Lely (1994) offers two possible explanations: SLI impairs either (1) the canonical schemas for linking semantic and syntactic roles or (2) the information necessary for identifying the syntactic frame that is then linked to the semantic roles. In other words, SLI is either a problem in the rules that make syntax and semantic visible to each other—a control problem—or a disruption in the syntactic information itself that feeds the syntactic-semantic interface—a logic problem. Advancing the latter, van der Lely argues that if SLI involved a disruption in syntactic-semantic schemas, there should be symmetric breakdown and problems in both forward and reverse linking. But since the impairment is asymmetric, it appears that the deficit lies in the syntactic component itself—*in the representations that feed the interface*. It is thus a deficit in "the syntactic representation which specifies the relationship between the verb and the argument positions" (van der Lely 1994, p. 64), which results in the inability to analyze the verb-argument frame in sufficient detail in order to report it to the semantics. So what might look like an interface disorder is really a deficit in what precedes the interface: syntactic-semantic control is preserved here, but syntactic logic is impaired.[8]

While we have contrasted SLI with control disorders to push the issue of logic vs. control, there is nonetheless an important feature common to both deficits. *Neither is a performance problem*: the impaired behavior does not surface as a consequence of resource or processing difficulties (although Leonard 1989, 1998 offers some counters). Van der Lely (1994, p. 62) expressly dismisses nonlinguistic processing demands. Hadley and Rice (1996, p. 237) do also for their SLI subjects, whose delayed emergence of auxiliary *be* they trace to morphological deficits ramifying on the syntax, not to performance or processing factors.

This not to say that these disorders have no performance effects—only that we must be clear about where and how performance factors enter the explanation (Leonard 1998, p. 237–268 makes a similar point). If SLI and control disorders are caused by problems in the structure of the algorithms of mentalese, then there will most certainly be performance effects since the subjects engage in real-time processing in the experiments. But performance itself is not the cause. An interesting result from Gopnik and Crago's (1991) study supports this line of reasoning. On the tests where normals and SLI individuals differ,

[8]One might still argue, however, that this syntactic deficit is in fact an interface disorder and not an issue of the logic of syntactic algorithms per se. Assuming that the control rules across syntax and semantics are bidirectional, what might be affected here is just the syntax-to-semantics mapping, with rules in the other direction preserved. While such an account seems to merely restate the findings, not explain them, it does raise the empirical issue of whether cross-domain computation is bidirectional and, if so, whether it can be impaired in only one direction (some theoretical work on this can be found in DiSciullo 1997).

the latter also spend significantly more time on task, a processing effect. But on cases where they are not different, as in the test of argument structure (since SLI individuals in their study have normal syntax), there are no differences in time on task. Insofar as time on task is a performance matter, we should expect a processing deficit associated with the algorithmic deficit—after all, the processes are operating over impaired representations. But on tasks where neither group has an algorithmic deficit (i.e., normal), the processing and performance factors level out.

4. MORPHOLOGY IN WILLIAMS SYNDROME (AND SLI): A CONTROL EXPLANATION

Now that we have established the relevance of the logic/control distinction to language breakdown (and hence to mental architecture generally), we turn our attention to a fuller computational-control explanation for one cross-domain disorder that has played a crucial role in the ongoing arguments about mental architecture: morphology in Williams syndrome. The logic-control structure of mental algorithms has a clear place in Aronoff's (1994) view of morphology and the use of this theory to explain aspects of the control disorders discussed above. For Aronoff, morphology is an autonomous component that interfaces with both syntax and phonology. There are two kinds of principles that operate entirely within the domain: (1) those that determine abstract word classes (rules of the formation of paradigms) and (what he calls) realization pairs (statements that relate word classes to their rules for surface form, or distribution statements); (2) those that handle the interaction of morphology with syntax and phonology (morphosyntactic rules that correlate syntactic and lexical forms; realization rules that translate morphological form into surface phonology).

In terms of Kowalski's model of computation and algorithms, word classes and realization pairs constitute the logic of morphological algorithms. For example, with gender, the logic of morphological algorithms states the gender classes of a language and the distribution of these classes with respect to surface form. Morphosyntactic and realization rules, in contrast, constitute control because they manage the flow of morphological data across domains. For gender, these control rules state how the syntactic notion of gender is visible to morphology and point to the surface phonological forms of the classes themselves. Moreover, morphological control is bidirectional: there are cases where morphology both determines and is determined by syntax and phonology. This entire picture can be schematized as follows:[9]

[9]Interestingly enough, this picture and its division between logic and control give a computational rendering of a universal condition on the relationship between morphology and phonology generally. Languages with a lot of bound inflectional morphology on nouns appear to have innumerable declensional paradigms as a result of the many surface manifestations for case, gender, number, person, etc. But Carstairs (after Aronoff 1994, p. 64ff.) observes that there is a limit to the number of paradigms, what he calls paradigm economy: the total number of inflectional classes in the language equals the highest number of morphophonological manifestations of any particular word form. That is, if one morphological form can surface at most five different ways, then the language will have at most five abstract word classes.

Note how paradigm economy is in fact a statement about the overall visibility and cross-indexing of morphological and phonological information. Essentially, the principle is a meta-control statement: the maximal realizability of any particular form is reapplied to the classes of forms themselves—what holds most for one form holds for the abstract classes themselves. This sounds very much like call-by-value parameter passing in functions. The highest numerical value of the output of any single realization function is then passed as a value back into the morphological component as the value for the number of classes themselves. This suggests, moreover, that the flow of information between morphology and phonology is a kind of loop, where the information returned to morphology enters at a place in the computation different from where the loop started. Hence, Carstairs's principle of paradigm economy looks very much like a straightforward control statement in a programming language.

