

MENTAL COMPUTATION AND LANGUAGE BREAKDOWN: CLARIFICATIONS, EXTENSIONS, AND RESPONSES

WILLIAM FRAWLEY

Department of Linguistics and Cognitive Science, University of Delaware

This paper is a response to commentaries on my target paper for *Computational Intelligence*, “Control and Cross-Domain Mental Computation: Evidence from Language Breakdown.” In this response, I acknowledge certain errors in my initial construal of control and dismiss unwarranted criticisms. I then reexamine both control and certain language disorders in light of the explicitness of cross-domain communication and the visibility of representations to each other. In the end, I reassert the validity of the logic/control (visibility) distinction in mental computation and argue that the contrasts between Specific Language Impairment and Williams syndrome parallel this distinction.

Key words: explicit/implicit knowledge, representation, visibility, language disorders.

In my target paper for this special issue, I make the following arguments. The distinction in computation between logic and control is relevant to mental computation, and empirical evidence for this difference can be found in the way two language disorders affect linguistic behavior. The language difficulties associated with Williams syndrome appear to be cross-domain and hence problems with control. In contrast, the disorders associated with Specific Language Impairment appear to be within-domain and hence problems of logic. These differences come through especially clearly in an analysis of morphological theory and experimental performance. Language loss thus cuts broadly across mental computation—disorders that affect within-domain data structures (logic) and those that affect cross-domain information management (control). The latter disorders, moreover, are not generalized performance difficulties but specific to the control aspect of mental computation. These claims dovetail with a variety of recent advances in cognitive science and linguistics—e.g., proposals about the explicitness of cross-domain information and the implicitness of within-domain information and theoretical arguments about the organization and interaction of linguistic modules.

The careful and thoughtful responses of the commentators—even the critical ones, as I hope to show—suggest to me that the foregoing arguments withstand scrutiny. Indeed, I think the claims are strengthened by a number of corrections and clarifications that emerge in the commentaries. My responses fall into four categories: (1) acknowledgment of errors and direct responses to what I see as misguided criticisms; (2) modifications of and progress on the idea of control; (3) modifications of and progress on accounts of language disorders; and (4) incommensurables—points where the commentators and I are at potentially irresolvable odds—though I try to reach out on these points for common ground.

1. PLEAS AND DENIALS

1.1. Errors

Commentators point out three errors in the paper that I must acknowledge (although *nolo contendere* is perhaps the correct plea). First, Cooper, Chandrasekaran, and, to a certain

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

extent, Kovordányi argue that my view of control is too narrow (and, paradoxically, too literal (Cooper) and too metaphorical (Chandrasekaran)). Borrowed from classic, symbolic programming languages, my view misses the variety of other control proposals in the computational and cognitive science literature. I think this criticism is correct. In my defense, let me say that I chose this version of control because I thought it was reasonably well understood in the computational literature and thus could provide common ground for discussion.

Cooper's, Chandrasekaran's, and Kovordányi's criticisms show that control must be understood as a version of the more general computational issues of communication and coordination, which, in turn, manifest themselves differently depending on the architecture and system-internal ways of drawing the data/management line. Certainly Soar and ACT are very much built on considerations of dataflow, and lessons from these models on how to manage interfaces and representational domains need to be appreciated more by modularity-adherent linguists like myself. (Linguists in general need to listen more to those who know how running systems and real minds work—and real languages for that matter!—but this is a thorny and, perhaps, irreconcilable issue in the sociopolitics of the field.) As Cooper and Chandrasekaran observe, neither Soar nor ACT is in principle at odds with modularity.

Kovordányi's point that my view of control fails for within-domain processing is a related criticism, but less clear to me as one that is on target. The only neurally plausible architecture—neural nets—lacks control (at least within-domain) because of cascading and hence no resource management or monitoring. Must there always be resource management and monitoring in control? The connection of control to these executive functions depends on how control is implemented (as Chandrasekaran and Perner and Dienes note). It may be implicit in the data structures affected, hence forcing representations to read each other directly across domains, or it may be explicit, hence forcing representations to communicate via an independent (and thus resource-intensive) mechanism. I explore this difference more fully below in reconsiderations of control, but I want to point out that neural nets do not in principle eliminate control issues. Something like the former idea—direct reading of one representation by another—may be lurking in neural nets. Ruling out data management in neural nets conflates instantaneousness with simultaneity: if mental events are instantaneous (which they cannot be in principle and so this remains a metaphorical way of talking), there is indeed no dataflow issue; but if mental events are simultaneous (which is the essence of parallel processing), then control issues are not ruled out. I see and appreciate Kovordányi's criticism that control applies only to high-level domain interactions, where such things as explicit sequencing are relevant—e.g., in the connection of language to motor output in speech or in the link of spatial representation to motor output in grasping. But I remain suspicious about applying shudder quotes to within-domain “control” because, as I argue below, submodular control is not “just processing” and different views of the coding of control clarify both mental computation and the proper account of language disorders. (See the papers in Monsell and Driver (2000) for more detail on many of these issues, especially high-level and low-level control, processing costs, and control in connectionist networks.)

A second charge, related to the first, is that I am too wedded to the Turing model and symbolic computing. Cooper, Chandrasekaran, and Kovordányi make this point directly, as do Thomas and Stanlaw, though more obliquely, with Thomas committed to something like dynamical modeling and Stanlaw, coming from a contextualist position on mind, leery of any computational mentalese. I am supportive of symbolic computation because I think it is absolutely required for an explanation of the mind. But I acknowledge that any suggestion

on my part that the Turing world is the only world is misguided. Again, in my defense, I have to say that I have pushed rhetorically for an argument grounded in what I thought was in common consciousness, but this may have incorrectly allowed the inference that symbolic computing is the only game in town, which it is not.

Still, symbolic computing *is* one of the games in town and, I believe, a necessary one. There are some aspects of mental computation that cannot be captured without symbolic, rule-like representations (however they might be implemented). Statistical, associationistic computation is not—and cannot be—the only knowledge mechanism. I do not want to rehearse the furious debate in cognitive science over this matter, but only re-assert that we need symbolic computation—or, better, that statistical computation is insufficient—for many well known, robust phenomena: e.g., in language alone, the detection of phonological structure above the segment (see Marcus (2001) for a thorough critique, showing, for example, that networks can never freely generalize in the way necessary for language learning).

Another way to put this is to say that in the real world, not the technological computational one, where programmers deliberately enrich the code, data do not enter the system as types, but as tokens. No one hears a noun, and NP is not some second- or third-order distribution over transitional probabilities. Even those who otherwise advance statistical computation in other areas of language acknowledge its limits when applied to anything more than simple sequences (Aslin, Saffran, and Newport 1998, p. 324). Statistical distributions occur over tokens. How does the mental system construct non-statistical types from token input? Or better, how does the system use statistically distributed tokens as evidence for types?¹ Why does the mental system operate independently of the statistical token-bath in which it lives, producing things it never hears, not producing things it hears all the time, disregarding high frequency tokens as evidence for certain types, and so on (see Cowie 1999)? I am not saying that an associationistic system cannot deliver something like symbolic types—only that if it does, it must evince symbolic behavior and must create variables. And we should also admit that a symbolic system can evince probabilistic behavior—modules of symbolic rules could probabilistically interact, for example. In my view, the best mental architecture is a hybrid one (e.g., Sun, Merrill, and Peterson 2001) because it captures both explicit and implicit knowledge (and raises all sorts of interesting control issues).

This criticism of my over-attachment to the Turing model actually uncovers a more central issue for cognitive science raised by the entire idea of mental dataflow, storage, management, and breakdown. A foundational problem in the study of mental computation is not whether you are Turing-blinded, statisticophilic, or chaotic-dynamic, but how you get and coordinate *as effects of the neural wetware* what both associationism and symbolism stand for: large-scale memory-intensive, brute force storage of nonanalyzed data connected by probabilistic hot links; computationally-intensive running of deterministic rules over a small set of stored, but analyzed, abstractions. Both are required for morphology, the crucial area of data for the arguments of the target paper (Pinker 1999). Indeed, Aronoff's comments show that the right analysis of the complex morphological structure of a language requires that

¹The brain seems to separate these two effects. The syntactic behavior of split-brain patients, for example, reveals that the right hemisphere is attuned to statistical properties of small whole-sentence tokens in a distribution while the left hemisphere is sensitive to analytical types (Gazzaniga, Ivry, and Mangun 1998, p. 335; Hauser 1997, pp. 150–151 reports startlingly similar findings for zebra finches). Thus one might say that the brain collects statistical information about sentences, as it does for all sorts of input (thanks to William Idsardi for this observation), but it is not clear how this statistical information is used, if at all, for type-construction.

the system have procedures that operate over abstract morphological classes—declensions—and that these abstract types be the locale of the control processes in morphology. Declensions, it seems safe to say, are not properties of the input data.

To put it another way, given what the neural architecture is and the tasks it needs to perform, how do you get these two (and other) classes of behavior as consequences? Neither symbolism nor associationism (nor dynamism) is faithful to the neural hardware (contrary to connectionist evangelizing, biology does not speak self-evidently through neural nets: Aizawa 1992). These are models of the knowledge level and make predictions at that level. Some models rely more on linkages to the implementation level, some on such linkages to the ecological level. But the mere invocation of a model as self-evident because it looks like a network is at best just a rhetorical strategy.

At the risk of speaking functionalism and thus calling down further aspersions on myself, I want to underscore a definitional point. *What the mental computation system is* does not have to be isomorphic with *what the mental computation system does*. If it turns out that mind *is* a chaotic system, then we still must explain how it *does* symbolic computation for some parts of morphology (and other domains). If it *is* a symbolic system, then we must explain how it *does* probabilistic association for certain aspects of morphology (and other domains).

What does this have to do with control in mental computation? Another effect of the wetware is control. If the mind is a statistical learning device or if mentalese is something like object-oriented programming, we must explain how it does dataflow management (a point I think Cooper nicely articulates). Consequently, even though I am admittedly stuck in a Turing world, and thus open to criticism for it, I am happy here because the issues arising with classic, functional and procedural programming languages simply do not go away just because other approaches to computing capture certain things about the mind. There must be symbolic effects (as well as associationistic effects). Control lives—but with a much more subtle identity than I have portrayed. (See Monsell and Driver (2000) on the functional realization of control vs. its explicit realization.)

