

Contemporary Debates in Philosophy

In teaching and research, philosophy makes progress through argumentation and debate. *Contemporary Debates in Philosophy* presents a forum for students and their teachers to follow and participate in the debates that animate philosophy today in the western world. Each volume presents opposing viewpoints on contested themes and topics in the central subfields of philosophy, and is edited and introduced by an expert in the field. The opposing essays, commissioned especially for the volumes in the series, are thorough but accessible presentations of opposing points of view.

1. *Contemporary Debates in Philosophy of Religion* edited by *Michael L. Peterson and Raymond J. VanArragon*
2. *Contemporary Debates in Philosophy of Science* edited by *Christopher Hitchcock*
3. *Contemporary Debates in Epistemology* edited by *Matthias Steup and Ernest Sosa*
4. *Contemporary Debates in Applied Ethics* edited by *Andrew I. Cohen and Christopher Heath Wellman*
5. *Contemporary Debates in Aesthetics and the Philosophy of Art* edited by *Matthew Kieran*
6. *Contemporary Debates in Moral Theory* edited by *James Dreier*
7. *Contemporary Debates in Cognitive Science* edited by *Robert J. Stainton*

Forthcoming *Contemporary Debates* are in:

Metaphysics edited by *Ted Sider, Dean Zimmerman, and John Hawthorne*
Philosophy of Mind edited by *Brian McLaughlin and Jonathan Cohen*
Social Philosophy edited by *Laurence Thomas*
Political Philosophy edited by *Thomas Christiano*
Philosophy of Language edited by *Ernie Lepore*

Contemporary Debates in Cognitive Science

Edited by

Robert J. Stainton

© 2006

 **Blackwell**
Publishing

CHAPTER THIRTEEN

Locating Meaning in the Mind (Where it Belongs)

Ray Jackendoff

1 Premises of Conceptualist Semantics

This chapter is drawn from a larger project, *Foundations of Language* (Jackendoff, 2002), whose goal is to reconstruct linguistic theory in a way that permits a better integration of linguistics with psychology, neuroscience, and evolutionary theory – and a better integration of the parts of linguistics with each other. In particular, unlike recent mainstream generative linguistics (e.g., Chomsky, 1995), it is intended to include all of language, not just an artificially stipulated core. The reader is referred especially to chapters 9 and 10 for more detail of the material discussed here.

The most basic aspect of generative grammar, the tradition of linguistics within which I work, is that it is *mentalist*: it takes linguistic theory to be modeling the structural organization in the brain of a language user. It is this ground rule that places interesting constraints on the theory of grammar: First, language must be stored and processed in the brain. Hence theoretical elegance equals not mathematical elegance but rather “brain elegance,” however that may come to be characterized. Second, language must be learned. Hence we need a theory of the human language capacity, or Universal Grammar, and how it interacts with experience to create adult language competence. Third, the language capacity must have evolved since humans diverged from chimpanzees about 5 million years ago, which places constraints on what can be attributed to Universal Grammar.

One missing part in standard generative grammar and most of its offshoots is a good compatible theory of meaning, one that is comparably mentalistic. To take the two most influential approaches to meaning within linguistics, Cognitive Grammar (Lakoff, 1987; Langacker, 1987) is mentalistic but mostly isn't made to be compatible with generative grammar; and formal semantics (Chierchia and McConnell-Ginet, 1990) is mostly apsychological, so that it certainly doesn't connect with psycholinguistics and language acquisition. I've been trying to work out a more satisfactory

approach for about 30 years. Here I want to go through some aspects of it. Some date back to older work (Jackendoff, 1976, 1983, 1987); some are new to *Foundations of Language*.

I take the basic problem of a mentalistic semantic theory to be:

- How can we characterize the messages/thoughts/concepts that speakers express/convey by means of using language?
- How does language express/convey these messages?

I leave the terms "messages/thoughts/concepts" and "express/convey" deliberately vague for the moment. Part of our job is to sharpen them. In particular, we have to ask:

What makes these mental entities function the way meanings intuitively should?