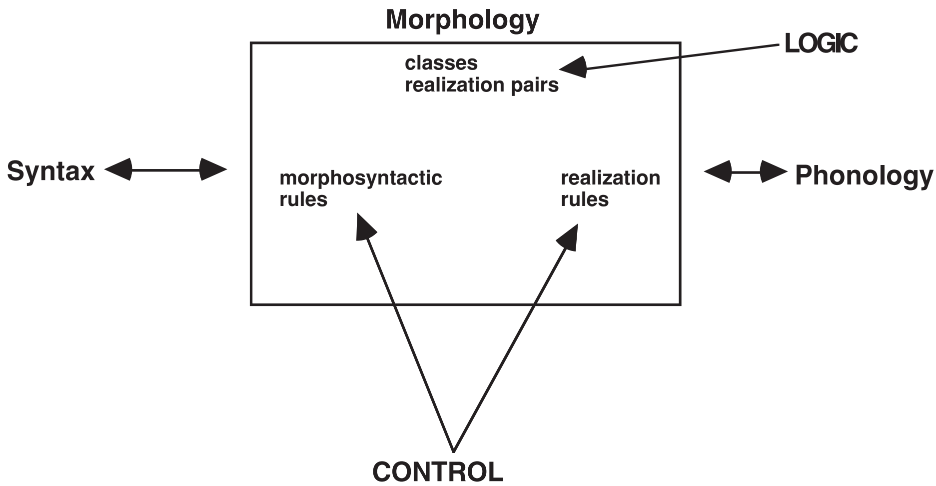

FIGURE 1. The logic and control of the morphological module.

Aronoff's account of morphology via logic and control not only gives a computational rendering of part of the language faculty, but also provides a framework for the Williams syndrome data and its contrast with SLI. Recall that Karmiloff-Smith et al. (1997) found that French-speaking Williams children fail at two points in gender assignment: They do not mark the whole NP for gender, and they have difficulty using nominal endings alone to infer the gender of the noun. Both these problems are traceable to inverse realization rules, not to the core of morphology itself, which remains intact (unlike in SLI).

In Karmiloff-Smith et al.'s (1997) study, Williams children are much worse than normals on agreeing the entire NP containing adjectives and nonce nouns (after Karmiloff-Smith et al. 1997, p. 250):

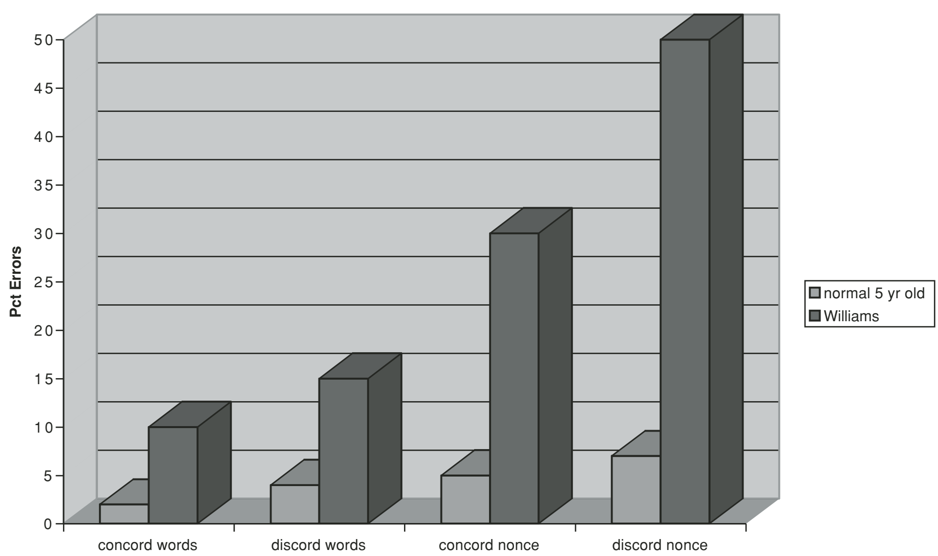

FIGURE 2. Errors on phrasal agreement.

However, most of this failure comes from errors on agreeing the adjective itself (after Karmiloff-Smith et al. 1997, p. 253):

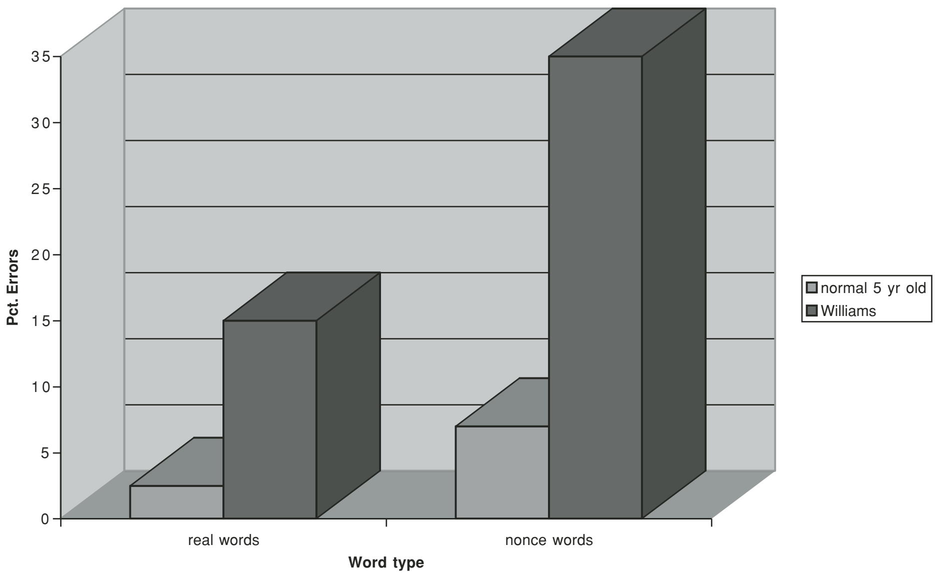

FIGURE 3. Errors on adjective agreement.

They appear to be able to infer the gender from the noun—which they should be able to do, if their ability to construct abstract inflectional classes is preserved—but they fail to spread it throughout the phrase. In this respect, they are very much like second-language (L2) learners of French. Vuchic (1993) has found that L2 learners know the gender of French nouns and understand the spreading outward of this feature within the phrase, but they spread it best leftward and are markedly deficient at spreading to the postnominal adjective (*la maison vert* or *pommes vert*).[10] In Aronoff's account, spreading of the agreement is a morphosyntactic process, a matter of how the adjective inherits the value of the suffix for the syntax. It is thus a deficit in the flow of information into morphology from syntax, not a purely within-domain morphological problem.

