

(2) *Objectivity*. Logic, geometry, and mathematics are not uninterpreted formal systems that people happen to universally assent to regardless of which community they inhabit. Formal interpretations of physical phenomena permit predictions concerning the behavior of objective reality even in contexts vastly beyond the scope of actual (or possible) human experience. How then does mathematical reasoning manage to preserve truth about distant contexts if mathematical objects are merely psychological data structures with local inferential features? In other words, quite apart from its universality, how, in the psychologicist account, does mathematics come by its objectivity (cf. Smith 1996)?

(3) *Error*. It is tempting to account for the validity of logical inference in terms of the way that (normal, healthy) cognitive systems actually reason. But we can make mistakes regarding the properties of abstract objects. Even professional mathematicians occasionally draw false inferences about mathematical objects. And a real feeling of surprise and discovery can accompany mathematical innovation, that moment when humanity discovers that we have been conceiving of some mathematical construct incorrectly all along. The intuition that mathematical objects can have properties quite different from those imputed to them, even by professionals, fuels Platonist intuitions (Dummett 1978). Validity cannot merely consist in conformity with the way people actually reason; it is a property of arguments that conform to the way we *ought* to reason. How psychologism can account for this remains uncertain.

Jackendoff (pp. 330–32) suggests several mechanisms of social “tuning” that can serve to establish (universal) norms within a community – norms against which error may be judged and the appearance of objectivity can arise. So when Joe mistakes a platypus for a duck (p. 329), his error is relative to the impressions of the rest of his community. “Objective” fact and the appearance of universality is established by community consensus. Unfortunately, this account does quite poorly with logic and mathematics. A mathematical or logical discovery happens when one member of the community realizes that something is wrong with the way the community conceptualizes some aspect of the field, and demonstrates that error to the other members of the community. The issue here is how a whole community can be shown to be in error when the objective reality against which the error is judged is mere community consensus. Platonism has an obvious solution to this issue, but CS will have to work for one.

We are by no means arguing that universality, objectivity, and error cannot be accommodated by CS. But Jackendoff does suggest that CS can provide insight into the appeal of formal approaches to semantics. Before it can explain the success of its rival, it must itself account for the nature of the logical apparatus on which formal work rests. We suspect that this can indeed be done. But until it is, CS remains incomplete in an important way.

ACKNOWLEDGMENT

We thank Robert Stainton for his helpful comments and discussion.

Delegation, subdivision, and modularity: How rich is Conceptual Structure?

Damián Justo^a, Julien Dutant^b, Benoît Hardy-Vallée^c, David Nicolas^d, and Benjamin Q. Sylvand^e

^{a,b,c,d,e}Institut Jean Nicod, Centre National de la Recherche Scientifique – École des Hautes Études en Sciences Sociales – École Normale Supérieure (CNRS – EHESS – ENS), 75007 Paris, France; ^aÉcole des Hautes Études en Sciences Sociales, 75006 Paris, France; ^{b,e}U.F.R. de Philosophie, Sorbonne Paris IV, 75005 Paris, France; ^bÉcole Normale Supérieure – Lettres et Sciences Humaines, 69342 Lyon, France; ^cDépartement de Philosophie, Université du Québec à Montréal, Montréal, Québec H3C 3P8, Canada. aledam@noos.fr Julien.Dutant@paris4.sorbonne.fr benoithv@iquebec.com dnicolas@gmx.net Benjamin.Sylvand@paris4.sorbonne.fr <http://www.institutnicod.org> <http://benoithv.free.fr> <http://d.a.nicolas.free.fr/research>

Abstract: *Contra* Jackendoff, we argue that within the parallel architecture framework, the generality of language does not require a rich conceptual structure. To show this, we put forward a delegation model of specialization. We find Jackendoff’s alternative, the subdivision model, insufficiently supported. In particular, the computational consequences of his representational notion of modularity need to be clarified.