Unfortunately, the intellectual politics begins right here: this is not the way everyone construes the term "semantics." Rather than engage in arguments based on terminological imperialism, I will use *conceptualist semantics* as a term of art for this enterprise. (My own particular set of proposals, which I have called Conceptual Semantics (Jackendoff, 1990), is an exemplar of the approach but not the only possible one.) Above all, I don't want to get trapped in the question: Is this enterprise really a kind of semantics or not? The relevant questions are: Is this enterprise a worthwhile way of studying meaning? To what extent can it incorporate intuitions and insights from other approaches, and to what extent can it offer insights unavailable in other approaches? It is important to see that the messages/thoughts/concepts conveyed by language serve other purposes as well. At the very least, they are involved in the cognitive processes illustrated in figure 13.1.

Linguists spend a lot of time accounting for the combinatoriality of phonology and syntax. But it is assumed (although rarely articulated) that these systems serve the purpose of transmitting messages constructed from an equally combinatorial system of thoughts: a sentence conveys a meaning built combinatorially out of the meanings of its words. This combinatorial system is represented in figure 13.1 by the component "formation rules for thoughts," which defines the class of possible thoughts or conceptual structures. (An important terminological point: one use of the term *syntax*

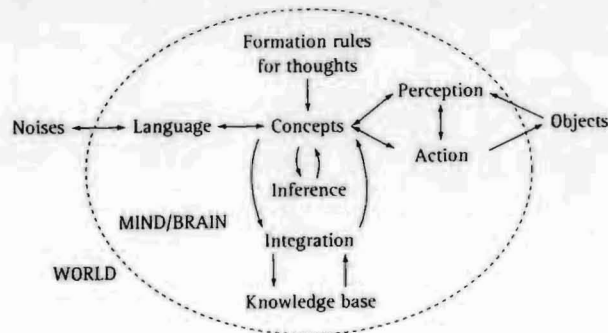


Figure 13.1 The place of conceptualist semantics in the architecture of mind.

pertains to any sort of combinatorial system. In this sense phonology has a syntax, music has a syntax, chess has a syntax, and of course thought has a syntax. In the use favored by linguists, however, *syntax* refers specifically to the combinatorial organization made up of such units as noun phrase and verb phrase. In this sense thought is definitely *not* syntax.)

In the framework of *Foundations of Language* (Jackendoff, 2002), the combinatorial structures for thought are related to the purely linguistic structures of syntax and phonology by so-called *interface rules* (double arrows in figure 13.1). In particular, different languages have different interfaces, so that the same thought can be mapped into expressions of different languages, within tolerances, allowing for the possibility of reasonably good translation among languages. An important part of the interface rules is the collection of *words* of the language. A word is a long-term memory association of a piece of phonology (its pronunciation), a set of syntactic features (its part of speech and contextual properties such as subcategorization), and a piece of conceptual structure (its meaning). Thus each word in an utterance establishes a part of the utterance's sound-grammar-meaning correspondence; other parts of the correspondence are mediated by interface rules that map between syntactic structure and combinatorial structure in semantics.

These two parts of figure 13.1 – the combinatorial system of meaning and its interfaces to linguistic expression – are closest to what is often called "linguistic semantics." Now consider the other interfaces. The use of thoughts/concepts to produce further thoughts/concepts is what is typically called "inference" or "reasoning." Since we are interested in the study of real people and not just ideals, this interface must include not only logical reasoning but also making plans and forming intentions to act, so-called "practical reasoning" (Kahneman et al., 1982; Bratman, 1987; Gigerenzer, 2000) and "social reasoning" (Fiske, 1991; Tooby and Cosmides, 1992).

We also must account for the integration of thoughts conveyed by language with previous knowledge or beliefs. Part of previous knowledge is one's sense of the communicative context, including one's interlocutor's intentions. Thus the work of this interface is closest to what is often called "pragmatics."

The interfaces to the perceptual systems are what permit one to form thoughts based on observing the world. In turn, by using such thoughts as the input to language production, we can talk about what we see, hear, taste, and feel. These interfaces operate in the other direction as well: language can be used to direct attention to some particular part of the perceptual field. The interface with the action system is what permits one to carry out intentions – including carrying out intentions formed in response to a command or request.