Now consider the second morphological deficit observed by Karmiloff-Smith et al. (1997). Williams children perform much worse than normals on inferring the gender of a nonce lexical item from the ending. That is, given the model like *deux faldines*, they fail to see that the ending patterns with other feminine nouns (after Karmiloff-Smith et al. 1997, p. 253):

[10]Karmiloff-Smith et al. (1997) also find that Williams children are deficient in inferring the gender of a noun from an isolated plural model and so conclude that they have a morphological deficit, i.e., cannot deduce morphology on its own terms since the plural neutralizes the cues. But their measure of failure is again on the full NP response: it is not clear whether these failures are carried by the postnominal adjective (as in the case above) or article or both. Moreover Vuchic (1993) has found that gender and number are independently acquired in L2 French, with gender acquired first. So it is not clear what effect the number manipulation may have had on performance. In any case, inferring the gender on the basis of the phonological form of the ending alone is a morphophonological task, not a purely morphological one, so we again have a case of the interaction of levels, not a within-domain strategy.

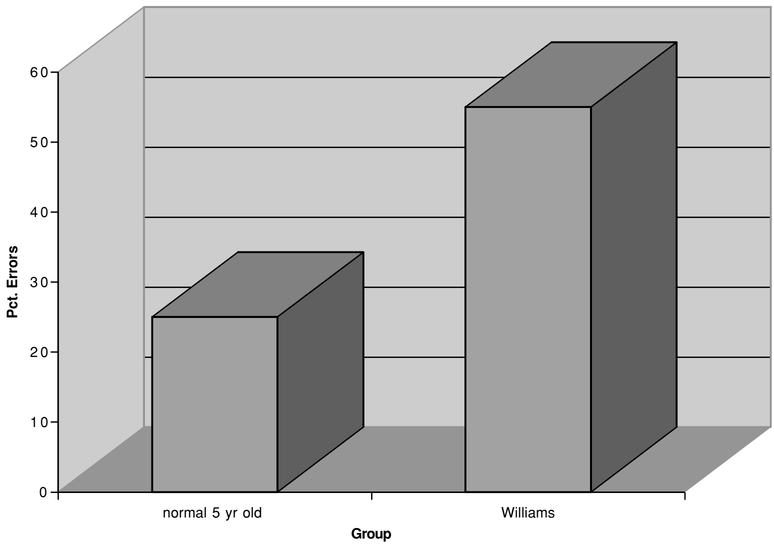

FIGURE 4. Failure to make use of cues.

Again, this is a deficit in the interface, not abstract morphology. This task requires that the children take a phonological manifestation and translate it into a morphological class. It is thus an inverse realization rule, to use Aronoff's term, because the flow of computation is from the phonological module into the inflectional class. The Williams children have trouble with this while the normals do not. In many ways, this deficit is the other side of the previous one. Whereas the failure to spread can be traced to a disruption in the in-flow of information from syntax into morphology, the failure here is a consequence of disrupted in-flow of phonology into morphology. The phonetic clues are not visible to the abstract classes or morphological arrays. Thus Karmiloff-Smith et al.'s (1997) data on morphological failure in WiIliams children are the manifestation of a *single problem*: they reflect deficits in in-flow of information into the morphological component, one from syntax and the other from phonology.[11]

Interestingly enough, independent corroboration of this view of a morphological control deficit can be found on tests of the morphological performance of individuals with neurological impairment. French agrammatic aphasics tend to perform well on local gender assignment (articles) but fail on nonlocal assignment (pronouns). This is taken as support for the claim that agrammatism is not a deficit in abstract representation since the errors appear to be a function of failure to compute gender only non-locally (Jakubowicz and Goldblum 1995; Jarema and Friederici 1994).

Moreover, Badecker et al. (1995) have shown, from data on Italian and French anomics, that there is a distinction between a mental *lexical* level, where morphosyntactic features are assigned, and a phonological-orthographical one, where surface form is computed. Their subject Dante, who has no access to phonology and thus cannot name words, nonetheless knows their gender. This suggests that in some sense gender is an independent morphological feature. But another subject, GM, has both levels intact and performs worse!

[11]Vuchic (1993) also observes that L2 learners have a correct representation for gender agreement but still have an incorrect surface string; the last stage in the acquisition of this type of agreement is to get the phonological features correct. This suggests another correspondence between Williams children and L2 learners of French—both have the morphological representation intact but problems in the management of morphophonological features throughout the whole NP.

If Dante has morphology and phonology severed, then there is no competition for output because there is no interaction of levels. So Dante has "pure" access—morphology by itself. But GM has the interactions preserved, if defective in their form, so the interactions are the problem: whereas Dante has good morphology because his lexical component is freed up, GM has poor performance because he has good phonology. The Williams children in Karmiloff-Smith, Grant, and Berthoud's (1997) study behave more like GM than Dante: They have depressed performance because they are trying to manage levels, not because they have a within-domain deficit: if anything, this is a cross-domain deficit.

In short, other experimental results indicate that the Williams children perform very much like L2 learners, agrammatic aphasics, and anomics with preservation of morphological-phonological interaction. This suggests either a performance problem or a problem in the correlation of domains. Now these may be two ways of saying the same thing, but I doubt it. Certainly, the correlation of domains has capacity and execution effects and so looks like a performance issue. But relatively constant results emerge on experiments that vary widely on the performance pressures placed on the subjects: from immediate verbal or nonverbal response to delayed such response, to response that can elicit help from the experimenter. If these results are performance problems, then they must be so in a very unusual sense of *performance*.

We might instead cast them as problems in interface management. Since these patterns emerge only where the subjects have to control information across domains, they may be the result of representations that guide computational control, one function of which is to regulate cross-domain computations. How else should we explain an ostensible within-domain deficit when the individuals otherwise have normal core word-formation rules? These results are again the exceptions that prove the rule.

Needless to say, this explanation contrasts sharply with that given for morphological impairment in SLI. Gopnik, Crago, and others repeatedly find an inability in SLI individuals to construct abstract inflectional paradigms. Because they treat lexical items as unanalyzed wholes, their core morphological representations are not rich enough to subserve grammatical and phonological interactions. Thus, whereas Williams children have preserved morphological representation and defective report out of this information to other domains, SLI has defective morphological representation and preserved report of this defective information. These problems can give the appearance of a similar disruption, but in fact the Williams children, in their language at least, might be more accurately said to suffer from a data-management problem rather than one of abstract analysis.