The third charge of error has to do with use of the literature. I was very glad to learn of other studies of control heretofore unknown to me: in particular Cooper's work on object-oriented programming and Chandrasekaran's useful work on generalizing control at the knowledge level. These are important advances, and more on them below. But I obviously have no access to the many papers cited that are in press or forthcoming. In the interest of not delaying this commentary further, I have chosen not to secure those papers, read them, and include them in this response. Readers should know the time course of this paper and the special issue to get a sense of why I am worried about delay. The paper has been in circulation for review since the fall of 1999 and then slightly updated when it was to be sent to commentators; final comments did not arrive to me until the summer of 2001, and I wrote the present response in August, 2001. Thus there is some criticism—e.g., Thomas's arguments that Williams children have a within-domain disorder for lexical representation and a syntactic deficit—and some support—e.g., Perner and Dienes's distinction between vehicle control and content control—that I simply cannot address in any detail here. I hope that in a future issue of *Computational Intelligence*, we can continue discussion of these latest ideas in order to further the entire argument.

1.2. Rejections

There are also three criticisms that I think are misguided. Two of these come very close to being incommensurables, so I will try to walk the line between counter and dismissal.

1.2.1. *Development.* Thomas raises what might be called *foundational objections* that development—not static mental-computational structure—is the true cause of the language disorders I discuss; his remarks dovetail with several of Stanlaw’s, so I treat them together.

Thomas’s argument runs like this. Williams syndrome and SLI are developmental disorders in that they arise as a consequence of the trajectory of cognitive development—with all its interacting variables in the learning process—not from damage to some pre-programmed information. This is because Nature has made humans as learning systems, which must go through a process of development from initial states in order to achieve final adult knowledge states. The structure of the final states looks quite different from that of the initial mental program, making inferences about causes from the initial and intermediate states quite limited, unless you have a full developmental theory. As a consequence, the best language for explaining developmental disorders is that of machine learning, or an account of developmental learning devices, in particular some kind of dynamic modeling because that most accurately captures both the intrinsic developmentalism of the disorders and Nature’s arrangement of the ecology.

In some ways I agree completely with Thomas but fail to see the impact of what he has to say. I do know that he and I differ on what it means to learn or have damage or infer a cause. These differences make our conversations borderline incommensurable, but let me try to respond.

Thomas criticizes me for using developmental disorders as the source of evidence for the logic/control claim because the conditions are dynamic, and the intermediate developmental states at which I am sampling will be unrecognizable from the final states. I would be better off, he says, using data from adult aphasias, where, presumably, the logic/control issue is settled and so the inferences would be more stable.

It is easy to turn this argument on its head. If developmentalism were correct, it would be *best* to study children or others at intermediate states of development because in these individuals, the structures in question can be more easily teased apart. In the adult, the full integration and restructuring of the initial mental programs has taken place, thus forever rendering impossible the discovery of how particular components contribute to the disordered behavior. In adult aphasias, the logic-control distinction has become blurred beyond recognition, thus making adults the *worst* candidates for inference about causes. In children, who are going through the process of learning to manage and integrate logic and control, you can still see the independent effects.

I have to confess that I do not believe any of the foregoing because I think that careful examination of the data can reveal the phenomena in both adults and children. But I want to underscore the potentially lethal precariousness of adherence to developmentalism. If the mental programs are in continual update and integration, you have no guarantee of a clean look anywhere at any time. Van der Lely says: “Logically, if a mechanism is impaired then it will impinge on the representations that develop from or are served by this mechanism.” In such a case, the earlier the look at that impaired mechanism, the better.

Take a noncognitive developmental disorder, like progeria (where children, because of nonfamilial genetic mutation, age rapidly). Certainly their initial states undergo substantial change to their final states in the trajectory of their development. Does this development force us to wait until they die to infer causes? Or do we say something like we would say as developmentalists about Williams syndrome: that what causes the language disorders in Williams syndrome is “just language delay” (a consequence of the developmental trajectory)? This seems to me to be like saying that the strokes and heart

disease in a six-year-old with progeria are caused by “just rapid aging.” What causes those things? Causes and consequences seem to me terribly conflated here.

Invoking an entire developmental trajectory and arguing that the final state differs radically from the non-final states risks disregard for the structural conditions on the initial states—as if by adding up 5, 6, and 7, you change the nature of addition by the time you get to 18. Maybe Thomas would reply that, in the case of mental development, the cognitive equivalents of addition do change over the course of development, but this then depends on what you see to be learning and the proper account of machine learning: more on this below.

Selective impairments (a phrase that Thomas seems not to like) do not vitiate effects on the whole system and on developmental trajectory. The selective spatial and cardiac impairments in Williams syndrome—and *they are selective*—affect the performance on general cognitive development and cardiovascular performance precisely because the mind and the body are systems, *où tout se tient* (and, as with the mind as a system itself, much depends on *comment tout se tient*). But none of this has an impact, as far as I can see, on trying to infer particular, discrete factors at any stage of the developmental trajectory. Visual knowledge is developmental. Do edge-detecting mechanisms grow to unrecognizability in adults? Auditory knowledge is developmental. Do the mechanisms subserving the detection of voice onset time become forever murky in adults?

Arguments for developmentalism come precariously close to Stanlaw’s almost fatalistic objection to research on disorders—that you cannot really infer anything about mental structure from a failure. Or, as he says, taking a resistor out of an old radio, thus making it squawk, does not allow the inference that the resistor is an antisquawking device. To this, I just have to say: a resistor *is* an antisquawking device—once you are not fooled by the semantics. Squawk is non-good-sound. With the resistor in, the radio produces good sound (i.e., non-non-good sound (= non-squawk). And so, functionally, the resistor is an antisquawk device, though tortuously analytically so. The morphological component in the mind is arguably an anti-bad-wordform mechanism. SLI produces an anti-anti-bad-wordform mechanism: i.e., a disorder.

This way of arguing does not fall into Stanlaw’s trap that you cannot infer the whole structure of the radio from a component: “but we can’t, in general, tell how a radio works by looking at a broken component, just as we can’t, in general, tell the whole from its parts.” You most certainly can infer *something* about the function of a system from a component. Nature cuts mind-brains at their joints. If we can’t infer structure from the consequent failures of components, then we should be suspicious about inferring structure from successful performance. Maybe there is no structure to infer. If so, we are committed to accounts of one-time behavioral occurrences—of which there can be no science, an epistemological ledge on which, I think, radical developmentalism teeters.

At this point, the whole issue of learning and the proper machine model thereof comes to the fore. This is another furious and extended debate that I do not want to rehearse but simply counter. For Thomas,

Nature does not produce a *pre-programmed* cognitive system but constructs a *learning system*, in which much of the detailed structure emerges through a development process requiring interaction with the physical and social environment. . . . Developmental disorders . . . are more properly considered within the framework of *machine learning*. . . [where] impairments are not the result of damaging specific functional components in a pre-programmed system (whether logic or control); rather they are the result of an atypical trajectory of development.

Well, it all depends how you look at it.

In my view, language and other biologically necessary aspects of the human system are not learned in the sense that Thomas uses the term—actively testing hypotheses against evidence. Language, face knowledge, space knowledge, musical knowledge, etc. *happen to devices like us in this kind of environment*. They are not things we do or construct in some traditional-learning sense. Certainly, the “detailed structure emerges through a development process,” what a linguist would call *parameter-setting*, but the details miss the forest for the trees. We learn language as we learn surfaces or mathematical incrementation. They are triggered, not tested, by the “interaction with the physical and social environment.” They are no more actively constructed in a folk-learning sense—hence rendering their initial state completely unrecognizable in their final state—than is the mental route of the northward journey to the mid-Atlantic states of Red Admiral butterflies who have never made the journey before. Migration is something that happens to these butterfly devices in this kind of environment. Language is learned in the same way that rats “learn” that cats are natural triggers of fear (LeDoux 1996, pp. 143–144). If one still wishes to call this experience of natural triggers *hypothesis testing* because of minimal exposure to data (apologies to Fred Adams for the phrasing of this counter), then that is fine with me, but it looks more like hypothesis confirmation and instantiation than testing (see, e.g., Gallistel 1990; Cowie 1999 presents a skeptical, but still supportive, view).

What does this have to do with control? Well, it all depends how you look at it. Machine learning is generally understood as a triple of computable input operated on by an evaluation procedure that outputs a program that solves the data such that any new input is rendered accountable to the output program. There is a point where learning stops and no truly new data are encountered. A system that goes on reorganizing forever in the face of new data, rather than just tweaking, does not match the facts: when have you heard a brand new phoneme or seen a brand new edge? The “detailed structure emerges [continually] through a development process,” and so there are new tokens, but after a point—early in development I would say—there are no new types. The more that is built into one part of the machine-learning triple (or as Thomas says, “Every inductive learning system must incorporate biases that constrain the hypotheses”), the less that is built into others. So if the input had all the structure it needed, the system would minimize evaluation and program output and would thus be an absorption device (not a realistic picture). If the data were terribly impoverished and the system had to construct everything, it would have to be over-pre-programmed and would in effect learn nothing (also terribly unlikely).

Here is a picture of (human) machine learning. The human device receives computable input (token inputs that impinge on the device and the device recognizes because of what it is built to recognize) and its evaluation procedure comes pre-set with at least two biases—data structures and management. (It is entirely plausible that control would be a part of the biological machinery, given that the brain has a massive representational coordination problem. See the next section.) The system then constructs program outputs for recognizable inputs following at least two “commands”: organize the information and coordinate/communicate it (not the same as process it). (See below.) A disruption in aspects of this pre-set evaluation procedure for certain kinds of recognizable inputs is one cause of disorders like Williams syndrome and SLI. There are certainly developmental effects, and the kind of “learning” that ensues from this disruption of the evaluation procedure does lead to program output that looks very different from the initial state. But this is because particular aspects of the pre-set evaluation procedure are defective.

Thomas claims that he gets success in his own computational modeling of these disorders by manipulating the “phonological representations, involving a reduction in

similarity and redundancy.” How does the system know what is similar and what is redundant if not by a biased evaluation procedure? This is a deficit in the system’s pre-programming, even if a limited one—a well-known and unavoidable criticism of every associationist model (the tyranny of analogy—similar on what basis? redundant on what basis?). None of this “ignore[s] development as a [my emphasis, WF] cause of the pattern of impairments in the developmental disorder.” But why also ignore the effects of what Nature and the connectionist programmer have done to the initial state? *Control is something that happens to devices like humans when they are put in these kinds of environments.*

1.2.2. Lumping. Thomas claims that in considering disorders, I have lumped too many things together and so missed important details. He says, “it is not clear to me what is to be gained by throwing such a range of disparate developmental disorders into a single bag (WS, Turner syndrome, spina bifida with hydrocephalus, autism) when their underlying causes may be quite different.” The reason for lumping them together is that at the knowledge level, they exhibit similar effects. The value of linking ostensibly disparate disorders at some level of analysis is that it may provide a unified explanation and, in the case of disorders, potentially new ways of doing clinical interventions.² Functionalism again speaks up. All these disorders do have different causes, but functionally share a control deficit. Must we limit our explanation of every disorder by its idiosyncratic link to its cause (presuming we even know the cause)? Or is such narrowness the legacy (obligation?) of eliminativism? Are there no functional classes of disorders? Is anything to be gained by seeing mental phenomena as functionally convergent?