In order to make possible the kinds of interactions just enumerated, all these interfaces need to converge on a common cognitive structure. If we look at thought through the lens of language alone, we don't have enough constraints on possible theories. A richer, more demanding set of boundary conditions emerges from insisting that thought must also make contact with inference, background knowledge, perception, and action.

In fact, this view of thought permits us to make contact immediately with evolutionary considerations as well. Suppose we erase the interface to language from figure 13.1. We then have an architecture equally suitable – at some level of approximation – for nonlinguistic organisms such as apes. They too display complex integration of

perception, action, inference, and background knowledge, in both physical and social domains (Köhler, 1927; Goodall, 1971; Byrne and Whiten, 1988; Hauser, 2000). They just can't talk about it. It makes evolutionary sense to suppose that some of the fundamental parts of human thought are a heritage of our primate ancestry (Hauser et al., 2002; Pinker and Jackendoff, 2005).

To presume that we can invoke evolutionary considerations, of course, is also to presume that some of the overall character of thought is determined by the genome. I see at least three major domains of thought that suggest an innate genetic basis. The first is the understanding of the physical world: the identification of objects, their spatial configurations with respect to each other, the events in which they take part and interact, and the opportunities they offer for action on and with them. The second is the understanding of the social world: the identification of persons, characterization of the beliefs and motivations of other persons (so-called "theory of mind"), and the ability to understand the social roles of oneself and others, including such issues as kinship, dominance, group membership, obligations, entitlements, and morals (not as a universal system, but as the underpinning for all cultural variation in social understanding) (Jackendoff, forthcoming). The third domain that I think must be innate is a basic algebra of individuation, categorization, grouping, and decomposition that undergirds both the two systems just mentioned as well as many others.

In short, conceptualist semantics should aspire to offer a common meeting ground for multiple traditions in studying cognition, including not only linguistic semantics but also pragmatics, perceptual understanding, embodied cognition, reasoning and planning, social/cultural understanding, primate cognition, and evolutionary psychology. A high aspiration, but certainly one worth pursuing.

2 Attempts to Integrate the Commonsense View of Reference with Mentalist Linguistics

A crucial part of semantic theory is to explain how reference and truth-value are attached to linguistic expressions. Common sense tells us that linguistic expressions say things about the world. When I say *Russell was a great philosopher*, I am referring to Russell, a real person in the real world. Likewise, sentences intuitively have truth-values by virtue of how they relate to the world. *Snow is green* is false because snow isn't green.

The predominant traditions in Anglo-American semantics and philosophy of language take this commonsense position for granted. They therefore consider it the task of semantic/pragmatic theory to explain how linguistic expressions say things about the world and have truth-values based on their relation to the world. But it is not so easy to fit this commonsense position on reference and truth into a mentalistic theory of language.

Let us first think about how to understand the term "language." Frege and much of the tradition following him take language to be independent of its human users: it relates directly to the world. This combines with the commonsense view of reference to create a view like figure 13.2. More recent variants on this position (e.g., Kripke, 1972; Lewis, 1972; Montague, 1973; Stalnaker, 1984) substitute "possible worlds" and/or "the model" for "world," leaving the basic story unchanged. Language doesn't really




Figure 13.2 An "objectivist" or "realist" view of language.




Figure 13.3 The mind grasping "language in the world."

have much to do with the mind. Some in formal semantics, such as David Lewis, are very explicit about this, others are more or less agnostic.

How can this realist view of language be reconciled with the mentalist approach of generative grammar? One approach would be to jettison the mentalism of generative linguistics, but retain the formal mechanisms: to take the position that there is an objective "language out there in the world," and that this is in fact what generative grammar is studying. Some people have indeed taken this tack (e.g., Katz, 1981). But it disconnects generative linguistics from all sources of evidence based on processing, acquisition, genetics, and brain damage. And it forces us to give up the fundamental motivation for positing Universal Grammar and for exploring its character. I personally think that's too high a price to pay.

