categories rather than to perceptual similarity (e.g., light green may be perceptually closer to yellow than to dark green but has to be sorted with the greens); or matching objects by type rather than perceptual appearance, as for example when an analogue clock has to be matched with a digital clock rather than a (visually more similar) compass; or matching images of objects with their characteristic sound (Vignolo 1990). By contrast, matching tasks that require consideration of variations within a category such as, for example, matching of individual faces, do not crucially depend on left hemisphere integrity (Benton & Van Allen 1968).

There are symptoms of LBD, which on first sight, do not fit into a left-hemisphere dominance for extraction and combination of finite elements. These are "high level" disorders of motor control traditionally termed "apraxia." These symptoms have led to the proposal that left-hemisphere dominance concerns primarily motor control. Attempts to deduce language dominance from motor dominance have either emphasized the motor demands of speaking (Kimura 1983) or postulated that language evolved from gestural communication (Corballis 2002). Recent research suggests that apraxia has more to do with the application of combinatorial systems of finite elements than with motor control. Apraxia affects three domains of actions: imitation of gestures, performance of meaningful gestures on command, and use of tools and objects. Evidence has been provided that LBD patients fail imitation of novel gestures because they cannot reduce them to combinations of a limited number of defined body parts (Goldenberg 1996; Goldenberg & Strauss 2002). They have similar problems when this body part coding is required to match photographed gestures (Goldenberg 1999) or to replicate gestures on a mannequin (Goldenberg 1995), although motor control is trivial for pointing to photographs and very different from imitation for manipulating a mannequin. By contrast, the exclusive role of LBD is mitigated or vanishes completely when imitation puts fewer demands on body-part coding and requires instead fine-grained distinctions within one category of body parts (e.g., the fingers of one hand). Performance of meaningful gestures to command is frequently tested by asking for a pantomime of object use (e.g., "Show me how you would use a toothbrush"). Here the crucial difficulty of LBD patients seems to concern the demonstration of the object and its use by selecting distinctive features of the motor action associated with that use (Goldenberg et al. 2003). Use of tools and objects poses demands on many cognitive functions and can be impaired by brain lesions in many locations (Schwartz et al. 1999), but one component which is exclusively bound to left hemisphere integrity is the inference of possible functions from structural properties of objects. For example, LBD patients may fail to discover that a hook can be fixed to a ring by inserting it (Goldenberg & Hagmann 1998). Such failures can be attributed to an inability to detect a limited number of functionally relevant features and to solve mechanical problems by reducing them to basic functional relationships.

There is controversy concerning whether the co-occurrence of these difficulties with aphasia in LBD patients is a result of similarities between the affected functions or of anatomical contiguity between their neural substrates, but this opposition may be ill-conceived. Anatomical contiguity is unlikely to have arisen from arbitrary placement of unrelated functions. Presumably it reflects a deeper affinity of their neural substrate. It may be more fruitful to ask for the functional properties corresponding to this neural commonality. I propose that this commonality is to be sought in the ability to recognize a limited number of finite elements in manifold perceptual entities, and to combine them for reconstructing manifold entities. In this account, the neurally designed predisposition for language acquisition is not specific for language but also supports a range of nonverbal capacities.

Jackendoff's conceptualism

James Higginbotham

School of Philosophy, University of Southern California, Los Angeles, CA 90089-0451. **higgy@usc.edu**

Abstract: In this commentary, I concentrate upon Ray Jackendoff's view of the proper foundations for semantics within the context of generative grammar. Jackendoff (2002) favors a form of internalism that he calls "conceptualism." I argue that a retreat from realism to conceptualism is not only unwarranted, but even self-defeating, in that the issues that prompt his view will inevitably reappear if the latter is adopted.

In *Foundations of Language: Brain, Meaning, Grammar, Evolution* (henceforth *Foundations*), Jackendoff is sympathetic – more sympathetic than I, for one, would have expected him to be – to the view that the theory of meaning in empirical linguistics should link language to human action and communication, and that the notions of reference and truth are indispensable both as explaining relations of sentences to one another, as in implication, and their relations to their subject matter and conditions on their use. Jackendoff holds, however, that the proper implementation of this view requires the adoption of a variety of irrealism about what we refer to, and what makes what we say true or false. In Part III of *Foundations* he offers a variety of reasons for this irrealism, or conceptualism, as he calls it. None of these seem to me effective; I will consider a few below. More than this, however: Jackendoff's irrealism threatens to be self-defeating, in that the problems that he discerns for realist accounts are bound to return, in just the same form, under the interpretation of reference that he offers.