Consider an illustration from a different area of cognition: emotion. There is a disorder of emotional processing called *alexithymia* that is very much the emotional equivalent of blindsight. Alexithymics give physiological evidence of processing emotions, but exhibit little or no behavioral output of emotions, and they insist that they do not experience emotions on a conscious level. Interestingly enough, alexithymia can have both a clear physical cause (a lesion in control pathways between the lower level, emotional processing mechanisms and the anterior cingulate gyrus: Lane, Ahern, Schwartz, and Kazniak 1997) and non-evident physical cause (e.g., psychic trauma, which surely leaves its own physical damage: Taylor, Bagby, and Parker 1997). The functional consequences of each kind of alexithymia are astonishingly similar: bland affect, failure to recognize emotional stimuli, denial of emotional experience. At the computational-knowledge level, one can develop a unified account. They both are traceable to disruptions in the way the submodules of emotional processing report out the results of their computations in explicit form (Frawley and Smith 2001). As a consequence of this unified explanation, we can see more subtle types of alexithymia and develop recommendations for intervention.

The value of lumping is the hope that if there is a split between disorders of logic and disorders of coordination, we can find new and different subtypes of these deficits, different assessment batteries, and different intervention techniques. This seems to be the right research goal.

²When I was working on Williams syndrome some years ago, I remarked to the M.D. overseeing some of the cases that Williams and Turner kids both had specific spatial integration problems. He said he had never noticed that and that, indeed, it was true and this might lead to a different consideration of clinical practice. Williams and Turner have very different causes, but remarkably similar functional effects. Should they therefore have completely different interventions at the knowledge and behavioral levels?

1.2.3. *The Right Language.* In several commentaries, I am criticized for not using the right explanatory language or for using a language that implies more than it should or is vague when pushed and so fails to lead to proper inferences. The first such complaint concerns the biological plausibility of the computational-control vocabulary. Cooper wonders about this directly, and it is certainly on the minds of both Thomas and Kovordányi. Cooper says: “[this way of talking] suggests that the neural hardware rigidly adheres to a predetermined instruction sequence. (Although that instruction sequence may contain conditional ‘if-then-else’ statements, it remains a rigid instruction sequence.) Where does this sequence come from?” This worry is addressed by Kovordányi and Perner and Dienes, who point out that control across clearly separate neural domains needs coordination such as an instruction sequence because the outputs have to unfold in real-time. Kovordányi says, “somewhere along the chain of processing from perceptual input to motor output, the parallel flow of information must be sequenced. In other words, a series of choices has to be made regarding which part of information should be used for the production of output.” Sequencing requires explicit coordination and hence something like what Perner and Dienes call *content control*.

Cooper also raises the following concern: “[this way of talking] begs the question of how an operating system complex enough to perform the necessary instruction following and allocation of cognitive and effective resources might have evolved within the neural hardware.” It is pretty clear that such mechanisms have evolved. A major evolutionary advance is the development of a “central holding facility” (Bownds 1999, p. 35) that mediates sensory stimulation and response, which requires monitors to control delivery of stored representations for a task (see also Carstairs-McCarthy 1999, citing Calvin’s work in a similar vein, to support the development of managerial function in the evolution of linguistic capacity). Thus, we see in the mammal brain a variety of monitoring and regulation locales (Gazzaniga, Ivry, and Mangun 1998, pp. 267–283; 423–464). Deacon (1997, p. 158) argues that the human brain had to evolve managerial structures to oversee the coordination of cortical areas taken over for new tasks as brain and language co-evolved: “We can consider this managerial work to be analogous to the ‘higher functions’ of the brain, those functions that direct, coordinate, and monitor input-output functions.” (See Monsell and Driver (2000) for many neural realizations of control.)

The reverse of this criticism is whether purely biological systems that are apparently using control are thereby computing. Chandrasekaran observes:

Certainly, biological systems are replete with biochemical processes in which one subsystem dumps some chemical in some place which controls the behavior of another subsystem, or even cases where more than two subsystems coordinate their behaviors in this manner. Many of the issues of control are similar to the issues on intermodular control in software or cognition. We can ask, in the case of a malfunction, whether the problem was due to subsystem 1 failing to generate the chemical, deposit the chemical, subsystem 2 absorbing the chemical, or it failing to make use of the chemical, exactly similar to the author’s analysis of language disorders. But there is no automatic need to think of the chemical processes as ‘programming language of the body,’ or that biological systems are just executing algorithms.

Here it is important to consider what it means to control and the kind of system to which control applies, and Chandrasekaran’s (1994) own work is illuminating in this regard. He offers a hierarchy of controllers, from fast, limited-condition control that is relatively easily analyzed in terms of subtasks and plans (e.g., control in a thermostat) to slow, multi-condition control with complex subtasks and plans (control in *deliberative engines*).

Essentially he is constructing a gradient of representation-using systems and their attendant control requirements. Thermostats are not representation-using: they act because they use indications, not representations. Thermostats' insides always directly reflect the facts. Because they have no false representations, their control processes are simplified and one might think of their dataflow (if, indeed, they have data) as direct transfer. But humans can have false representations: a fact in itself suggesting the need for monitoring (Is what I am thinking correct? What should I expect? See Dennett 1991.) Hence, control in deliberative engines such as humans requires analysis via complex subtasks and plans.

The important issue that emerges from this consideration of the right language for biological systems is that control in the sense of domain management may be relevant for only representation-using devices. Only a very intelligent blender would be passing representations of food to its blades for a proper liquefying subroutine—and monitoring this dataflow. Mere uptake of the output of some process by another does not require control. I would thus argue that in the chemical example Chandrasekaran gives, there are no representations, so control in the sense I am using the term does not apply.

The second correct-language complaint concerns the cognitive-behavioral plausibility of the control vocabulary. Chandrasekaran wonders what is gotten cognitively by the proliferation of (to him, often arbitrary) control types—“it was not at all clear to me that distinctions such as ‘machine-level control’ versus ‘high-level control,’ not to mention ‘statement-level control,’ transcend the specific computational framework”—and the extent to which we can push the computational vocabulary—“Of course I make mistakes in my thinking, but crashes? What are they? I do crash when I have had too much to drink or too tired, but I am not aware of cognitive behavior that corresponds to crashing in the same sense in which we use for computer hardware or software.”

In some ways, this criticism is the opposite of the lumping complaint. It is the splitting complaint—too many terms. I am not suggesting, as Chandrasekaran seems to imply (“When we use terms like ‘You seem pretty lost in solving this problem, let us reboot,’ this is a metaphor”) that we infuse our explanations with computational phraseology for its own sake. Even though I can easily imagine this strategy supporting a successful advertising campaign for a new anti-anxiety drug—*Got a glitch? Bug in your interface? Reboot with Felixia*—nothing is gotten by gratuitous technical talk: *Wow! I am so glad to have bound my variables!* But there is value in using technical ideas as what Dennett would call *an intuition pump*: see if the technical distinctions that spin off from a larger one apply to the cases in question, and this will push the inquiry by affording both correct turns and instructive wrong ones. (See the computational papers in Monsell and Driver (2000) for more on the use and limits of computational vocabulary.)

There are different ways of implementing control, given the architecture and the language, and these place different demands on memory and resources. Some control strategies are more or less explicit, and some manage dataflow by partially linking or fully linking representations across domains. Do these distinctions apply cognitively? I think some do (see more below). For example, the idea of specifying data and operations together in object-oriented programming may be a computational way of talking about what Perner and Dienes call *vehicle control*. But no one is then requiring that we flesh out all of object-oriented programming language in cognitive terms: *so where's the iterator?*

When a program (ahem. . .) crashes, the system recovers in a small number of ways and, indeed, these recovery types seem to be instantiated in humans. Indeed, if this way of talking “raises the possibility of both program and operating system crashes” (Cooper), then you might in fact find disorders like dementia to be more like operating system (ahem. . .) crashes, compared to, say, amusia, which does look like an application-specific (ahem. . .) glitch.

1.3. Lessons Learned

These pleas and denials raise a number of questions for of continuing discussion:

- **Control**
What defines control? Must all control require resource management or be memory-intensive (Monsell and Driver 2000)? Are answers to these questions illuminated by a distinction between clear, cross-domain coordination (such as in language-to-speech output) and submodular coordination (such as syntax-semantics coordination)?
- **Developmentalism**
What is the role of developmentalism in explaining cognitive-computational phenomena? Does the developmental nature of disorders require a certain kind of account? How do versions of the machine-learning paradigm apply in the context of developmentalism? What are the nature and role of the initial state?
- **Functional Realizations**
Statistical and symbolic processing are *both* behavioral realizations of the mental computation system. How do these manifest themselves, given task and architecture? Given task and architecture, how is control functionally realized?
- **Extending the Explanation**
How far can we push the biological and cognitive plausibility of the computational vocabulary? Why and how do devices with brains—apparently only those things that also move (Llinás 2001)—develop internal regulatory and representational coordination schemes? How might cognitive management supervene on biological management? Do we see cases of the evolutionary emergence of a data management system?³

2. CONTROL REVISITED

2.1. Coordination, Communication, and the Form of Control

Four important lessons about control emerge from the commentaries of Cooper, Chandrasekaran, Kovordányi, and Perner and Dienes. The first is that control is best understood generally as the means by which a computational-cognitive system coordinates its representational domains. The level at which collections of representations constitute the domains of the system is an open question, as is the “empirical warrant” of the communication system (Pylyshyn 1996). But, as Newell observed, insofar as any system must trade and track data, it has to have a way of not simply running the data structures, but of managing the way the processing unfolds (with systems and their designers more or less explicit about the control assumptions). Control thus gives way to the larger issue of representational communication across such domains as the system realizes functional modularity—the interaction of semi-autonomous subsystems (Cooper 1995) or, as Aronoff observes for morphology, their partial correlation.