An alternative tack might be Frege's (1892): language is indeed "out in the world" and it refers to "objects in the world"; but people use language by virtue of their grasp of it, where "grasp" is a transparent metaphor for "the mind holding/understanding/making contact with" something in the world. Figure 13.3 might schematize such an approach. Generative linguistics, it might then be said, is the study of what is in the mind when it grasps a language. This way we could incorporate all the mentalistic methodology into linguistics while preserving a realist semantics.

But what sense are we to make of the notion of "grasping" an abstract object? We know in principle how the mind "grasps" concrete objects: it constructs cognitive structures in response to inputs from the senses. This is a physical process: the sense organs respond to impinging light, vibration, pressure, and so forth by emitting nerve impulses that enter the brain. But what inputs give rise to the "grasp" of an abstract object? An abstract object by definition has no physical manifestations that can impinge on the nervous system. So how does the nervous system "grasp" them? Without a careful exegesis of the term – which no one provides – we are ineluctably led toward a quasimystical interpretation of "grasping," a scientific dead end.




Figure 13.4 Concepts in the mind that are “about” objects in the world.

One way to eliminate the problem of how the mind grasps language is to push language entirely into the mind – as generative grammar does. We might then arrive at a semantic theory structured like figure 13.4. This is Jerry Fodor’s position. I think (Fodor, 1975, 1983, 1990, 1998): for him, language is a mental faculty that accesses combinatorially structured concepts (expressions in the “Language of Thought”). In turn, concepts have a semantics: they are connected to the world by virtue of being “intentional.” The problem for Fodor is to make naturalistic sense of intentionality. But intentionality suffers from precisely the same difficulty as “grasping” language in figure 13.3: there is no physically realizable causal connection between concepts and objects.

The upshot of these explorations is that there seems to be no way to combine a realist semantics with a mentalist view of language, without invoking some sort of transcendental connection between the mind and the world. The key to a solution, I suggest, lies in examining the realist’s notion of “objects in the world.”

3 Problems with the Commonsense View: “Objects”

Of course, “objects,” “states of affairs,” and “things out in the world” have an entirely intuitive construal. “Objects” invites us to think of Bertrand Russell, Noam Chomsky, trees, refrigerators, and so on. You can see them and touch them; you can count them, tell two of them apart, and in some cases move them around. “States of affairs” invites us to think of snow being white, dogs being animals, Russell’s having been a philosopher, and so on: observable facts about observable objects.

But consider: we refer routinely to all sorts of “objects” that are not so simple to put our hands on. A wide range of examples appears in (1)–(6). I leave it to the reader to see the difficulties each of them raises.



Figure 13.5a



Figure 13.5b

- 1 Fictional and mythical characters
 - (a) Sherlock Holmes
 - (b) the unicorn in my dream last night
- 2 Geographical objects
 - (a) Wyoming
 - (b) the Mississippi River
 - (c) the distance between New York and Boston
- 3 Virtual objects
 - (a) the square formed by the four dots in figure 13.5a
 - (b) the horizontal rectangle that goes behind the vertical one in figure 13.5b
- 4 Social entities
 - (a) the value of my watch
 - (b) the first dollar I ever earned
 - (c) Morris Halle’s PhD degree
 - (d) your reputation
 - (e) General Motors
 - (f) the score of tomorrow’s Red Sox game
- 5 Auditorily perceived objects
 - (a) Beethoven’s Seventh Symphony
 - (b) the words *banana* and *despite*
- 6 Other
 - (a) the set of all possible worlds
 - (b) the best of all possible worlds

Here is the point of these examples: the commonsense view of reference asserts that we refer to “objects in the world” as if this is completely self-evident. It is self-evident, if we think only of reference to middle-sized perceivable physical objects like tables and refrigerators. But as soon as we explore the full range of entities to which we *actually* refer, “the world” suddenly begins to be populated with all sorts of curious beasts whose ontological status is far less clear. For each of the types of entities cited above, some more or less elaborate story can be constructed, and some of them have indeed evoked an extensive philosophical literature. But the effect in each case is to distance the notions of reference and “the world” from direct intuition. The cumulative effect of considering all of them together is a “world” in which direct intuition applies only to a very limited class of instances.