Having remarked, in my view rightly, that the signal contribution of generative grammar was to take for the subject of linguistics not the formal properties of language but rather the basis for human knowledge and capacity for language, Jackendoff is wary (to the point of abhorrence) of saying that languages themselves are abstract objects whose properties we know (or "cognize," to use Chomsky's suggestion of a more neutral terminology). He is wary of this, not because he rejects the notion of implicit or tacit knowledge, but rather because he thinks that, once we say that languages are abstract, we have cut ourselves off from the psychological investigation that is to be the core of the enterprise (p. 297). He is also repelled (p. 299) by the idea that these abstract objects have always been lying around, waiting for people to "grasp" them. Abstract objects in general, he thinks, must be "human creations."

The conflicts here are illusory, however. What comes to hold only through human organization and activity is not the existence of abstract objects, but empirical identities: That language L has property P, may be a fact on a par with the truths of arithmetic; but that Higginbotham's language or Jackendoff's language = L, and therefore that Higginbotham's language or Jackendoff's language has property P, is a psychological contingency, to which all the available evidence, about them and other humans, is relevant. I suppose we may agree that a primitive mechanism of "grasping" is, if true, a counsel of despair. But how is the slogan that abstract objects are "human creations" supposed to help? Everyone knows on a moment's reflection that to enclose the largest area with a piece of string, you should form it into a circle. Supposing that circles are human creations brings us no closer to an explanation of why this should be so.

Jackendoff opposes what he calls common-sense realism about reference – according to which (simplifying only a bit) words refer to things – to his own conceptualist account, according to which speakers judge words to refer to things in "the world as conceptualized" by them. The basis for the substitution of the conceptualist view for the standard one is a variety of questions about reference given in Chapter 10, section 3, (pp. 300–303). All our old friends are there: Sherlock Holmes, the unicorn in my dream, the value of my watch, virtual squares, "politically constructed entities" such as Wyoming, and so forth. There is no space here to consider all of these, but I make two remarks.

https://doi.org/10.1017/S0140525X03360157 Published online by Cambridge University Press

First, Jackendoff ignores the point that nominal reference must be considered, not in isolation, but in the context of a sentence. Thus, take the value of my watch, or the distance between New York and Boston. The things are identified with "mixed numbers," \$50 for my watch, 204 miles for the distance. However, as observed originally by Carnap (1926), any mystification about them disappears when we observe that the reference is to the value of my watch in dollars, which is just the number 50, or the distance between New York and Boston in miles, which is 204. The virtual square formed by four properly placed dots, Jackendoff says, "isn't there physically." True enough, there are no lines, only the dots that are to be construed as the vertices of the square. There is, however, within the limits of perceptual accuracy, exactly one square of which they are the vertices, so that to say that the square is "formed by the four dots" indicates, not that there is no square, but rather how we are to understand the notion "x is *formed by* y."

Second, and more critically, Jackendoff urges (p. 304) that "reference" need not be to things in the "*real* world" (his emphasis). So statements about Sherlock Holmes or the unicorn in my dream can be taken on a par with statements about Derek Jeter or the unicorn in the garden. But this, I think, conceals a mistake: The distinction between names like *Sherlock Holmes* on the one hand, and *Derek Jeter* on the other, is a distinction that is made *within* our speech, not outside it. If you think there is a serious question whether Sherlock Holmes ever visited Mongolia, or that what is responsible for the truth of the statement that he lived in London is the same sort of thing that is responsible for the truth of the statement that Derek Jeter lives in New York, then you don't understand the name; for it is part of understanding the name *Sherlock Holmes* that Sherlock Holmes is a fictional character. The case is similar with dream-objects, but rather more interesting, since one could believe that one saw them, interacted with them, and so forth; but for us anyway there is common recognition that statements about some of the contents of dreams are made true or false in virtue of our representations alone, so that their superficial grammatical form is not a guide to their truth conditions. But the truth conditions are known, and known to be different from those of apparently similar statements. It is, therefore, no advance, and in fact an obscuring of the issues, to adopt for these reasons a conceptualist semantics. Jackendoff's thought seems to be that, if we are casual enough about objects of reference, on the ground that they are in the merely conceptualized world, the problems of distinguishing the truth conditions of fictions, some statements about dreams, and so forth will go away. But they won't.

The latter part of *Foundations* involves often very interesting discussion of semantic phenomena, both lexical and combinatoric. None of these, so far as I can see, require making a distinction such as he envisages between the "conceptualized world" and – the world.