The second lesson is that communication and coordination are a function of the architecture and the coding constraints in the system. It may not be possible to tell from the outside what kind of information is static knowledge and what kind manages coordination (Chandrasekaran), and the latter function may take on a variety of forms in the context of the system. Thus, the classic, sequential, imperative language in which I

³I have long thought that there would be value in examining the evolution of phonological rehearsal (even in nonlanguage tasks) in human working memory. In my view, this would be an interesting way of looking the emergence of a species-specific, linguistically mediated control system.

phrased the initial proposals is but one way of accomplishing coordination, which may be differently and more subtly realized via different architectures. Production systems, for example, place a premium on control processes since they organize the problem space. Object-oriented programming builds communication into the data structures by defining objects as data plus a statement of functionality. So when we speak of communication and coordination of representations, we need to do so relative to an architecture and tasks the system has to perform.

Third, coordination of representations may be more or less direct. As Chandrasekaran observes, representations might communicate with each other directly such that one is designed to read the output of another—"what module M1 writes only module M2 can read." Alternatively, communication may be more indirect, with one set of representations reading the output of another only partially or via some other mechanism. In short, coordination can manifest itself on a scale of directness of linkage of the representations that have to be managed.

Related to the matter of directness is the fourth lesson (mostly from Perner and Dienes): the information that constitutes the linkage in communication may be more or less explicit in form, depending on how the interacting representations need to use the information that is passed. If one set of representations must check what another module is doing, needs to track data, or requires information "about a detail that was not passed on automatically," then the form of the communication must be explicit (*predication explicit*, to be precise); in contrast, information that is passed directly can remain implicit. What exactly is made explicit or implicit depends on the logical constraints on explicitness (which Perner and Dienes outline), the nature of the computational system, and the tasks to be performed. But the important overall lesson is that communication and the form of representational linkage fall on a gradient of explicitness.

Directness of read between two representations correlates with implicitness of form. Within-domain coordination, preferentially implicit, comes to look like what Perner and Dienes call *vehicle control*, and as a consequence stays beneath awareness. But the more explicit the form of representational linkage, as in cross-domain coordination, the more indirect the read: *fact explicit* content (*content control*) and explicit representational coordination are more likely to rise to awareness, which is why coordination across very different domains of representation, say vision to action, can rise to awareness.⁴

These four lessons come together in a single one. Let us look at the tasks a system has to perform to see how, relative to that system, coordination may be accomplished more or less directly/explicitly. This might in turn tell us something about the form and locale of the coordination mechanisms relative to a task or domain. If there is a gradient of explicitness of the form of coordination mechanisms and directness of read, depending on the task architecture and the way communication is to be realized and managed, then we may see the system- and task-specific coordination mechanisms in some of the data. Language provides an interesting case here.

There is evidence from many processing studies that the language system must evince both modular and interactive effects for context, depending on the representations to be

⁴A related issue is the directionality of communication. Do modules read each other uni- or bi-directionally? Chandrasekaran remarks, "As far as I can understand, he [Frawley] does not make a distinction between problems with reporting to versus problems with reading from WM." But I do touch on this in the discussion of van der Lely's findings about the asymmetry of breakdown at the syntax-semantic interface in SLI. Directionality is important for its intrinsic computational interest and its role in disorders like conduction aphasia and output deficits. The papers in Monsell and Driver (2000) are instructive on the kinds and costs of controller-controllee relations.

processed and the tasks at hand. These different effects pose different forms of representational coordination between context and the representations processed, themselves traceable to the explicitness gradient.

Consider a well-known example from the lexicon. Phonology has to interact with (feed) the lexicon: lexical forms provide a context for phonology. In turn, the lexicon must interact with (feed) syntax, which is a context for lexical forms. Given that the linguistic system must resolve the relationship of representations to their contexts, we can now ask: how do lexical representations, as the context for phonology, communicate with phonology? How do lexical representations, as contextualized by syntax, communicate with the representations of syntactic context?

It turns out that there is not a single answer to these questions. When these different communication relationships are processed, the mental system shows contrasting behavior. The phonology-to-lexicon-relationship is interactive: word structure apparently directly and immediately affects phonological processing and segment choice. However, the lexicon-to-syntax relationship is modular: syntactic structure does not facilitate lexical choice directly, but instead generates a number of options from which the appropriate choice is selected after generation (Tanenhaus and Lucas 1987; Tanenhaus, Dell, and Carlson 1987).⁵

There is a certain rationale to this division of labor. The representational communication issues follow from the representational domains themselves and the tasks the system has to perform. In exhibiting interactive effects the phonology-to-lexicon interface allows the representations of lexical context to communicate directly with the phonological representations during processing, selecting at the point of generation the proper segment for the slot in the word. However, the communication between lexicon-to-syntax relationship is different, requiring the mental system to generate all lexical forms, hold them for a short time, and thereafter allow syntactic context intrude to select the proper one.

This analysis returns us to the matter of explicitness and directness of read. For the phonology-to-lexicon case, the system generates representations that are *readied to the task*: by their form, segments can be read directly by the lexical context. In this case, communication between form and context is implicit and looks remarkably like Perner and Dienes' vehicle control. The system has no need for lexical representations to check what they inherit because there is no representational detail that needs to be resolved on grounds other than what the phonology provides. In this case, the phonology and lexicon work on the same representation (as Tanenhaus and Lucas (1987) remark, the phonology-to-lexicon relationship is one of part-whole).

The opposite holds for lexicon-to-syntax communication. In this case, the lexicon and syntactic context have to communicate indirectly and via explicit form. The lexicon overgenerates forms, a result that requires syntactic context to check expressly for the item that is relevant. That is, the resolution of the relationship between word and sentence structure cannot be resolved by one module alone: the representation of syntax needs to check explicitly what the output of the lexicon is in order to ensure that the right lexical form is merged with the context (as Tanenhaus and Lucas (1987) remark, the relationship between lexicon-to-syntax is one of set membership).

We can put this another way: the way the system allows dataflow between phonology and lexicon and between lexicon and syntax depends on the tasks to be performed and the

⁵This account again brings up the issue of directionality. Are certain directions of flow more likely to require direct visibility?

way context needs to check or inherit forms into its own processing. Phonology-to-lexicon might be understood as a kind of “blind inheritance” of data whereas lexicon-to-syntax is a kind of “checked inheritance.” (Are there varieties of inheritance from computational theory that might elucidate these findings?) Alternatively, we might think of this in terms of the linguist’s notion of visibility: segments are wholly visible to lexical context while lexical items are opaque to syntax. According to linguistic theory, many structures are built by the passing and checking of features across representational domains. Is this feature-monitoring the way language implements computational control?

This kind of dual behavior in the lexicon surfaces in many other areas of language. Similar findings hold, for example, in lexical naming, where there are contrasting findings about how lexical representations communicate with semantic and phonological representations (Ellis and Humphreys 1999, pp. 361–362). Analogous results surface in the parsing literature, where there is a well-known split of evidence showing that sentence processing can be both modular and interactive. Indeed, it appears that, given the task and performance capacities of individuals, “human sentence processing can exhibit both modular and interactive effects” (Lewis 1996, p. 77).

In his review of the parsing evidence, Lewis (1996) shows that a single computational system (Soar, in his case) can produce both these effects, given the nature of the representations involved, the tasks to be performed by the system with respect to these representations, and the way control is implemented in the system. Basically, the system generates encapsulated syntactic structures but varies in how the production system operators build an updated model of the utterance. Lewis (1996, p. 80) says, “If the relevant semantic or contextual associations are in place to guide the selection of u-constructors. . . , then NL-Soar will behave as an interactive system, with all knowledge sources immediately participating.” Alternatively, if there is a direct read between the syntax and the other domains, there is no need to explicitly check the fit of the structure and the utterance model. But “if the relevant semantic or contextual associations are not in place. . . then NL-Soar will behave as a modular system, resolving the local ambiguity without recourse to these knowledge sources” (Lewis 1996, p. 80). In such a case, there would have to be an indirect read between syntax and context, with the latter explicitly checking the fit. Importantly, “NL-Soar is modeless—it is not waiting for its ‘interactive switch’ to be turned on. Rather, the behavior it exhibits will span a continuum of modularity and interactivity depending on the current state of its operator knowledge” (Lewis 1996, p. 80).

The lesson here is the crucial position in a (mental) computational system of the way control manages communication across representational domains.

The fine-grained control over the syntactic module is a direct result of embedding the processing in Soar’s control structure. Because control in Soar is obtained at the level of the operator (decision cycle), it makes the prediction that syntactic processing is controllable at that level. Temporally, this means on the order of 100 ms—at least at the word level, as opposed to, say, the phrasal or clausal level. (Lewis 1996, p. 81)

This is an important observation, given recent advances in linguistics. Minimalist Theory proposes that syntax works from the word up, with the system selecting a lexical form and generating structure outward therefrom, with features passed across domains to check the construction, ultimately allowing the structure to merge with both phonological and logical form. The fact that a single computational system can produce both modular and interactive syntactic effects by applying *control at the word-level* is a convergence of great

interest because it links both the logical structure of language with a possible implementation.

Similar results can be found in areas of cognition other than language (see Armstrong, Gleitman, and Gleitman (1983) for the classic paper on concepts), but these findings as a whole provide a larger set of arguments about the role of control and representational communication. First, there is a clarification. One might read the foregoing and say, as Kovordányi does, that this is “just processing.” Does recasting modular and interactive effects as communication and control issues yield anything that could not have been gotten otherwise by the simple processing vocabulary?

I finally have to come clean on a basic point: I no longer know what *processing* really means. It is the same confession I once gave for *learning*, which I used to think I understood. Seeing that there are so many ways that a knowledge system can input, store, update, change, and output, I have come to see *learning* as a cover term for a variety of procedures, phenomena, and states that work in concert to produce learning as an effect. Learning itself is best understood in more specific terms: untutored learning, triggering, hypothesis testing, explicit learning, meta-learning, and so on. I now feel that way about *processing*. Many factors conspire to produce the phenomenon we call *processing*: procedures for building representations (including the system’s intrinsic capacity for incremental construction and retrieval), the nature of representational formats (rules and abstractions or shallow code), the number of representations available at any one time (parallelism/sequentiality), the relative encapsulation of representations (modularity/interactivity), the visibility of representations to each other (control), and so on. What I am saying is that the directness/explicitness characteristic of representational communication is neither the whole of processing nor a restatement of the interactivity/modularity distinction (itself a separate processing parameter). To process, a mental computational system has to produce a lot of representations at once or one at a time (parallel/sequential parameter); these representations can intrude on each other or not (interactive/modular parameter), and these representations have to talk to each other in certain ways (direct (implicit)/indirect (explicit) control parameter). All kinds of in-principle systems fall out of these parameters. There could be a parallel, modular system with implicit or explicit control, for instance. Interactivity and implicit control seem natural allies, as do modularity and explicit control, in the same way that parallelism and interactivity are natural allies, but these are not necessary relations.