4 Pushing “the World” into the Mind

To sum up so far: The commonsense position on reference, which standard approaches to semantics take as their starting point, suffers from two complementary problems. First, if language is in the minds of language users, it is necessary to invoke some mystical connection from the mind to the world, either at the level of language, as in figure 13.3, or at the level of the concepts the language expresses, as in figure 13.4. Second, the notion of “objects in the world” is itself a serious problem.

I propose to deal with these problems by abandoning the unexamined notion of "objects in the world," and, for purposes of the theory of reference, pushing "the world" into the mind of the language user too, right along with language.

What does it mean to "push the world into the mind"? The change is in how we're supposed to understand statements about reference. (7)–(8) give the two opposing alternatives.

7 Commonsense realist theory of reference:

Phrase P of language L, uttered in context C, refers to entity E in the world (or possible worlds).

8 Conceptualist theory of reference:

A speaker S of language L judges phrase P, uttered in context C, to refer to entity E in [the world as conceptualized by S].

That is, in the conceptualist theory, the speaker's judgment and conceptualization play a critical role.

As initial motivation for exploring the conceptualist position, let's observe that a language user cannot refer to an entity without having some conceptualization of it. Consider an example like (9).

9 I don't know what that was, but here it comes again!

In order to utter (9) (and mean it), the speaker must have conceptualized some relevant entity, though certainly without a full characterization. That is, conceptualization of a referent is a necessary condition for a speaker to refer. However, being in the *real* world is not a necessary condition: speakers can refer without difficulty to entities like *Sherlock Holmes* and *the unicorn in my dream last night*. Furthermore, an entity's being in the real world is not *sufficient* for reference either: one has to conceptualize it in at least some minimal way. In short, an entity's being in the real world is neither a necessary nor a sufficient condition for a speaker's being able to refer to it. Rather, the crucial factor is having conceptualized an entity of the proper sort.

Still, I would not blame the reader for being a bit suspicious of this expression "the world as conceptualized by the language user." It smacks of a certain solipsism or even deconstructionism, as though language users get to make up the world any way they want, as though one is referring to one's mental representations rather than to the things represented. And indeed, there seems little choice. Figure 13.1, the conceptualist position, has no direct connection between the form of concepts and the outside world. On this picture our thoughts seem to be trapped in our own brains.

This outcome, needless to say, has come in for harsh criticism, from many different quarters, for example:

But how can mapping a representation onto another representation explain what a representation means? ... [E]ven if our interaction with the world is always mediated by representation systems, understanding such systems will eventually involve considering

what the systems are about, what they are representations of. (Chierchia and McConnell-Ginet, 1990, p. 47)

... words can't have their meanings *just* because their users undertake to pursue some or other linguistic policies; or, indeed, just because of any purely *mental* phenomenon, anything that happens purely "in your head." For "John" to be John's name, there must be some sort of real relation between the name and its bearer... something has to happen *in the world*. (Fodor, 1990, pp. 98–9)

But we can know the Markerese translation of an English sentence [i.e. its conceptual structure] without knowing the first thing about the meaning of the English sentence: namely, the conditions under which it would be true. (Lewis, 1972, p. 169)

How is it possible to escape this attack? I think the only way is to go deeper into psychology, and to deal even more carefully with the notion of thought. Think about it from the standpoint of neuropsychology: the neural assemblies responsible for storing and processing conceptual structures indeed *are* trapped in our brains. They have no direct access to the outside world. A position like Fodor's says that the "Language of Thought" is made of symbols that have meaning with respect to the outside world. I would rather say that conceptual structures are *not* made of symbols – they don't *symbolize* anything – and they don't *have* meanings. Rather, I want to say that they *are* meaning: they do exactly the things meaning is supposed to do, such as support inference and judgment. Language is meaningful, then, because it connects to conceptual structures. Such a statement is of course anathema to many semanticists and philosophers, not to mention to common sense. Still, let's persist and see how far we can go with it.