5. THE ARCHITECTURE REVISITED: MENTAL MODULES, EXPLICIT VS. IMPLICIT REPRESENTATIONS, AND WORKING MEMORY

The claims of this paper—that the algorithms of mental code have both logic and control and that different language disorders are associated with each—in the end return us to larger issues of mental architecture. I want to close this paper by looking at two proposals: (1) Jackendoff's (1997) claims about mental architecture and cross-module information transfer and how the foregoing arguments find a place in overall considerations of the language faculty; (2) Dienes and Perner's (1999) useful recasting of the declarative/procedural distinction as explicit vs. implicit representation and how this bears on both the formal structure of cross-domain information management and the nature of working memory—the latter often implicated in explanations of the disorders reviewed in this paper.

5.1. Crossing Domains in the Language Faculty

Jackendoff (1997) argues that the mental language module consists of three autonomous submodules, phonology (PF), syntax (SS), and semantics (LCS), all linked by interfaces. This means that there are two general types of representation—those that constitute domain-specific information itself and those that make information from one domain visible to another. While the details of Jackendoff's proposal are a matter of some dispute—e.g., not everyone believes that semantics is equivalent to lexical conceptual structure (LCS) or that logical form and the lexicon are interfaces—Jackendoff's basic point about the organization of the architecture is essentially uncontroversial (among modularists, that is).

Since the point of this paper is that control disorders disrupt cross-domain communication (in contrast to SLI), what Jackendoff says about interfaces and the nature of interface rules is of particular interest because the management of information flow across interfaces is a problem of computational control. What do linguistic interfaces require? Here is a short list of properties of interfaces and how they might be recast as computational control: again, caveat lector—this discussion is meant only as a way to initiate a vocabulary for explaining cross-module transfer, not the final word (the ideas that follow owe much to helpful discussion with William Idsardi).

(1) Modules are mismatched and the interfaces impose homomorphism (there is no isomorphism). As a consequence, different modules preserve or modify different things in their interactions. Consider morphology as a domain that must interact with syntax and phonology. In receiving information form syntax, morphology must be structure-sensitive and essentially compress the detailed trees it receives into usable word-level structures. But on output, where morphology feeds phonology so that structures can be said in real-time, morphology must be sequence-sensitive. These two different tasks in turn suggest that the procedures that relate morphology to each of these interacting modules has its own structure and manifests different patterns of breakdown.

    This picture of interface organization, it turns out, is a standard computational problem: how do you manage incongruities across program units? There are a variety of mechanisms to implement this process—from simple identity checking to the passing and change of parameters across program units. And there are more or less subtle means of managing this process. Some languages have an explicit marker (known as a *semaphore*) in the implicated program units to turn procedures on and off and synchronize processing. Others have a dedicated, separate program unit (a *monitor*) overseeing shared data and access. Still others have the program units signal each other as they operate (a *rendezvous*), allowing a kind of interleaving of cross-domain communication. We know little about what computable mentalese does in this respect, but it would seem to be safe to say that the second possibility, a monitor, is not likely. The empirical literature suggests that cross-domain interaction is structurally integrated, not located in a separate control device because breakdown does not seem to be as severe and complete as a damaged monitor would require. Nature thus seems to have engineered something like structured programming into mental code.

(2) The passing of information across modules is not a derivation but a correspondence. Derivations (i.e., progressive manipulation and change of a representation) happen *within modules*; across modules, there is *indexing*. This prohibition of derivation across modules is probably the result of processing requirements. Derivation requires that two representations be held simultaneously in working memory, and this could be accomplished within the processing demands of a circumscribed module. But across domains, it would be too demanding on resources to hold two different representations in memory *and* hold the data structures and operations of each of the different program

chunks. Consequently, mental architecture avoids this processing drain by data sharing, again a classic issue of computational control. Jackendoff's proposal that data is shared by linking is a standard (though not the only) computational solution to control.

(3) Modules do not relate deeply, but by outputs. Some information in a module is not visible to another module. This is a way of preserving autonomy in a system of interfaces. Again, the matter of morphology arises as a classic illustration. Morphological theory generally recognizes two kinds of word-forms relevant to phonology: non-derived and derived forms. This difference crucially affects the pronunciation of stems and affixes: *electric* (with a second syllable stress and final /k/) becomes *electricity* (with a stress shift and change of /k/ to /s/). How is this derivational information made visible to the phonology? Note that these forms as output need simply to be tagged as derived, not loaded with the details of how the derivation transpired and not simultaneously cross-referenced from module to module as the derivation unfolds. The receiving module of phonology need "see" only the tag to operate, not access details of what is derived and how (although recent proposals from Optimality Theory may cast this control function differently: see Kager, van der Hulst, and Zonneveld 1999). Computational control also is managed within this condition. Even in the most interactive communication across program units—concurrence, where one program unit partially activates another but neither suspends operation fully—some information in each program unit remains unavailable or irrelevant to the other. The computational issue is how to manage the interleaving of the visible information. Indeed, the morphology example looks surprisingly simple—explicit tag matches explicit tag.

(4) One module can enrich the input from another. Once information from one domain is made visible to another and then passed into the receiving module, the latter can do what it will to that information within its representational demands. Thus cross-module communication is not the mere pass-through of information but the active reprocessing of it in terms of the representations and procedures of the receiving domain. The control issue here is very much like what happens in parameter passing, where functions that share data across program chunks can call parameters by value (and preserve the value) or by address (and hence change the value of the parameter called).

A look at how these principles of cross-domain processing are fleshed out in practice with actual language data reveals a consistent solution. When representations are passed out of a module, they typically have features explicitly marked for visibility and, hence, action in the receiving domain. Erteschik-Shir (1998), for example, proposes annotated structural descriptions as a means to send focus structures from syntax into both phonology (for the correct prosody) and semantics (for the correct informational contrast). Similarly, Avrutin (1999) proposes explicitly annotated representations passed from syntax into discourse to manage nominal and pronominal structures sent from syntax into the file system of discourse representation. Indeed, many of the theoretical disputes in this area concern the mechanisms that make structural properties explicit and the extent of explicitness in the checking of features. These proposals lead to a more general question: what does explicitness have to do with it? Explicitness and the role of control in working memory seem to be the key; if so, then we can come to a nonperformance account of why individuals with such disorders manifest performance-like deficits.