Four challenges for cognitive neuroscience and the cortico-hippocampal division of memory

Harry Howard

Department of Spanish and Portuguese, Tulane University, New Orleans, LA 70118. **howard@tulane.edu** **http://www.tulane.edu/~howard/**

Abstract: Jackendoff's criticisms of the current state of theorization in cognitive neuroscience are defused by recent work on the computational complementarity of the hippocampus and neocortex. Such considerations lead to a grounding of Jackendoff's processing model in the complementary methods of pattern analysis effected by independent component analysis (ICA) and principle component analysis (PCA).

Jackendoff elaborates four challenges for cognitive neuroscience whose consequences reverberate throughout his book, *Foundations of Language* (Jackendoff 2002). In a nutshell, if spreading activation (SA), firing synchrony of neural units (FS), and multilayer perceptrons trained by back-propagation of error (BP) constitute the apogee of current neurotheory, then it has a long way to go to reach even the lowest echelons of descriptive adequacy for human language. In this commentary, I briefly review a neurologically realistic alternative to the SA/FS/BP trio that meets most of Jackendoff's challenges. Known as the Complementary Learning Systems (CLS) model, it was first developed by McClelland et al. (1995), and has been refined several times since then (see O'Reilly & Norman 2002). Its computational principles have been applied to a wide range of learning and memory phenomena (impaired and preserved learning capacities with hippocampal lesions in conditioning, habituation, contextual learning, recognition memory, recall, and retrograde amnesia) across several species (rats, monkeys, and humans). To explain how CLS works, we start at the end, with Jackendoff's fourth challenge.

Jackendoff sees it as contradictory that a compositional phrase, such as "lift the shovel" should be encoded in short-term memory via SA or FS, while a structurally-equivalent idiomatic phrase such as "kick the bucket" should be stored in long-term memory by the slow modulation of synaptic weights via BP. The CLS literature implicitly raises a comparable objection, which is resolved as the computational difference between hippocampal and neocortical function. By way of illustration, let us call on a linguistic example that neither Jackendoff nor CLS discusses, but which has considerable empirical depth (see Bowerman 1996; Bowerman & Choi 2001; Choi & Bowerman 1991).

Imagine a child viewing two events, one in which a cassette is put in a bag and another in which the same cassette is put in its case. Korean, in contrast, lexicalizes the events with separate items, namely, the verbs *nehta,* "put loosely in or around" for the former and *kkita,* "interlock, fit tightly" for the latter. Thus, the brain must keep both events separate, and presumably with their full complement of real-world detail, in order to account for the specificity of Korean. Nevertheless, both events have overlapping parts, such as the cassette, the motion of bringing two things together, and maybe even the person performing the motion. The brain must therefore ensure that parts shared among events do not interfere with one another. CLS asserts that these characteristics define episodic memory and the function of the hippocampus: the fast and automatic learning of sparse representations. Nevertheless, if all events are kept separate, there will be no way to generalize across them. Yet humans do indeed generalize; witness the fact that English uses the same spatial vocabulary for both events, namely, the preposition *in,* with the aid of the motion verb *put.* The brain must therefore be able to integrate events so as to abstract away from their specifics and encode the overall statistical structure of the environment. CLS asserts that these characteristics define semantic memory and the function of neocortex: the slow and task-driven learning of overlapping representations.

Returning to as "lift the shovel" versus "kick the bucket," we may conclude that the hippocampus makes the initial, short-term binding of the disparate features of the phrases from which the neocortex extracts any statistical regularities, such as the parallel [V [Det N]] structures. The idiomatic phrase sports an additional regularity, namely, the fact that it has a noncompositional reading which presumably can *only* be learned by the slow (i.e., multiple exposure) modulation of synaptic weights. The conclusion is that CLS avoids any inconsistent treatment of compositional and noncompositional phrases.

Turning to the first challenge, Jackendoff cites the multiple embedding of linguistic entities as leading to temporal incoherence for any solution to the feature-binding problem that relies on FS. The CLS model has made this same criticism from a more neurologically informed perspective. Its alternative is to return to the notion of conjunctive features, with a twist. The twist is to avoid the combinatorial explosion of units encoding a single feature conjunction by distributing the conjunctions across many units (O'Reilly & Busby 2002),

> where each unit encodes some possibly-difficult to describe amalgam of input features, such that individual units are active at different levels

https://doi.org/10.1017/S0140525X03360157 Published online by Cambridge University Press