In Jackendoff’s framework in *Foundations of Language* (2002), understanding the meaning of a sentence consists in constructing a representation in a specific cognitive structure, namely, Conceptual Structure (CS). CS is not dedicated to language, though. It is the structure that carries out most of our reasoning about the world. According to Jackendoff, this follows from what we call the Generality of Language Argument (GLA):

1. Language allows us to talk about virtually anything.
2. Every distinct meaning should be represented within CS.
3. CS must contain our knowledge about everything it represents.
4. Hence, CS contains large bodies of world knowledge: CS is “rich.”

For instance, if the difference between “to murder” and “to assassinate” is that the second requires a political motive, then CS contains knowledge about what it is to be a political motive (Jackendoff 2002, p. 286).

GLA excludes the idea that there is a specifically linguistic level of semantics, containing only a “dictionary meaning” as opposed to “encyclopedic information” (Jackendoff 2002, p. 285). It also excludes a minimal view of CS. We call *minimal* a CS that is able to represent all distinct meanings, but is not able to carry out computations other than the logical ones. A minimal CS could represent the meanings of “x is an elephant” and “x likes peanuts,” but would not be able to infer the second from the first.

We think that GLA is wrong: The generality of language is compatible with a minimal CS. Indeed, it is a viable possibility within Jackendoff’s general architecture of the mind. Consider the sentence: “The elephant fits in the mailbox.” To know that it is wrong is to represent its meaning and judge it to be false. Jackendoff would say that these two steps are carried out by different structures, namely, CS and Spatial Structure (SpS). Since only CS interacts directly with language, the sentence has to be translated into CS. From there it can in turn be translated into a representation in SpS. This would be done by dedicated interfaces. SpS is the place where the sentence is found false, for it is impossible to create a well-formed spatial representation of an elephant in a mailbox. We regard this as an instance of a delegation model:

(DM) Domain-specific computations are carried out outside CS, but their result is represented in CS, and may thus be expressed in language.

In this case the computation is very simple. It consists of checking whether an adequate SpS representation can be formed. Nevertheless, it is done outside CS. CS only represents its result, namely that the elephant does not fit in the mailbox.

It is a priori possible that DM applies to all the computations involved in our knowledge about physical objects, biological kinds,

other minds and so on. The resulting CS would be minimal. Hence, premise (3) is false: CS *could* represent meanings without containing world knowledge.

Jackendoff does not address this question. Instead, he directly proposes an alternative model for specialization. For instance, he takes social cognition as involving a specialized mental structure. But he claims that this is a substructure of CS, a “sub-specialization” (Jackendoff 1992a, Ch. 4). We call this the *subdivision model*:

(SM) Domain-specific computations are carried out within parts of CS, and can thus be expressed in language.

If most of our reasoning about specific domains has to be carried out within parts of CS, then CS has to be rich. But why should it be so? Jackendoff could put forward two distinct hypotheses.

The *computational unity hypothesis* claims that CS is a computational module, with a *unique processor*, and that sub-specializations are representational modules, that is, knowledge bases about specific domains.¹ On this hypothesis, domain-specific inferences are construed as logical inferences based on domain-specific premises and effected by a single processor, and this is why they are part of CS. However, such a claim is far from being uncontroversial. Many cognitive psychologists argue that putative “sub-specializations” such as Theory of Mind, carry out their computations independently of each other in a relatively autonomous way, and are possibly situated in distinct, dedicated neural structures (Leslie 1994; Segal 1996). Moreover, if the processor were damaged, it seems that one would lose all propositional computational abilities at once. But this pathology has not been observed.

A weaker hypothesis is that of a *unique representational format*. Jackendoff (2002, p. 220) seems to endorse it. It merely claims that all sub-specializations of CS share a common, propositional format and that all corresponding computations are of a quantificational-predicational character. Their computations need not be carried out by a common processor. However, we do not think that this view has any more plausibility than the hypothesis that some sub-specializations have their computations carried out in *sui generis* formats that are designed for the tasks that they solve. Our understanding of each other’s minds plausibly involves propositional representations, but this may be the exception rather than the rule. Moreover, it is not clear whether CS would, in this view, constitute a module in any interesting sense, or whether the hypothesis really differs from generalized delegation and a minimal CS.