Parallel processing is consistent with both lexical-form cases above. There is no reason why the phonology could not generate all forms in parallel and why the lexicon could not do the same for syntax. The visibility issue is how these multiple representations generated in parallel in their respective domains coordinate and communicate in order to merge. A form may be directly visible to context, but the timing of the uptake of this representation by context may vary—at the point of generation (interactive) or later (modular). The issue is how, given this running of the system, communication between representations is structured and managed. In one case there is a direct read via implicit form; in the other there is an indirect read via explicit form. In both cases the running is parallel.

So when we want to claim that certain results are a consequence of processing, I think it is important to ask: *what aspect of processing?* In the same way that those who study experimental phonetics investigate stress as the consequence of such factors as length, amplitude, and pitch of vowels, I hope to have shown by this paper that the visibility of representations is a distinct feature of the complex concept *processing* that deserves attention.

Now let us consider some predictions that would seem to fall out of this rethinking of control as representational visibility and communication. In cases of direct visibility,

where one representation is inherited blindly into another domain and so the form of control is implicit, the processing ought to be faster and the effects ought to look more performance-like because the formal requirements have been minimized. (It might be interesting, in this respect, to study the performance effects of the morphological implementations that Aronoff describes as default inheritance.) Direct visibility is thus functionally a within-domain communication matter because the two domains share data directly. Interestingly enough, the processing of phonology-to-lexicon representations evinces characteristics of general performance—overall statistical effects, like frequency. But in lexicon-to-syntax relationships, processing has a distinct time delay and the effects of frequency are militated. This would seem to be the result of the fact that explicit communication must reference the form of the representation, or be content-driven, because it is distinctly cross-domain.

These considerations suggest a new picture of the way information is communicated across representational domains: direct, or matched/overlapping representations where the communication is implicit and hence potentially visible to another representation at the point of generation; indirect, or formally mediated representations where the communication is explicit and hence potentially visible to another representation after generation. We might sketch this as follows:

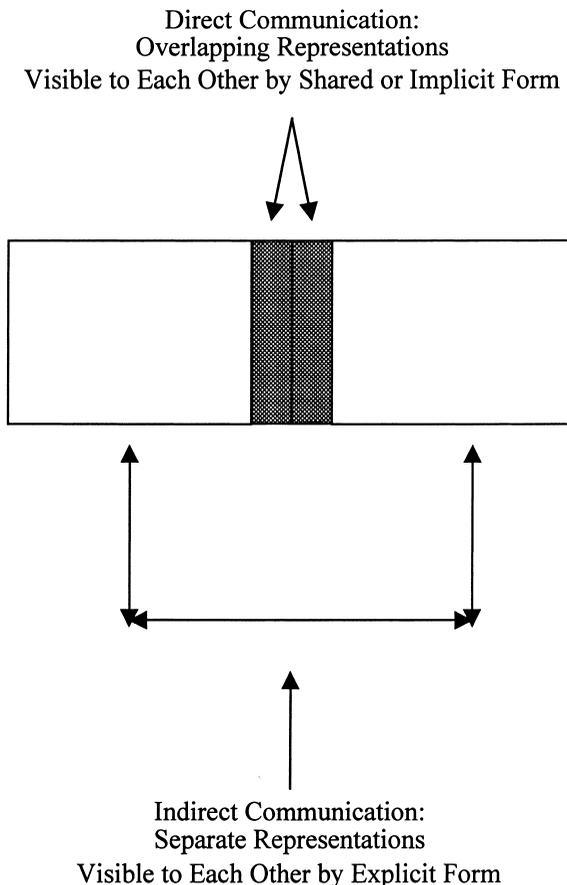


FIGURE 1. Communication across representational domains.

The usefulness of this new picture of control can be seen in its application to another area of mind, emotions. Emotional processing has five logical domains of representations that must communicate with each other in order for successful emotional processing to occur (see, e.g., Oatley and Jenkins 1996): (1) elicitors: causes and emotional cues; (2) appraisal: evaluation of elicitors and representation of feelings; (3) secondary appraisal: inferences about significance of emotion; (4) action readiness: plans for affective response; (5) execution: overt emotional behavior. That is, emotions are triggered by recognizable elicitors, which are then appraised for their emotional content; these appraised representations are then judged for more specific significance and these secondarily appraised representations feed plans for action, which in turn link to motor plans for execution. Interestingly, there are kinds of alexithymia (emotion blindness: see above) associated with disconnects across each component: i.e., failure to make information from one component visible to another (Frawley and Smith 2001). The new picture, above, of representational visibility constrained by explicitness clarifies some details with respect to alexithymia types since linkage across the representational components may vary in directness.

The clearest cases of alexithymia in the clinical literature are those with a disconnect between elicitors and appraisal and between secondary appraisal and action/execution. The harder cases to discern—but they *are* discernible—are those that delink appraisal and secondary appraisal. Given the gradience of explicitness of control, it would seem that the linkage of elicitors and appraisal and of secondary appraisal and action require that the system deliver representations across these components in explicit form. Effective communication across these domains of representation would require the inheriting component to work on a representation of a different kind from the feeding component and hence the need to check features. For example, the transition from an emotion-invoking stimulus to the appraisal of that stimulus necessitates the transformation of information in quantitative form (intensity) into qualitative form (fear) (what LeDoux has called the low-road/high-road distinction in emotional processing): indeed, there are separate neural pathways for these procedures. Similarly, on output, the system must take emotional appraisals and convert them to representations compatible with the motor system, again along a special pathway that is keyed to emotive response (there are different processing routes for smiling vs. mere teeth baring). Thus, one might say that to cognize an emotion and to generate proper emotional behavior, the system has to manage representational communication across divergent representational domains, which in turn would call for maximal explicitness of representational form.

However, in making inferences on appraisals about the significance of the appraisal to the context or the self, the emotional system is working on the same representation it constructed from the input. Secondary appraisal requires the direct passing of an evaluation into a further evaluation and hence no check for validity. Secondary appraisal blindly inherits the appraisal. The representations of appraisals are directly visible to the mechanisms that will infer the significance of them.⁶

⁶In more philosophical terms, we might say that the distinction between explicit and implicit linkage in the emotional system contrasts the semantics and pragmatics of mentalese. When appraisal constructs a representation from an elicitor, it has to check the truth (semantics) of the elicitor; similarly, on output, when motor plans inherit the representations of secondary appraisal, they need to check the semantics of the secondary appraisals to judge sheer applicability and the very basis on which they will operate. But when secondary appraisal operates over appraisals to add significance to the appraisal, the system does not need to check the truth of the appraisal any more than an implicature needs to reference the truth of a proposition in order to occur. Secondary appraisal simply inherits appraisal. One might say, then, that the relationship between appraisal and secondary appraisal in mentalese is one of pragmatics. An intriguing theoretical question arises: what do the pragmatic aspects of mental representations need to make explicit?

When we get disorders of emotional processing, the effects seem to correlate with these observations about directness. Input and output disconnects in alexithymia are much more likely to rise to awareness than are appraisal-secondary appraisal disconnects. Insofar as explicitness correlates with awareness, these consciousness effects would seem to be a function of how the emotional system manages the visibility of information across functional submodules.

This recasting of issues into representational visibility also leads to important reconsiderations of the language disorders that have formed the basis of much of the evidence in this paper. It would suggest various manifestations in breakdown: control losses that look to be within-domain because of the disruption of indirect communication; losses that look to be cross-domain because of disruption of indirect, explicit processes. Before I turn to those, I would like to offer some additional questions that are implied by the analysis and might be the basis of further discussion.

2.2. Lessons Learned

- **Visibility**
What does it mean for one representation to be visible to another? What features of a representation can be made explicit or left implicit? What is at stake in the difference between blind and checked inheritance? For example, what syntactic features might be blindly inherited from one domain to another, or must be rendered into a particular form for recognition by phonetic form or logical form? Or to use the phrasings of Perner and Dienes: what kind of representational processing requires two modules to work on the same entity (implicit form) or on different entities (explicit form). Do we get kinds and degrees of linked systems and subsystems?
- **Languages**
Aronoff observes that complex representational communication seems to arise with complex morphology. Partial correlations across subsystems happen with cross-checking via abstract morphological types, like declension classes. Can we trace the development of these communication/visibility systems? When pidgins “complexify” into creoles, do the representational visibility mechanisms also change? Aronoff (personal communication) has pointed out to me that morphological development follows what appear to be the rules of grammaticalization (my term, not his) in that semantically motivated morphological phenomena become formalized. Which kinds of formal phenomena become tracked for visibility? Kirby’s (1999) analysis of the evolutionary constraints on language might also be relevant here.
- **Effects**
Are there specific processing effects associated with each kind of communication? Or better, what are the behavioral effects of the difference between the visibility of representations to each other vs. the timing and uptake of these representations? Explicitness ought to correlate with memory intensiveness and delay, for example.
- **Development**
How does a mental representation system get these kinds of visibility? One could postulate these distinctions as learned phenomena, as Lewis’s (1996) findings seems to suggest is possible for parsing, where the system is ready to communicate but is built with enough flexibility to adapt to direct or indirect, depending on the system’s history. Or the distinction might be part of the system’s specifications. Is there any developmental evidence that children work first from maximal separation of

representations and then develop representational overlaps? Must increased capacity of working memory be in place for increased kinds of representational communication?

3. DISORDERS REVISITED

3.1. SLI *versus* Williams Syndrome

Whatever the theoretical sympathies of the commentators, their responses to my analysis of SLI and Williams syndrome suggest to me that the split between these two kinds of disorders remains a clear one. What is the nature of this split and how do we account for it? I ultimately want to reassert the logic/control parameter, understood much more productively as an issue in visibility and explicitness of form, as essential to the explanation.

Van der Lely and Gopnik show convincingly that SLI is a well circumscribed disorder that affects submodules and so has distinct subtypes: relatively pure grammatical SLI, relatively pure morphological SLI, and so on. Moreover, each kind of SLI cannot be traceable to a disruption in another module (e.g., SLI is not, at base, an auditory or phonetic problem). Morphological SLI is a problem with domain-specific word formation mechanisms: “SLI subjects simply do not have the underlying, automatic rules for relating words or affixes” (Gopnik); grammatical SLI is a problem with the computational rules that underlie syntactic dependency: “the core deficit responsible for G-SLI grammar involves ‘Movement’” (van der Lely). The strength of Gopnik’s argument is the range of evidence, cutting as it does across languages, neuropsychological tests, and familial distribution. The power of van der Lely’s argument is the narrowing of the data: grammatical SLI is traceable to only certain, well-circumscribed aspects of syntactic knowledge. Surely there is something specific to the S in SLI! The conclusion is simply unavoidable that SLI is distinct, core-computational, and domain-specific, a logic disorder—just the sort of thing that would go haywire in the ways it does when Nature’s errors range over certain parts of the factory-installed equipment.