5.2. Why Explicitness Matters

Dienes and Perner's (1999) recent theory of explicit and implicit knowledge gives us a way of talking about the nature and management of explicitly annotated representations across processing domains. In their theory, the relative explicitness of knowledge is a function of the overtness of the component and subcomponent representational states that comprise a knowledge structure. All representations have three components: *content* (what the representation is about), *attitude* (stance toward the representational content, "know," "believe," etc.), and the *holder of the attitude* (the subject). These components have representational subcomponents: e.g., content is comprised of a property predicated of an individual—P(x)—with factivity (truth or falsity at some time). Thus, any claim about whether knowledge is implicit or explicit is a claim about whether any of the components or subcomponents is overtly represented. By this account, what is standardly known as declarative knowledge turns out to be just one type of explicit representation, one where the subcomponents of content—the property predicated, P(x), and the factivity of this predication—are overtly represented: declarative knowledge expressly states its application to particular cases. In contrast, procedural knowledge is one type of implicit knowledge, where predication and factivity are left covert: procedural knowledge applies to cases generally by not expressly stating its applicability to a particular case.

This view of implicitness and explicitness bears directly on the form of information passed across processing domains. Representations that are within domains are preferentially implicit and procedural because they must apply generally (i.e., no explicit factivity). Hence, implicit procedural knowledge is a natural ally of module-internal knowledge. By contrast, representations reported out of a domain have to be maximally explicit in order to be checked and used. Thus, in explicitly representing predication and factivity, declarative knowledge signals *by its form* its application to particular cases: only in this explicit form can the representation be fully usable by another domain because its applicability and validity can be tracked. To put it another way, if you have no access to the occasioning of a piece of information, you may "have the implicit feeling" that it did happen, but you do not know *that* it did. Implicit knowledge is thus active, but by its form does not leave itself open to manipulation and revision. Explicit knowledge is therefore the preferred form for cross-domain processing, particularly for information sent to central processing since explicitness best serves hypothesis-checking and inference: significantly, many researchers see one of the core deficits of the syndromes reviewed above as the inability to make inferences and integrate modular and central processing.

The implications of this analysis for control and the disorders reviewed in this paper are straightforward. Control might be understood as one way that computational-cognitive systems make information explicit so as to be visible across processing domains. Indeed, these control disorders might be a function of the failure to explicitly annotate representations for delivery to working memory for use in other processing domains (vs. logic disorders, like SLI, which are disruptions of implicit domain-internal representations). If so, then the frequently argued conclusion in the cognitive and clinical literature that these disorders are best understood as disruptions of resources and space in working memory becomes substantially clearer.

Avrutin (1999) argues this working-memory deficit in his considerations of disruptions of the syntax-discourse interface. Sullivan and Tager-Flusberg (1999), citing Karmiloff-Smith et al.'s (1995) point that Williams children's modular knowledge fails to feed social computation, have claimed that Williams children falter on second-order belief tasks because of a working memory deficit that prevents them from integrating information across mental domains. A similar point is made by Levy (1996) in examining what she sees as Williams children's alternate routes of cognitive access and the availability of cognitive resources. Bishop et al. (2000) report that some forms of Turner syndrome (45, X^m) have subtle long-term verbal deficits, but these deficits, unlike those of amnesiacs, e.g., are a function of difficulties in reaccessing the represented information, presumably through explicit representations in working memory, not failures to represent the verbal information initially. Pani, Mervis, and Robinson (1999) even claim that the classic global spatial deficits in Williams syndrome are not traceable to inability in these individuals to represent the global configuration of objects, but a "general weakness in planning and in organizing information in working memory" (p. 457).

But what does it mean to claim that individuals lack resources, have a weakness in working memory, or are unable to manage on-line information in a temporary working buffer? Received theory, in fact, holds that working memory consists of submodules that process temporary sensory-specific representations and a central executive (Baddeley 1997). Which parts of working memory are implicated in these resource losses?

When resource claims apply not to all information in working memory, but to domain-specific representations, they cannot signal a generalized resource loss, a cross-domain performance deficit, or a broad working memory disruption. These working memory deficits must be caused either by a deficit in those representations that are delivered out by the representational processors or those monitored by the executive. In each case, it is the explicitness of representations that matters: working memory is the place where the outputs of computations are temporarily stored in a form accessible to other domains. If anything, claims about resource limits or working memory deficits are really claims about these domain-specific annotated structural descriptions.

Indeed, it is clear that the language disruptions of the syndromes reviewed in this paper are not symptomatic of general processing difficulties, i.e., they are not performance deficits, problems of real-time execution generally, or generalized working memory losses. The poor gender performance by French Williams children—and the association between phonology and morphology in French gender—is not one of working memory load or resource management, but representations. Karmiloff-Smith et al. (1997) *expressly eliminate* these performance factors in their study of morphological knowledge in Williams children. Their results are not traceable to difficulty in the circumstances of assessment, a typical performance issue: "it is not because people with WS have difficulty with formal test situations" (1997, p. 255); nor are they traceable to unusual attentional demands, also a classic performance factor: "it is not because the WS participants cannot cope with nonce stimuli" (p. 255); nor are they traceable to excessive demands on working memory or capacity limitations of the subjects, again a characteristic performance factor: "perhaps memory problems force them to omit the terms. . . our analysis shows this not to be the case. . . their short-term memory test age [is] well above the age of the controls. . . Memory problems do not explain the present results" (p. 255). What, then, accounts for the findings? "Their problems lie in the assignment of grammatical concord across article, noun, and adjective" (p. 255)—in short, cross-domain representational processes in morphological competence.[12]

[12]This is not to say that there are no such generalized performance or working memory deficits. Performance-based accounts of mental retardation portray the condition as a deficit in overall working memory capacity (Ferretti and Cavalier 1991). Significantly, in problem solving tasks with proper instruction and a change in the distribution of capacity, mentally retarded individuals can *markedly improve* their performance. Moreover, tests of their knowledge of problem solving strategies show that they can mentally represent the problem space and solutions, even under failure of execution. This is clearly a performance deficit because the manipulation of the capacity demands enhances performance irrespective of the representational knowledge of the subjects.

Performance must be understood as a cover term for real-time limitations and limitations intrinsic to the mind/brain as a general computing system, including such things as access, seriation, attention, allocation of resources, and speed (see Halford, Wilson, and Phillips 1998 for a good survey of general processing limits). This is not to say that these factors are unaffected by competence or representational issues. Rather, it means that manipulation of capacity or resource demands should have effects independent of manipulations of the representations themselves. Performance ought to get worse under capacity and resource stress and improve under resource and capacity facilitation. Competence should stay the same under conditions of good or bad performance, as we see with the Williams children.

