> for different inputs, and many such units are active for each input (Hinton et al. 1986). Therefore, the input is represented by a complex distributed pattern of activation over units, and each unit can exhibit varying levels of sensitivity to the featural conjunctions present in the input. (O'Reilly & Busby 2002)

The conclusion is that Jackendoff's objection is vindicated, and resolved.

The "problem of 2" presumably disappears with a distributed encoding, since each repeated item would be represented by different levels of activation of the neural population encoding the feature conjunction according to the item's context. However, the current CLS literature does not address this issue directly, leaving the reader uncertain whether the approach will scale up correctly.

Finally, the derivation of typed variables is a goal of CLS, in the guise of the learning of relational structures. Its supporters echo Jackendoff's reiterated protestation that freely-combining typed variables are fundamental to higher cognition. Unfortunately, the CLS simulations in which a relational structure is learned suffer from the general opaqueness of distributed encoding. That is to say, the network appears to have learned the relations that it was exposed to, but from the text of the reports, one does not understand how it is done.

In this respect, the work of two other researchers is highly relevant. Jackendoff cites Shastri and Ajjanagadde (1993) as one of the few computational models that grapples with the representation of typed variables. Shastri's more recent model SMRITI (Shastri 2002) takes these ideas a step further by specifically attributing to the hippocampus the responsibility for creating role-entity bindings, so that it will assign to an event in which John gives Mary a book in the library on Tuesday, the representation *GIVE: giver = John, recipient = Mary, give-object = a book, location = library, temporal-location = Tuesday.* This is tantamount to the relational structure of first-order predicate logic, if not Event Semantics (see Parsons 1990). In a slightly different vein, Pulvermuller (2002) reviews and expands on the concept of syn-fire chains as a neurologically plausible mechanism for serial order in language. Such chains explicitly encode the relational structure of syntax, though Pulvermüller does not localize them to any particular cortical area. In fact, Pulvermüller provides a fascinating neuroscientifically grounded foil to Jackendoff, and reading the two of them together is a rewarding intellectual exercise. More to the point, both Shastri and Pulvermüller wind up invoking freely-combining typed variables in a way that is more transparent than CLS.

So what of the other 418 or so pages of Jackendoff's text? The various modules of the grammar-based processing architecture presumably reflect independent clusters of statistical regularities in the linguistic input, learned by the gradual adjustment of synaptic weights. This suggests a further distillation of the CLS: The hippocampus performs independent component analysis (ICA) on its input patterns in order to orthogonalize them, that is, to remove their common features and so make them maximally unrelated (Kempermann & Wiskott 2004). At the very least, this would separate a linguistic pattern into its phonological, syntactic, and conceptual components, and then into the independent subcomponents thereof, such as Shastri's role-entity bindings in the conceptual module. The drawback of ICA is that it separates linguistic patterns into an enormous number of dimensions, for example, *temporal-location* = Monday, temporal-location = Tuesday, and so on. It is the function of the neocortex to reduce these dimensions to the most relevant or informative ones, which suggests that neocortex performs principal component analysis (PCA, see Simoncelli & Olshausen 2001) on the hippocampal output. For instance, the independent temporal components mentioned two sentences ago could be reduced to the single principal component *temporal-location* = day of the week, thereby synthesizing a typed variable.

Psychologism and conceptual semantics

Luke Jerzykiewicz[a] and Sam Scott[b]
[a]*Cognitive Science Department, Carleton University, Ottawa K1S 5B6, Canada;* [b]*Philosophy-Neuroscience-Psychology Program, Washington University in St. Louis, St. Louis, MO 63130.* **ljerzyki@uwo.ca**
sscott@artsci.wustl.edu **http://www.carleton.ca/~ljerzyki**
http://www.artsci.wustl.edu/~sscott

Abstract: Psychologism is the attempt to account for the necessary truths of mathematics in terms of contingent psychological facts. It is widely regarded as a fallacy. Jackendoff's view of reference and truth entails psychologism. Therefore, he needs to either provide a defense of the doctrine, or show that the charge doesn't apply.

Jackendoff's vision of the language faculty in *Foundations of Language: Brain, Meaning, Grammar, Evolution* (2002) is impressive in scope and rich in insightful detail. However, his account of abstract objects (sect. 10.9.3) requires substantial elaboration and defense before it can safely avoid the fallacy of psychologism.

In Jackendoff's account, conceptual structures within the generative semantic module are not themselves interpreted – they do not *have* a semantics. They just *are* the semantics of natural language. The fine-grained data-structures that inhabit the semantic module interface richly with perceptual modalities and with motor outputs, while individually not necessarily representing anything in the world as such (cf. Brooks 1991). The familiar appearance that words refer to entities and events can be explained – for concrete referents, at least – in terms of the relationship between semantic constructs and the outputs of perceptual faculties. It is these outputs that we experience as our "world." Now, in the case of abstract objects (like beliefs, mortgages, obligations, and numbers), which manifestly lack perceptual features, the theory makes only slightly different provisions: The data-structures that encode them possess inferential rather than perceptual features. Interfaces to syntax and phonology treat all conceptual structures similarly, regardless of whether their constitutive features are exclusively inferential or, in part, perceptual. Thus, Jackendoff's naturalistic theory of concepts rejects Platonism and identifies abstract objects with the cognitive structures that express them.

The paradigm cases of abstract objects are mathematical and logical entities. Oddly, Jackendoff does not discuss these cases explicitly. Yet if the Conceptual Semantics (CS) account of abstract objects is to work at all, it must work for them. The trouble is that CS entails *psychologism,* the view that the necessary truths of mathematics and logic are to be accounted for in terms of contingent facts about human cognition. According to psychologism, $2 + 2 = 4$ is a fact of human psychology, not a fact that is independent of human beings. Frege (1953) raised seminal objections to this doctrine and today psychologism is typically viewed as a patent fallacy (Dartnall 2000). There is room for discussion, however. Haack (1978) points out that it is far from obvious whether Frege's objections continue to apply. Frege's target was the introspectionist psychology of the day, and Jackendoff (1987; 2002) carefully avoids this approach. But to get off the ground, a psychologistic account of abstract entities must cope with three challenges:

(1) *Universality.* Some norms derive their authority from community standards. Those norms are no less real for their conventional nature (traffic rules come to mind), but they are only true by agreement. By way of contrast, norms governing the behavior of abstract logical and mathematical entities are *universal* (a point stressed by Nagel 1997). Community standards derive their authority from these norms, and not vice versa. Even people with untutored intuitions can come to recognize the truth of a law of logic or mathematics, though they may require quite a bit of reflection to do so. CS needs an explanation of how some abstract objects (which are supposed to be mental entities) come to possess these inferential features. Are they innate? If so, Jackendoff's rejection of Fodor loses some of its bite. Are they learned? If so, the poverty of stimulus problem rears its ugly head.

(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 psychologistic 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;* [a]*École des Hautes Études en Sciences Sociales, 75006 Paris, France;* [b,e]*U.F.R. de Philosophie, Sorbonne Paris IV, 75005 Paris, France;* [b]*École Normale Supérieure – Lettres et Sciences Humaines, 69342 Lyon, France;* [c]*Département 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,