5 A Simple Act of Deictic Reference

Consider about the simplest act of using language to refer to a "middle-sized object": a use of referential deixis such as (10).

10 Hey, look at that! [pointing]

The deictic pronoun *that* has almost no intrinsic descriptive content; its semantics is almost purely referential. In order to understand (10), the hearer not only has to process the sentence but also has to determine what referent the speaker intends by *that*. This requires going out of the language faculty and making use of the visual system.

Within the visual system, the hearer must process the visual field and visually establish an individual in it that can serve as referent of *that*. The retinal image alone cannot do the job of establishing such an individual. The retina is sensitive only to distinctions like "dark point in bright surround at such-and-such a location on retina." The retina's "ontology" contains no objects and no external location. Nor is the situation much better in the parts of the brain most directly fed by the retina: here we find things like local line and edge detectors in various orientations, all in retinotopic

format (Hubel and Wiesel, 1968) – but still no objects, no external world. And this is all the contact the brain has with the outside world; inboard from here it's all computation.

However this computation works, it eventually has to construct a cognitive structure that might be called a "percept." The principles and neural mechanisms that construct percepts are subjects of intensive research in psychology and neuroscience, and are far from understood. The outcome, however, has to be a neurally instantiated cognitive structure that distinguishes individuals in the perceived environment and that permits one to attend to one or another of them. This cognitive structure that gives rise to perceived individuals is nonlinguistic: human infants and various animals can be shown experimentally to identify and track individuals more or less the way we do, so the best hypothesis is that they have percepts more or less like ours.

Of course percepts are trapped inside the brain too. There is no magical direct route between the world and the percept – only the complex and indirect route via the retina and the lower visual areas. Hence all the arguments that are directed against conceptualist semantics apply equally to percepts. This may bother some philosophers, but most psychologists and neuroscientists take a more practical approach: they see the visual system as creating a cognitive structure which constitutes part of the organism's understanding of reality, and which helps the organism act successfully in its environment (Marr, 1982; Koch, 2004). If there is any sense to the notion of "grasping" the world perceptually, this wildly complex computation is it; it is far from a simple unmediated operation.

And of course a visual percept is what is linked to the deictic *that* in (9) and (10), through the interfaces between conceptual structure and the "upper end" of the visual system. Thus language has indeed made contact with the outside world – but through the complex mediation of the visual system rather than through some mysterious mind-world relation of intentionality. Everything is scientifically kosher.

A skeptic may still be left grumbling that something is missing: "We don't perceive our percepts in our heads, we perceive objects out in the world." Absolutely correct. However, as generations of research in visual perception have shown, the visual system populates "the world" with all sorts of "objects" that have no physical reality, for instance the things in example (3): the square subtended by four dots and the "amodally completed" horizontal rectangle. So we should properly think of "the perceptual world" (or "phenomenal world" in the sense of Koffka, 1935) not as absolute reality but as the "reality" constructed by our perceptual systems in response to whatever is "really out there."

Naturally, the perceptual world isn't totally out of synch with the "real world." The perceptual systems have evolved in order that organisms may act reliably in the real world. They are not concerned with a "true model of the world" in the logical sense, but with a "world model" that is good enough to support the planning of actions that in the long run lead to better propagation of the genes. Like other products of evolution, the perceptual systems are full of "cheap tricks," which is *why* we see virtual objects: these tricks *work* in the organism's normal environment. It is only in the context of the laboratory that their artificiality is detected.

Thus *the perceptual world is reality for us*. Apart from the sensory inputs, percepts are entirely "trapped in the brain"; they are nothing but formal structures instantiated

in neurons. But the perceptual systems give us the sense, the feeling, the affect, of objects being out there. We experience objects in the world, not percepts in our heads. That's the way we're built (Dennett, 1991; Koch, 2004).

In short, the problem of reference for the intuitively clear cases is not at bottom a problem for linguistic theory, it is a problem for perceptual theory: how do the mechanisms of perception create for us the experience of a world "out there"?