One wonders if, in the face of these arguments, a radical developmentalist would object to lumping morphological SLI with grammatical SLI because they likely have different physical causes and different patterns of behavior, even though they share an explanation at the knowledge level. Moreover, finding that individuals with SLI also have motor problems, as Thomas observes, does not mean that the core of the disruption is not specific. I cannot think of any deficit that has no effects beyond itself.

Thomas presents evidence that such clear-cut argumentation for Williams syndrome (and SLI mainly by implication) is less compelling. Now that I have come to understand control in more specific terms, I partly agree with him, but my reasons for doing so are different from his. Thomas thinks that Williams syndrome is a developmental disorder with within-domain language deficits. I now believe that representational communication has two landing sites, cross-domain (indirect or explicit) and “functionally within-domain”(implicit or direct). So if Williams syndrome is a visibility disorder, we might expect it to surface in two forms, one of which looks like a within-domain disorder. This is the argument I hope to make when I reconsider some of the data. But first I want to turn attention to the issue of explicitness.

3.2. Explicitness

If SLI is classically within-domain and if within-domain knowledge is preferentially implicit in form (Perner and Dienes), then we should expect failures in SLI to correlate

with implicit knowledge format and, by implication, that SLI successes ought to be in tasks which require explicit knowledge. This is partly borne out. As Gopnik observes, individuals with morphological SLI fail at the automatic running of morphological rules but can make use of knowledge in explicit form to execute morphological tasks. They can use brute force memory of unanalyzed forms and apply metalinguistic knowledge, such as memorized rules, to override failures in the execution of the combinatorics of morphological logic.

Individuals with Williams syndrome evince the opposite behavior. They apparently can run the implicit rules (pace Thomas) and otherwise show deficits in areas where explicit knowledge seems required. Thus, the claims in the literature that Williams children may have a working memory deficit are, in my view, actually claims that these children have a deficit in the construction of domain-specific knowledge in explicit form for delivery to working memory and for overt checking of validity between domains of representations. I want to (re-)underscore the suggestiveness of Karmiloff-Smith et al.'s (1998) observation of Williams children's poor performance on explicit tests of syntactic knowledge.

So far so good, but a more subtle problem arises. Is the distinction between SLI and Williams syndrome entirely that of implicit vs. explicit knowledge? Does each condition have a broader, non-domain-specific processing deficit that translates into these specific effects? There is a clear reason why an explicit/implicit deficit cannot be the whole story—if it were, then we should see other, more generalized explicit/implicit failures. It is not as if SLI children fail at all tests of the automatic running of core computational rules, not even on those confined to language. Nor do Williams children exhibit a generalized explicitness deficit. What seems to be happening with these syndromes is the result of the interaction of the visibility of representations and explicitness. This would explain some of the interesting counters raised by Thomas in his discussion of additional data on Williams syndrome.

3.3. Within-Domain and Cross-Domain

There are two sets of counters to my analysis of the SLI and Williams syndrome data, heretofore not addressed, that I would like to raise now for discussion to clarify the disorders themselves and the mental-computational analysis of them.

3.3.1. Deficient Regulars in Williams Syndrome? Thomas reports on a re-running of experiments that led Clahsen and Almazan (1998) to argue that the morphological deficits in Williams syndrome are confined to irregulars (not a rule-construction deficit). I used their findings to illustrate the logic/control split. Because irregulars require visibility to phonology and are stored as whole forms with marking of their phonological patterns, they are control-like, essentially cross-domain, and ought to be disrupted in Williams syndrome. (Regulars, because they are combinatoric and entirely within-domain, ought to be disrupted in SLI.) Thomas finds that Williams children also exhibit problems with regulars, thus calling into question my argument that the regular/irregular split into SLI/Williams is justified by a logic/control split. However, with control now taking on two forms, and a more subtle analysis, we ought not to be surprised by Williams children also failing on regulars in the way they do: extension to novel forms.

Here is what the morphology-phonology interface looks like:

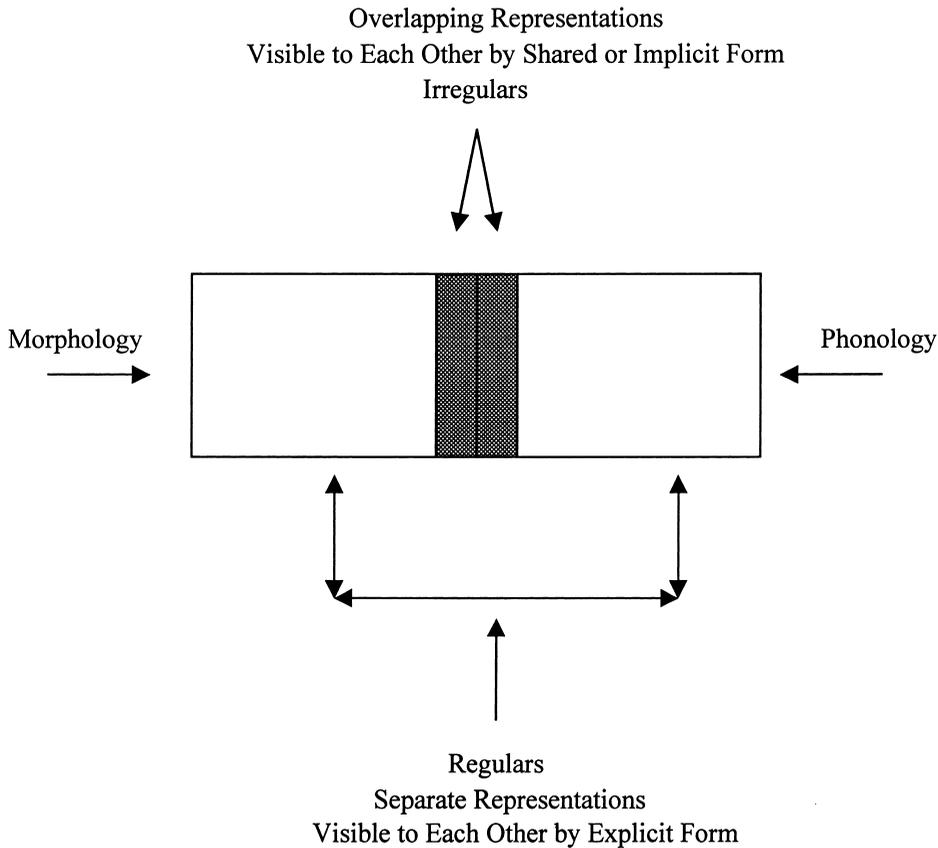


FIGURE 2. Morphology-phonology interface.

Regular forms, such as *want + s* or *want + ed*, are computed in abstract form via the combinatorics that merge stems and affixes and then sent to phonology for the running of phonological rules. For regulars, the phonological module receives from morphology something like WANT + ED, the form of which phonology must then explicitly check for applicability to the phonological rules. Phonology “asks”: what is the last segment of the stem? If voiced, add a voiced affix; if a fricative, insert a schwa; and so on. Regulars require explicit visibility of the forms to phonology for the forms to be pronounced.

In contrast, irregulars enter phonology readied for their pronunciation. An irregular form, such as *sang*, is encoded in the morphology in its phonological form with its associative “hot links” to other phonologically similar forms: [sæŋ]→[ræŋ]. Phonology does not need to check explicitly for the applicability of the form to the domain-specific rules because the visibility is direct and implicit in the form itself. Phonology blindly inherits the irregulars. (See Pinker (1999) for this kind of analysis.)

Visibility (control) thus takes two forms in the morphology-phonology interface. It can be explicit, where regulars are checked expressly. It can be implicit, where irregulars are

intrinsically flagged for their visibility. This leads to four possible disruptions of the morphology, based on the interaction of logic/visibility and explicit/implicit form:

- Logic Disruption, Implicit Form
- Logic Disruption, Explicit Form
- Visibility Disruption, Implicit Form
- Visibility Disruption, Explicit Form

The first is classic SLI: loss of combinatoric rules in implicit form; it is also what Williams children apparently *can* do (see more for an analysis of Thomas's result). The second does not exist: logic, as within-domain, comes in implicit form. The third is what Clahsen and Almazan found for Williams syndrome and what Thomas found in his computational simulation of Williams syndrome (and note that Gopnik remarks that SLI children do pretty well with unanalyzed morphological forms). The fourth is what Thomas found on the redoing of the Clahsen and Almazan data. In short, there are two ways that visibility can play a role in morphology, and Williams syndrome manifests both. Let me concentrate on the last two types.

Clahsen and Almazan find that Williams children are poor at irregular morphology and fine at their regular morphology. If irregulars have a direct read into phonology, then this irregular deficit is clearly a control loss, but one which looks as if it is within-domain. Because the form of the representations is shared across modules, the forms are in fact functionally within-domain because of the way the computational system organizes domains and their communication. I might agree, here, that Williams syndrome manifests a within-domain deficit: a within-domain control deficit. I do not know, honestly, what to make of Thomas' dismissal of the Clahsen and Almazan findings on the irregular performance as "just a reflection of language delay." What causes the delay? In any case, I note that he does not find that Williams children are good at irregulars, which would be surprising. Moreover, in his manipulation of computational simulations for irregulars in Williams syndrome, Thomas reports that a reduction in similarity and redundancy of phonological representations yields irregular pasts defective in a way consonant with the Williams data. Because the manipulation disrupts the very factors that make irregulars intrinsically visible to phonology, not the morphological structure of the forms themselves, this is not at all unexpected. Irregulars come in whole phonological forms and are linked by their patterned associations to whole forms. So what I see Thomas arguing here is that the reason Williams children manifest poor performance on irregular forms—and the way you can simulate this loss on a connectionist machine—is to disrupt the visibility conditions on these forms as direct reads into the phonology: Williams syndrome is a control disorder.

A similar, though more complex, story holds for Thomas's finding of a disruption of regulars in Williams syndrome. Regular forms are generated abstractly in the morphology and then sent to the phonology for explicit checking. If Williams syndrome is a visibility disorder, then we should see a loss *in this checking*. Williams syndrome ought to manifest a loss in regular morphology, not because of a breakdown in the combinatoric rules (SLI), but because of the failure of the visibility mechanisms to take well-formed structures and read them explicitly into the phonology. I do not have access to the data and so am proposing this analysis in the interest of expedience, but I want to suggest that this is what Thomas has found in the regular disruption is: poor uptake of well-formed structures (vs. SLI, with its good uptake of ill-formed structures). How can we be sure that this is what has happened?