Following the computational vocabulary of this paper, I would call this a control disorder. To use phraseology reminiscent of Kowalski, the annotations on structures in working memory tell processing domains how to use the declarative information from one domain to solve problems in another—how, e.g., derived or nonderived words are to be said. Control disorders, explicitness, and working memory thus come together as failures to manage the explicit annotations necessary for cross-domain processing, deficits in part of the structure of the algorithms of mentalese. This is why explicitness matters.

Indeed, if explicitness and working memory come together in control disorders, we might then make some progress on empirical and experimental work. Explicit knowledge is known to be directly testable. If control disorders result from failures to explicitly annotate structural descriptions, then individuals with these disorders ought to fail direct tests. In this light, Karmiloff-Smith et al.'s (1998) findings of poor performance by Williams children on explicit tests of syntactic knowledge is suggestive.

5.3. Banish GOTO

Importantly, as Jackendoff (1997) points out, interfaces are not conceptually necessary. One could imagine architectures constructed otherwise, with, say, direct intervention of one domain in another, or even without domains per se to interface. But the point I want to insist on is that they are *computationally necessary*—as control—since any architecture has to be computable, and that imposes a more stringent condition on its organization. Control is the regularizer of independently solved problems in separate domains. Logic gives the data structures; control prevents the system from crashing and allows recovery when it does crash. Control vs. logic gives a computational way of talking about what Nature has engineered in mental architecture.

Interestingly enough, what Nature appears to have engineered evolutionarily is what software designers have come to themselves: structured programming. Some years ago in computer science, there was a famous debate over the elimination of GOTO statements—control statements that let the flow of computation go anywhere in a program. Kowalski (1979b) pointed out that GOTO leads to spaghetti code, compiling disasters, programmer forgetfulness, and impenetrable bugs. One consequence of Kowalski's persuasive arguments has been the development of structured programming, where related program units are grouped in the coding itself so that the flow of computation is a property of the organization of the code itself, not a brute force imposition. Just as humans have banished GOTO from their engineered code, so has Nature decided that the best programming language of thought is one that is intrinsically organized, with no all-purpose escape mechanisms to let data flow anywhere and run the risk of crashing when compiled in the wetware. Nature—the original software engineer and structured programmer—thus appears to have given some thought to thought.

ACKNOWLEDGMENTS

Earlier versions of this paper were presented at the University of Pennsylvania and at the 1997 Linguistic Society of America meeting in Chicago. I thank the many members of the audiences who gave excellent commentary and criticism. Thanks also go to William Idsardi, Barbara Landau, and Colin Phillips for discussing some of these problems with me.

REFERENCES

ARONOFF, M. 1994. Morphology by Itself. MIT Press, Cambridge, MA.

AVRUTIN, S. 1999. Development of the Syntax-Discourse Interface. Kluwer, Dordrecht.

BACH, E. 1989. Informal Lectures on Formal Semantics. SUNY Press, Albany, NY.

BADDELEY, A. 1997. Working memory. *In* The Cognitive Neurosciences. *Edited by* M. Gazzaniga. MIT Press, Cambridge, MA, pp. 755–764.

BADECKER, W. et al. 1995. The two-stage model of lexical retrieval: Evidence from a case of anomia with selective preservation of grammatical gender. Cognition, **57**:193–216.

BELLUGI, U., A. BIHRLE, D. TRAUNER, T. JERNIGAN, and S. DOHERTY. 1990. Neuropsychological, neurological, and neuroanatomical profile of Williams Syndrome children. American Journal of Medical Genetics, **6**(Suppl):115–125.

BELLUGI, U., P. WANG, and T. JERNIGAN. 1994. Williams syndrome: An unusual neuropsychological profile. *In* Atypical Cognitive Deficits in Developmental Disorders. *Edited by* S. Broman and J. Grafman. Erlbaum, Hillsdale, NJ, pp. 23–56.

BENDER, B., M. LINDEN, and A. ROBINSON. 1989. Verbal and spatial processing efficiency in 32 children with sex chromosome abnormalities. Pediatric Research, **25**:577–579.

BISHOP, D. et al. 2000. Distinctive patterns of memory function in subgroups of females with Turner syndrome: Evidence for imprinted loci on the X-chromosome affecting neurodevelopment. Neuropsychologia, **38**:712–721.

BLANK, M., M. GESSNER, and A. ESPOSITO. 1979. Language without communication: A case study. Journal of Child Language, **6**:329–352.

BOTOLINO, U., and L. LEONARD. 2000. Phonology and children with Specific Language Impairment: Status of structural constraints in two languages. Journal of Communication Disorders, **33**:131–150.

BRAITENBERG, V. 1984. Vehicles: Experiments in Synthetic Psychology. MIT Press, Cambridge, MA.

BROOKS, R. 1991. Intelligence without representation. Artificial Intelligence, **47**:139–159.

CLAHSEN, H., and M. ALMAZAN. 1998. Syntax and morphology in Williams syndrome. Cognition, **68**: 167–198.

COURCHESNE, E. et al. 1994. A new finding: impairment in shifting attention in autistic and cerebellar patients. *In* Atypical Cognitive Deficits in Developmental Disorders. *Edited by* S. Broman and J. Grafman. Erlbaum, Hillsdale, NJ, pp. 101–137.

CRAGO, M., and S. ALLEN. 1996. Building the case for impairment in linguistic representation. *In* Toward a Genetics of Language. *Edited by* M. Rice. Erlbaum, Hillsdale, NJ, pp. 261–289.

CROMER, R. 1981. Developmental language disorders: Cognitive processes, semantics, pragmatics, phonology, and syntax. Journal of Autism and Developmental Disorders, **11**:57–74.

CROMER, R. 1994. A case study of dissociations between language and cognition. *In* Constraints on Language Acquisition: Studies of Atypical Children. *Edited by* H. Tager-Flusberg. Erlabaum, Hillsdale, NJ, pp. 141–153.

CULL, C., and M. WYKE. 1984. Memory function of children with spina bifida and shunted hydrocephalus. Developmental Medicine and Child Neurology, **26**:177–183.

CURTISS, S. 1981. Dissociations between language and cognition: Cases and implications. Journal of Autism and Developmental Disorders, **11**:15–30.

DAMASIO, A. 1994. Descartes' Error. Avon Books, New York, NY.

DENNETT, D. 1996. Kinds of Minds. Basic Books, New York, NY.