To conclude, within Jackendoff’s architecture of the mind, the generality of language is compatible with either a rich or a minimal CS. The choice of the former requires that the computational consequences of Jackendoff’s representational notion of modularity be at the very least clarified.

ACKNOWLEDGMENTS

Thanks to Roberto Casati for setting up a workshop on Jackendoff’s work, and to Ray Jackendoff for discussing issues related to the present argument. Justo acknowledges support by CONICET and Fundación Antorchas, Argentina.

NOTE

1. For further discussion of representational (or intentional) and computational modularity, see Segal (1996).

Neuropsychological evidence for the distinction between grammatically relevant and irrelevant components of meaning

David Kemmerer

Department of Audiology and Speech Sciences and Department of Psychological Sciences, Purdue University, West Lafayette, IN 47907-1353. kemmerer@purdue.edu

Abstract: Jackendoff (2002) argues that grammatically relevant and irrelevant components of meaning do not occupy distinct levels of the semantic system. However, neuropsychological studies have found that the two components doubly dissociate in brain-damaged subjects, suggesting that they are in fact segregated. Neural regionalization of these multidimensional semantic subsystems might take place during language development.

Jackendoff’s *Foundations of Language* is, without a doubt, a monumental achievement. It both clarifies and begins to fulfill the deeply pressing need for integration not only within linguistics but also between linguistics and the connected disciplines of psychology, neuroscience, and evolutionary biology.

Here I concentrate on the relation between linguistics and neuroscience. Although Jackendoff points out that a great deal has been learned about the functional organization of various aspects of language in the brain, he doesn’t devote much space to exploring how these findings can shed light on current issues in linguistic theory. To illustrate the potential applications of recent neuro-linguistic research, I present an example that bears directly on two theoretical topics that are near to Jackendoff’s heart: the syntax-semantics interface, and the basic architecture of the semantic system.

As Jackendoff observes, many linguists have been converging on the notion that grammatical constructions consist of morphosyntactic patterns that are directly associated with schematic meanings; and, in order for a word to occur in a given construction, its own meaning must be compatible with that of the construction (Goldberg 2003). Consider the well-known locative alternation:

- (1) a. *Sam sprayed water on the flowers.*
 b. *Sam dripped water on the flowers.*
 c. **Sam drenched water on the flowers.*
- (2) a. *Sam sprayed the flowers with water.*
 b. **Sam dripped the flowers with water.*
 c. *Sam drenched the flowers with water.*

The construction in (1) has the broad-range meaning “X causes Y to go to Z in some manner,” whereas the one in (2) has the broad-range meaning “X causes Z to change state in some way by adding Y”; each construction also has a network of more restricted narrow-range meanings that are essentially generalizations over verb classes (Pinker 1989). *Spray* can occur in both constructions because it encodes not only a particular manner of motion (a substance moves in a mist) but also a particular change of state (a surface becomes covered with a substance). However, *drip* and *drench* are in complementary distribution, for the following reasons. One of the narrow-range meanings of the first construction is “X enables a mass Y to go to Z via the force of gravity,” and this licenses expressions like *drip/dribble/pour/spill water on the flowers* and excludes expressions like **drench water on the flowers*. Similarly, one of the narrow-range meanings of the second construction is “X causes a solid or layer-like medium Z to have a mass Y distributed throughout it,” and this licenses expressions like *drench/douse/soak/saturate the flowers with water* and excludes expressions like **drip the flowers with water*.

According to the Grammatically Relevant Subsystem Hypothesis (GRSH), a fundamental division exists between, on the one hand, semantic features that determine the compatibility between verb meanings and constructional meanings, and on the other, semantic features that capture idiosyncratic nuances of verb meanings, for example, the featural distinctions between *drip*, *dribble*,