I think it is significant that he finds that the regular losses were restricted: "the WS group were less prepared to extend the regular past tense rule to novel items." (This same failure in application to novel instances also occurred in the Karmiloff-Smith et al. (1998)

data, which he notes as the only place where the deficit in that study appeared.) What is it about novel items? In novel regular morphological forms, it seems to me, the phonology has to be called explicitly for the forms to merge. Alternatively, for novel regular forms, the system could generate the morphological forms independently of the phonology. So what we may be seeing in the case of a deficit in regulars for Williams syndrome is well-formed generation of the words via abstract structures, but a failure for those structures to be able to be interpreted explicitly by the phonology. How would we know this? If we had two novel forms that put the same demands on the morphological system but different demands on the phonological system, then we ought to see a difference on these for Williams children. Alternatively, if we had two novel forms that placed different demands on the morphological system but the same demands on the phonological system, we ought to see differences for SLI. Some modifications of data from Pinker's (1999) wonderful study might help in this respect.

Take two novel forms, *ploasted* and *ploasped*. Each poses similar computational demands on the morphology. Each is regular because neither evokes the phonetic pattern of an irregular and so each is analyzed via abstract stem and affix: PLOAST + ED and PLOASP + ED. A sound morphological module ought to handle these equally well. However, these two forms differ in how the phonology can read them. English allows stems with syllable final clusters preceded by a long vowel only if the clusters are apical (use the tip of the tongue). Thus, *ploasted* is well formed in the phonology (*st* uses the tongue tip), but *ploasped* is not (*sp* violates the apical rule, with *p* produced by the lips). If the explicit visibility argument is correct, Williams children ought to do better on the former than the latter because *ploasped* places direct computational demands on the morphology-phonology interface: that is, the system would have to explicitly check the abstractions delivered out of the morphology, and the form taxing the checking ought to evince performance deficits. It would be interesting to see the data Thomas has on the Williams children's failures of regular to see what kinds of computational demands the novel forms might place on the way the morphology and phonology should be talking to each other.

In contrast, consider *desked* and *plicked*, which are both regular but in different ways. The former has an abstract tie to a noun root and so the abstract stem is a lemma, covering both noun and verb formations (like *sink* in *I sinked the dishes into the water* 'put into the sink' vs. *sank*: see Pinker 1999). But the latter is completely novel and arguably a simpler stem in the abstract. Neither, however, poses any particular problem for the way phonology has to read them. In this case, we might expect SLI individuals to do worse on the former because of the computational demand on the morphological system itself.

I admit that I have no idea about how these predictions might turn out in detail. I am making them only to suggest that there might be a way of sorting out empirically the claim that Williams syndrome has two manifestations of a visibility deficit by showing that cases where Williams children seem to have a combinatoric problem (and so are like SLI) may actually be illustrations of problems of representational linkage: faltering on novel regulars precisely where the phonology must explicitly check the combinatorially generated forms for visibility. So when Thomas argues that his simulations show that SLI and Williams syndrome "arise because a developmental process acts on phonological representations which are deficient in two different ways," I must agree with him. SLI originates from a rule deficit which sends corrupted data across an intact phonological interface. Williams syndrome originates from a visibility deficit, in two forms, which sends good data cross a corrupted phonological interface. The processing is indeed deficient in two different ways.

3.3.2. *Real Counters to Williams Syndrome and Visibility?* The foregoing shows that Williams syndrome can have what appears to be within-domain effects by virtue of how one

form of visibility is structured and managed via a direct read of representations. In other words, within-domain deficits in Williams syndrome are actually cross-domain, where the domains read each other directly, share the same representation, and so mimic within-domain loss. True within-domain loss—a deficit in the implicit, core computational rules—does not appear. But Thomas raises a variety of other critiques of the cross-domain argument by citing data from very recent studies. I will save direct response to those results for a later issue, but let me just describe how I might approach the data.

Thomas indicates that there is data to show that Williams children's lexical problems are not just cross-domain (e.g., semantic to pragmatic mappings) but also within-domain, "disordered (fuzzy) semantic representations." Williams children evidently evince behavior in the lexicon that is both within- and cross-domain. This sounds like the regular morphology problem. What is the theoretical view of the structure of the modules in question and how are the representations in them visible to each other? For example, if the lexicon is an abstract, formal structure (made of things like individual, predicate, functor, etc.) and not some rich set of conceptual types (THING, PERSON, etc.), then it may not be correct to analyze the disruption as "fuzzy": a fuzzy functor? The fuzziness might come from a perfectly intact formal lexicon interacting deficiently with other information. Indeed, on the theory that the lexicon is a set of pointers between word forms and concepts, one might call the performance of Wernicke's aphasics "fuzzy" because they have a preserved formal lexicon but faulty linkage of the forms to conceptual structure. We might also call the performance of Williams children on morphological regulars "fuzzy," but I hope to have shown that it is a visibility issue. The charge of "fuzzy" can be, well. . . , pretty fuzzy. Furthermore, semantics and pragmatics may share representations in the way that morphology and phonology do. Certain implicatures may be more or less directly read off the semantic form. Are we again in a position to reanalyze the data with a different theoretical apparatus, thus yielding a subtler picture of what counts for within- and cross-domain?

A similar question arises with Thomas's suggestion that Williams children have syntactic deficits, as evidenced by their failure on complex constructions, such as relative clauses. When I see that data, I will want to look closely at the cause of poor performance. Relative clauses have been well studied and are known to pose distinct difficulties in terms of how the mind has to merge the interpretation of the structure with the mental discourse file (a classic visibility issue, not one of computational construction per se). Relatives raise a particularly difficult problem of assertions and the tracking of individuals' identity; markedly better performance can often be elicited on experiments with these structures by the simple rephrasing of stimuli (Crain and Lillo Martin 1999).

So I am intrigued by new, purported, within-domain difficulties in Williams syndrome as evidenced in problems with the lexicon and relative clauses. But I am also very suspicious, if the lessons of two decades of experimental research on these structures are a guide.

3.4. Lessons Learned

- Closeness of Analysis

What exactly is disrupted and how in these disorders? Can a more detailed linguistic analysis of the data suggest different locales of loss or even commonalities of loss across apparently disparate phenomena? Van der Lely isolates G-SLI in the computational mechanism related to movement. Does M-SLI have a similar loss in the combinatoric rules rather than features? Does Williams syndrome evince losses in the syntax-phonology interface or the syntax-semantic interface analogous to the two kinds of visibility loss in the morphology-phonology interface?

- Explicitness
How much of a role does the explicit/implicit dichotomy play in these disorders? Can the failures in performance at one form suggest methods for intervention via the other—as SLI individuals can use explicit forms?
- Generalizability
Does the distinction between logic and visibility apply to other, non-language mental domains? Do we see similar disruptions in spatial knowledge, face knowledge, or musical knowledge? Does information in these domains interact with explicitness in the same way?

4. INCOMMENSURABLES

I want to close this response by looking at three places where I see little hope of progress and a common dialogue as long as the discussion stays in its present form. These three incommensurables are all from Stanlaw’s paper, but my saving them for this point in the response is not meant to single him out for incommensurability. These points of difficulty likely come from Stanlaw being an anthropologist and so, if not totally aligned with social and cultural views of language and knowledge, at least sympathetic to them and their attendant clash with abstract computationalism. Still Stanlaw’s commentary raises a number of worries that lurk either partially stated or unstated in other papers (e.g., in Thomas’s comments) and certainly will be in the minds of others who may read this debate. I will keep the analysis short and do my best to try to reach back to find a common ground.

4.1. Hardware/Software

Stanlaw remarks: “Even if we assume that the organic clinical problems demonstrated by Williams syndrome people manifest themselves as problems in mentalese, this is still not a software problem. For example, no amount of software tweaking can repair a bad chip or address. Thus, the analogy still fails.” In saying this, if I understand him, Stanlaw is worried that it is not possible to speak meaningfully of the software/knowledge level when the hardware/neural level is involved. Response to this kind of apples-and-oranges charge—brains are brains and minds are minds—can be very challenging because it sounds as if he is saying that you cannot make claims about (disrupted) mental structures caused by (damage to) the brain they run on *because the cause is physical*. If he is in fact saying this, then he and I have a very long road ahead of us, and it seems necessary for him and me to sit down and talk about what it means for a brain to be related to mind.

In that conversation, I would offer something like mild functionalism and supervenience. Mind is what a brain does. A change in mind requires a change in brain (so Williams syndrome involves a mind change and a brain change), but not vice versa. So mind states are not reducible to brain states, but exist in an independent level of analysis supported by and correlated to brain states. For some, the relative independence of mind and brain is not theoretical, but a necessary property of the universe—not a position I adhere to.

On this view, Williams syndrome is both a software problem and a hardware problem. The real lesson, however, is a mind-brain link: brain-constrain your mind claims. I admit to having done only a little of such constraining in the foregoing. But I have offered no less than anyone else, as far as I can see! And no one is the worse for wear.

If there can be no analysis at the knowledge level because there is a physical level involved (and how else could the world turn out?), then we really cannot go on talking. If we must talk in only physical terms because there is a physical level involved

(eliminativism), then we really cannot go on talking. (*Talk* is just too non-physical a term!) But if we can talk at the knowledge level with a more-or-less-sensitive eye to the physical level, then perhaps we can go on talking. I would contribute to that conversation with something like the following:

The language problems in Williams syndrome suggest to me that it is a deficit in the computational visibility of representations to each other, not a deficit in the representations themselves. I would have to work with a cognitive neuroscientist on locating the physical correlates of the visibility problem. (Incidentally, all neural correlates for the behavior of individuals with Williams syndrome remain quite murky.) But I do know that the brain has control devices and areas that function to manage dataflow. So it would not be unusual to find physical locales of visibility. Note in this regard that one common charge by connectionists against symbolic mental computation is that there is no physical correlate of symbol storage. But Marcus (2001, pp. 54–56) argues that intracellular processes could implement registers that store variables. Here is symbolic computation with a more-or-less sensitive eye to the physical level.

In the end, I feel as if we all are sort of like Kant. He argued hard for the categorical imperative and then failed to thoroughly live by it himself. Is Kantianism thereby corrupted?