DENNIS, M. et al. 1987. Language of hydrocephalic children and adolescents. Journal of Clinical and Experimental Neuropsychology, **9**:593–621.

DENNIS, M. et al. 1999. Congenital hydrocephalus as a model of neurodevelopmental disorder. *In* Neurodevelopmental Disorders. *Edited by* H. Tager-Flusberg. MIT Press, Cambridge, MA, pp. 505–532.

DIENES, Z. and J. PERNER. 1999. A theory of implicit and explicit knowledge. Behavioral and Brain Sciences, **22**:735–755.

DISCIULLO, A-M., ed. 1997. Projections and Interface Conditions. MIT Press, Cambridge, MA.

ELMAN, J. et al. 1996. Rethinking Innateness: A Connectionist Perspective on Development. MIT Press, Cambridge, MA.

ERTESCHIK-SHIR, N. 1998. The syntax-focus structure interface. *In* The Limits of Syntax. *Edited by* P. Culicover and L. McNally. Academic Press, San Diego, CA, pp. 211–240.

FERRETTI, R., and A. CAVALIER. 1991. Constraints on the problem solving of persons with mental retardation. International Review of Research on Mental Retardation, **17**:153–191.

FISCHER, A., and F. GRODZINSKY. 1993. The anatomy of programming languages. Prentice Hall, Englewood Cliffs, NJ.

FODOR, J. 1983. The Modularity of Mind. MIT Press, Cambridge, MA.

FRANGISKAKIS, J. M. et al. 1996. Lim-kinase hemizygosity implicated in impaired visuospatial constructive cognition. Cell, **86**:59–69.

FRAWLEY, W. 1997. Vygotsky and Cognitive Science: Language and the Unification of the Social and Computational Mind. Harvard University Press, Cambridge, MA.

GARMAN, M. 1990. Psycholinguistics. Cambridge University Press, Cambridge, UK.

GHEZZI, C., and M. JAZAYERI. 1987. Programming Language Concepts. Wiley, New York, NY.

GOPNIK, A. 1990. Feature blindness: A case study. Language Acquisition, **1**:139–164.

GOPNIK, M., and M. CRAGO. 1991. Familial aggregation of a developmental language disorder. Cognition, **39**:1–50.

HADLEY, P., and M. RICE. 1996. Emergent uses of *be* and *do*: Evidence from children with specific language impairment. Language Acquisition, **5**:209–243.

HALFORD, G., W. WILSON, and S. PHILLIPS. 1998. Processing capacity defined by relational complexity: implications for comparative, developmental, and cognitive psychology. Behavioral and Brain Sciences, **21**:803–864.

HAYES-ROTH, B. 1985. A blackboard architecture for control. Artificial Intelligence, **26**:251–321.

HEIM, I., and A. KRATZER. 1998. Semantics in Generative Grammar. Blackwell, Cambridge, UK.

JACKENDOFF, R. 1990. Semantic Structures. MIT Press, Cambridge, MA.

JACKENDOFF, R. 1997. The Architecture of the Language Faculty. MIT Press, Cambridge, MA.

JAKUBOWICZ, C., and M. GOLDBLUM. 1995. Processing of number and gender inflections by French-speaking aphasics. Brain and Language, **51**:242–268.

JAREMA, G., and A. FRIEDERICI. 1994. Processing articles and pronouns in agrammatic aphasia: Evidence from French. Brain and Language, **46**:683–694.

KACZMAREK, B. 1987. Regulatory functions of the frontal lobes: A neurolinguistic perspective. *In* The Frontal Lobes Revisited. *Edited by* E. Perecman. Erlbaum, Hillsdale, NJ, pp. 225–240.

KAGER, R., H. VAN DER HULST, and W. ZONNEVELD, eds. 1999. The Prosody-Morphology Interface. Cambridge University Press, Cambridge, UK.

KAPLAN, J., H. BROWNELL, J. JACOBS, and H. GARDNER. 1990. The effects of right hemisphere damage on the pragmatic interpretation of conversational remarks. Brain and Language, **38**:315–333.

KARMILOFF-SMITH, A. 1998. Development itself is the key to understanding developmental disorders. Trends in Cognitive Sciences, **2**:389–398.

KARMILOFF-SMITH, A. et al. 1995. Is there a social module? Journal of Cognitive Neuroscience, **7**:196–208.

KARMILOFF-SMITH, A. et al. 1997. Language and Williams syndrome: How intact is “intact”? Child Development, **68**:246–262.

KARMILOFF-SMITH, A. et al. 1998. Linguistic dissociation in Williams syndrome: Evaluating receptive syntax in on-line and off-line tasks. Neuropsychologica, **36**:343–351.

KOWALSKI, R. 1979a. Algorithm = logic + control. CACM, **22**:425–436.

KOWALSKI, R. 1979b. Logic for Problem Solving. North Holland, Amsterdam.

LEHNOFF, H. et al. 1997. Williams syndrome and the brain. Scientific American, **December**:68–73.

LEONARD, L. 1989. Language learnability and specific language impairment. Applied Psycholinguistics, **10**:179–202.

LEONARD, L. 1998. Children with Specific Language Impairment. MIT Press, Cambridge, MA.

LEVY, Y. 1996. Modularity of language reconsidered. Brain and Language, **55**:240–263.

MACDONALD, G. and D. ROY. 1988. Williams syndrome: A neuropsychological profile. Journal of Clinical and Experimental Neuropsychology, **10**:125–131.

MCCAULEY, E., T. KAY, J. ITO, and R. TREDER. 1987. The Turner syndrome: cognitive deficits, affective discrimination, and behavior problems. Child Development, **58**:464–473.

MCCLAMROCK, R. 1995. Existential Cognition. University of Chicago Press, Chicago, IL.

MCGLONE, J. 1985. Can spatial deficits in Turner's syndrome be explained by focal CNS dysfunction or atypical speech lateralization? Journal of Clinical & Experimental Neuropsychology, **7**:375–394.

MERVIS, C., et al. 1999. Williams syndrome: Findings from an integrated program of research. *In* Neurodevelopmental Disorders. *Edited by* H. Tager-Flusberg. MIT Press, Cambridge, MA, pp. 65–110.

MORROW, J., and T. WACHS. 1992. Infants with myelomeningocele: visual recognition memory and sensorimotor abilities. Developmental Medicine and Child Neurology, **34**:488–498.