4.2. Mind Without Mentalese

Stanlaw observes: “it is erroneous to claim that because the mind produces a language, its internal functions necessarily are—or even can be—represented by a language.” This comment has me troubled—because I worry that it may mask a belief that minds and brains do not run complex code. My worry is supported by a later claim of his: “Even if the brain is ultimately computational, there is no reason to assume it acts like C++ or even LISP. Fuzzy logics, analogic computation, genetic algorithms, or fault-tolerant computing all seem to me to be better equipped to mimic the plasticity and redundancy of the brain.” This array of computational alternatives suggests a rhetorical strategy that may ultimately deny computation. Stanlaw goes on: “Is the physical universe computable? If not, it is unlikely that the biological or mental universe would be.” Stanlaw’s argument would thus seem to be that it is possible (likely?) that there is no mental computation over mentalese because there is no computation.

This is again a hard position to reach out to. I would start the conversation by saying that even the most anti-mentalesers acknowledge computation, “reason-respected behavior” (Clark 2001, p. 15), and that the brain is a “natural computer” (Clark 2001, p. 18). Maybe it is one that runs code different in kind from the machines we make and manipulate, or maybe it implements abstractions or executes functions in “natural ways.” Maybe it is a powerful statistical analyzer, coding transitional probabilities and recoding those as probabilities (Aslin et al. 1998). Whatever the code, a code-runner it is.

What kind of code-runner? A mind that produces the kinds of “extraordinarily ordinary” behavior that we see in brain-equipped organisms must be implementing a code with abstractions, operations over these abstractions, and other such earmarks of language-like structure (Marcus 2001). Desert ants, possessed as they are of small brains, store variables; insects reset their flight courses and compute new ones (Gallistel 1998). Bats store a neural representation in terms of time, frequency, and amplitude of the sonar signal they emit for prey identification so they can compare it with the returning signal to accomplish echolocation. But they also they turn their ears off right before they emit the sonar signal because the loudness of it on emission would damage their ears (Hauser 1997,

pp. 160–169). So, as the basis of comparison with the returning signal from the prey, bats use a complex mental representation that they never, in fact, hear. They must have imagined it. They must have thought it all out beforehand. How do they do that?

4.3. All That Other Stuff

Stanlaw, referencing his earlier (much appreciated) positive review of my book *Vygotsky and Cognitive Science*, says: “I think what made his book so interesting is what this paper lacks: attention to Vygotsky’s beliefs that language and the higher mental processes have ultimately a social basis. This, I think, is Frawley’s ‘road not taken,’ as he has chosen not to look at context or mere ‘performance,’ focusing only on computation.” He is concerned that my emphasis mental computation has bypassed all sorts of other interesting and relevant material that bears on mind, language, and disorders:

Often their [Williams children’s] formal language skills are surprisingly good, if sociolinguistically inappropriate. This suggests that emotional or social difficulties might also be involved with these maladies, as much as problems in the computational code. These clinical symptoms, then, could be reflections of some other kinds of breakdowns besides failures of abstraction (which seems particularly true for autism). Frawley’s dismissal of ‘performance deficits,’ then, perhaps is premature.

If Stanlaw believes that in order to explain Williams syndrome, SLI, or other functionally analogous disorders, we must *essentially* make reference to social context, emotion, and performance, then we are indeed at a point of incommensurability. We do not need a computational, or any other, theory of emotions to account for Williams syndrome, I am sure; nor do we need a theory of performance or social mind to explain SLI. If I came down too hard on social context, emotion, and performance, then I am to blame for that. But, as I hope I have shown not only with the epilogue of my book (Frawley 1997), and the work cited previously on emotion (Frawley and Smith 2001), these factors are very relevant to the computational mind generally (and its losses). But it is critical to understand their structure and the way they support mental ecology.

How can we start a conversation that brings social context, emotion, and performance into the computational mind (presuming it is computational: see Section 4.2)? Here is a suggestion for each:

- Social context is computed on the basis of a small number of core variables that we genetically inherit, plus a large number of learned factors that we experience in our sociocultural niche. There is ample evidence that genetically related species carry around a shallow social structure and evidence that we begin the world with a similar schematic template (Tomasello and Call 1997). There is also evidence that in social mind deficits (e.g., autism), specific computational structures are implicated and the deficits proceed in if not predictable, at least understandable, ways. That is, social context is like language.
- Emotions are computed on the basis of a small number of core variables that we genetically inherit, plus a large number of learned factors that we experience in our sociocultural niche. There is ample evidence that genetically related species carry around emotional structure and evidence that we begin the world with a similar schematic template (LeDoux 1996). There is also evidence that in emotional deficits (e.g., alexithymia), specific computational structures are implicated and the deficits proceed in if not predictable, at least understandable, ways. That is, emotions are like language.

- *Performance* is a cover term that means a lot of things, much like *learning* and *processing*. (I might also be inclined to say that performance works on the basis of a small number of core variables that we genetically inherit, plus a large number of learned factors that we experience in our sociocultural niche.) One thing that *performance* covers is the way a synchronized collection of representational domains runs in real time, with distractions, memory pressures, and other factors contingent on the moment and individual. Each individuated system has different performance factors, and if performance difference are individual, can there ever be a science of performance? Are these differences that make a difference?

REFERENCES

- AIZAWA, K. 1992. Biology and sufficiency in connectionist theory. *In* The Symbolic and Connectionist Paradigms: Closing the Gap. *Edited by* J. Dinsmore. Erlbaum, Hillsdale, NJ, pp. 69–88.
- ARMSTRONG, S., L. GLEITMAN, and H. GLEITMAN. 1983. What some concepts might not be. *Cognition*, **13**:263–308.
- ASLIN, R., J. SAFFRAN, and E. NEWPORT. 1998. Computation of conditional probability statistics by 8-month-old infants. *Psychological Science*, **9**:321–324.
- BOWNS, D. 1999. *The Biology of Mind*. Fitzgerald Science Press, Bethesda, MD.
- CARSTAIRS-MCCARTHY, A. 1999. *The Origins of Complex Language*. Oxford University Press, Oxford, UK.
- CHANDRASEKARAN, B. 1994. Understanding control at the knowledge level. <http://www.cis.ohio-state.edu/~chandra/intelligent-control.pdf>.
- CLAHSEN, H., and M. ALMAZAN. 1998. Syntax and morphology in Williams syndrome. *Cognition*, **68**: 167–198.
- CLARK, A. 2001. *Mindware*. Oxford University Press, Oxford, UK.
- COOPER, R. 1995. Towards an object-oriented language for cognitive modelling. *In* Proceedings of the 17th Annual Conference of the Cognitive Science Society. *Edited by* J. D. Moore and J. F. Lehman. Cognitive Science Society, Pittsburgh, PA, pp. 556–561.
- COWIE, F. 1999. *What's Within: Nativism Reconsidered*. Oxford University Press, Oxford, UK.
- CRAIN, S., and D. LILLO-MARTIN. 1999. *An Introduction to Linguistic Theory and Language Acquisition*. Blackwell, Oxford, UK.
- DEACON, T. 1997. *The Symbolic Species: The Co-evolution of Language and the Brain*. Norton, New York, NY.
- DENNETT, D. 1991. *Consciousness Explained*. Little Brown, Boston, MA.
- ELLIS, R., and G. HUMPHREYS. 1999. *Connectionist Psychology*. Psychology Press, East Sussex, UK.
- FRAWLEY, W. 1997. *Vygotsky and Cognitive Science: Language and the Unification of the Social and Computational Mind*. Harvard University Press, Cambridge, MA.
- FRAWLEY, W., and R. N. SMITH. 2001. A processing theory of alexithymia. *Cognitive Systems Research*.
- GALLISTEL, C. 1990. *The Organization of Learning*. MIT Press, Cambridge, MA.
- GALLISTEL, C. 1998. Symbolic processes in the insect brain. *In* An Invitation to Cognitive Science, 4: Methods, Models, and Conceptual Issues. *Edited by* D. Scarborough and S. Sternberg. MIT Press, Cambridge, MA, pp. 1–51.
- GAZZANIGA, M., R. IVRY, and G. MAGNUN. 1998. *Cognitive Neuroscience: The Biology of the Mind*. Norton, New York, NY.
- HAUSER, M. 1997. *The Evolution of Communication*. MIT Press, Cambridge, MA.
- KARMILOFF-SMITH, A., et al. 1998. Linguistic dissociations in Williams syndrome: Evaluating receptive syntax in on-line and off-line tasks. *Neuropsychologia*, **36**:343–351.

- KIRBY, S. 1999. *Function, Selection, and Innateness: The Emergence of Language Universals*. Oxford University Press, Oxford, UK.
- LANE, R., G. AHERN, G. SCHWARTZ, and A. KAZNIAK. 1997. Is alexithymia the emotional equivalent of blindsight? *Biological Psychiatry*, **42**:834–844.
- LE DOUX, J. 1996. *The Emotional Brain: The Mysterious Underpinnings of Emotional Life*. Simon and Schuster, New York, NY.
- LEWIS, R. L. 1996. Architecture matters: What Soar has to say about modularity. *In Mind Matters*. Edited by D. Steier and T. Mitchell. Erlbaum, Hillsdale, NJ, pp. 75–84.
- LLINÁS, R. 2001. *I of the Vortex*. MIT Press, Cambridge, MA.
- MARCUS, G. 2001. *The Algebraic Mind*. MIT Press, Cambridge, MA.
- MONSELL, S., and J. DRIVER. (eds.). 2000. *Control of Cognitive Processes: Attention and Performance XVII*. MIT Press, Cambridge, MA.
- OATLEY, K., and J. JENKINS. 1996. *Understanding Emotions*. Blackwell, Cambridge, UK.
- PINKER, S. 1999. *Words and Rules*. Basic Books, New York, NY.
- PYLYSHYN, Z. 1996. The study of cognitive architecture. *In Mind Matters*. Edited by D. Steier and T. Mitchell. Erlbaum, Hillsdale, NJ, pp. 51–74.
- SUN, R., E. MERRILL, AND T. PETERSON. 2001. From implicit skills to explicit knowledge: A bottom-up model of skill learning. *Cognitive Science*, **25**:203–244.
- TANENHAUS, M., and M. LUCAS. 1987. Context effects in lexical processing. *Cognition*, **25**:213–234.
- TANENHAUS, M. G. DELL, and G. CARLSON. 1987. Context effects in lexical processing: A connectionist perspective on modularity. *In Modularity in knowledge representation and natural-language understanding*. Edited by J. Garfield. MIT Press, Cambridge, MA, pp. 87–108.
- TAYLOR, G. J., R. M. BAGBY, and J. D. A. PARKER. 1997. *Disorders of Affect Regulation*. Cambridge University Press, Cambridge, UK.
- TOMASELLO, M., and J. CALL. 1997. *Primate Cognition*. Oxford University Press, Oxford, UK.