NEVILLE, H., D. MILLS, and U. BELLUGI. 1994. Effects of altered auditory sensitivity and age of language acquisition on the development of language-relevant neural systems: Preliminary studies of Williams syndrome. *In* Atypical Cognitive Deficits in Developmental Disorders. *Edited by* S. Broman and J. Grafman. Erlbaum, Hillsdale, NJ, pp. 67–83.

PANI, J., C. MERVIS, and B. ROBINSON. 1999. Global spatial organization by individuals with Williams syndrome. Psychological Science, **10**:453–458.

PORT, R., and T. VAN GELDER, eds. 1995. Mind as Motion: Explorations in the Dynamics of Cognition. MIT Press, Cambridge, MA.

PYLYSHYN, Z. 1984. Computation and Cognition. MIT Press, Cambridge, MA.

REIGEL, D. 1993. Spina bifida from infancy through the school years. *In* Teaching the Student with Spina Bifida. *Edited by* F. Rowley-Kelly and D. Reigel. Paul H. Brookes, Baltimore, MD, pp. 3–30.

ROURKE, B. 1988. The syndrome of nonverbal learning disabilities. The Clinical Neuropsychologist, **2**: 293–330.

SCHNEIDER, W., and W. OLIVER. 1991. An instructable connectionist/control architecture: Using rule-based interactions to accomplish connectionist learning in a human time scale. *In* Architectures for Intelligence. *Edited by* K. van Lehn. Erlbaum, Hillsdale, NJ, pp. 113–145.

SCHOPLER, E. 1994. Neurobiologic correlates in the classification and study of autism. *In* Atypical Cognitive Deficits in Developmental Disorders. *Edited by* S. Broman and J. Grafman. Erlbaum, Hillsdale, NJ, pp. 87–100.

SIGMAN, M. 1994. What are the core deficits of autism? *In* Atypical Cognitive Deficits in Developmental Disorders. *Edited by* S. Broman and J. Grafman. Erlbaum, Hillsdale, NJ, pp. 139–157.

SILBERT, A., et al. 1977. Spatial and temporal processing in patients with Turner's syndrome. Behavior Genetics, **7**:11–21.

SLOMAN, A. 1993. The mind as a control system. *In* Philosophy and Cognitive Science. *Edited by* C. Hookway and D. Peterson. Cambridge University Press, Cambridge, UK, pp. 69–110.

SMITH, N., and I-M. TSIMPLI. 1995. The Mind of a Savant. Blackwell, Oxford, UK.

SOKOLIK, M., and M. SMITH. 1992. Agreement of gender to French nouns in primary and secondary language: A connectionist model. Second Language Research, **8**:39–58.

STEVENS, T., and A. KARMILOFF-SMITH. 1997. Word learning in a special population: Do individuals with Williams syndrome obey lexical constraints? Journal of Child Language, **24**:737–765.

SUCHMAN, L. 1987. Plans and Situated Action. Cambridge University Press, Cambridge, UK.

SULLIVAN, K., and H. TAGER-FLUSBERG. 1999. Second order belief attribution in Williams syndrome: Intact or impaired? American Journal of Mental Retardation, **104**:523–532.

TAGER-FLUSBERG, H. 1991. Semantic processing in the free recall of autistic children: Further evidence for a cognitive deficit. British Journal of Developmental Psychology, **9**:417–431.

TAGER-FLUSBERG, H. 1994. Dissociations in form and function in the acquisition of language by autistic children. *In* Constraints on Language Acquisition: Studies of Atypical Children. *Edited by* H. Tager-Flusberg. Erlbaum, Hillsdale, NJ, pp. 175–194.

TEMPLE, C. M., and R. A. CARNEY. 1993. Intellectual functioning of children with Turner syndrome: a comparison of behavioral phenotypes. Developmental Medicine and Child Neurology, **35**:691–698.

TEUFEL, B. 1991. Organization of Programming Languages. Springer, New York, NY.

THAGARD, P. 1996. Mind. MIT Press, Cambridge, MA.

THAL, D., E. BATES, and U. BELLUGI. 1989. Language and cognition in two children with Williams syndrome. Journal of Speech and Hearing Research, **32**:489–500.

THELEN, E., and L. SMITH. 1994. A Dynamic Systems Approach to the Development of Cognition and Action. MIT Press, Cambridge, MA.

TUCKER, G. R., et al. 1968. A psychological investigation of French speakers' skill with grammatical gender. Journal of Verbal Learning and Verbal Behavior, **7**:312–316.

TYLER, L., et al. 1997. Do people with WS have bizarre semantics? A primed monitoring study. Cortex, **33**:515–527.

UDWIN, O. 1990. A survey of adults with Williams syndrome and idiopathic infantile hypercalcemia. Developmental Medicine and Child Neurology, **32**:129–141.

UDWIN, O., W. YULE, and N. MARTIN. 1987. Cognitive abilities and behavioural characteristics of children with idiopathic infantile hypercalcemia. Journal of Child Psychology and Psychiatry, **28**:297–309.

UDWIN, O., and W. YULE. 1991. A cognitive and behavioural phenotype in Williams syndrome. Journal of Clinical and Experimental Neurology, **13**:232–244.

VAN DER LELY, H. 1994. Canonical linking rules: forward versus reverse linking in normally developing and specifically language-impaired children. Cognition, **51**:29–72.

VAN GELDER, T. 1998. The dynamical hypothesis in cognitive science. Behavioral and Brain Sciences, **21**:615–665.

VUCHIC, R. 1993. A study of noun phrase agreement in French as a second language: An autosegmental model. Ph.D. Dissertation, University of Delaware.

WHITE, B. 1994. The Turner syndrome: origin, cytogenetic variants, and factors influencing the phenotype. *In* Atypical Cognitive Deficits in Developmental Disorders. *Edited by* S. Broman and J. Grafman. Erlbaum, Hillsdale, NJ, pp. 183–195.

YAMADA, J. 1990. Laura: A Case for the Modularity of Language. MIT Press, Cambridge, MA.

YAMADA, J., and S. CURTISS. 1981. The relationship between language and cognition in a case of Turner's syndrome. UCLA Working Papers in Cognitive Linguistics, **3**:93–115.

ZHANG, J., and D. NORMAN. 1994. Representations in distributed cognitive tasks. Cognitive Science, **18**: 87–122.