I suspect some readers may find this stance disquieting. My late friend John Macnamara, with whom I agreed on so much, used to accuse me of not believing there is a real world. But I think the proper way I should have replied to him is that we are ultimately concerned with *reality for us*, the world in which we lead our lives. Isn't that enough? (Or at least, isn't that enough for linguistic semantics?) If you want to go beyond that and demand a "more ultimate reality," independent of human cognition, well, you are welcome to, but that doesn't exactly render my enterprise pointless.

6 Indexical Features

Let's continue in this direction a bit and see where it leads. A basic characteristic of a percept is that it constitutes a figure distinguished from the background on which it is located and in which it moves. Suppose you hear me say *Hey, look at that!* You stare at the patterned rug and see nothing; then suddenly a bug "pops out" in experience ("Oh god, there it is! Eeuww!"). Nothing has changed in "the physical world" or in your retinal image. All that has changed is the organization of percepts "trapped" in your brain: the bug-percept has emerged as figural. Let us call this figural characteristic the *indexical* feature of the percept. It gives the mind a "something" to which perceptual features can be bound. A speaker who utters (9) (*I don't know what that was, but here it comes again!*) is reporting a percept with an indexical feature but minimal identifiable perceptual features.

The mind establishes indexical features, among other things, in response to perceptual input. But once established, they need not go away in the absence of perceptual input: we intuitively sense that objects continue to exist when we're not seeing them. Indexical features can undergo "splitting," for instance when we break a lump of clay in half to form two distinct individuals. They can also undergo "merger," as when we mold two lumps of clay into an undifferentiated mass.

The indexical feature of a percept is the crucial feature for linguistic reference. If there is no indexical feature to which perceptual features can be bound, there is nothing that a deictic or other referring linguistic expression can be attached to either: there is no *that* there.

We can see how this account extends to some of the problematic entities mentioned earlier. Consider again virtual objects such as the square subtended by four dots. Although there is no square "actually out there," the visual system constructs a percept with all the right features for a "square out there." Hence there is something "out there" in the perceived world, and we can refer to it.

Returning to the standard literature on reference: Frege (1892) made the distinction between the sense and the reference of an expression by citing the well-known

example *The morning star is the evening star*. In his analysis, two senses are attached to the same reference. We can understand Frege's example in present terms as reporting the merger of indexicals associated with different perceptual features, on the basis of some discovery. So Frege's problem is not a uniquely linguistic problem; rather it lies in a more general theory of how the mind keeps track of individuated entities.

The reverse, indexical splitting, can also occur. I used to think there was one literary/cultural theorist named Bloom, until one day I saw a new book by Bloom and was surprised because I had thought he had been dead for several years. It suddenly dawned on me, to my embarrassment, that there were *two* Blooms, Allan and Harold – so my indexical came to be split into two. Think also of discovering that someone you have seen around the neighborhood is actually twins.

More generally, indexical features play a role altogether parallel to the discourse referents in various approaches within formal semantics such as Discourse Representation Theory (Kamp and Reyle, 1993), File Change Semantics (Heim, 1989), and Dynamic Semantics (Groenendijk et al., 1996). Hence many of the insights of these approaches can be taken over here, with the proviso that the “model” over which reference is defined should be a psychologically motivated one.

7 Entities Other Than Objects

Let's go a bit more into what a psychologically motivated model would include. Let's start by observing that indexical features do not always identify individual objects: these are just one sort of entity that can be identified by an indexical.

For instance, consider uttering *Hey, look at THAT!* in reaction to muddy water pouring out of a pipe. Here the intended reference may well be the *substance* rather than any particular bounded quantity of it. Moreover, various deictics can be used to refer to a wide range of nonmaterial entities, as seen in (11).

- 11 (a) Goodness! Did you hear *that*? [Sound]
 (b) [Dentist:] Did *that* hurt? [Tactile sensation]
 (c) Pro-prepositional phrase:
 Please put your coat *here* [pointing] and put your hat
there [pointing]. [Location]
 He went *thataway* [pointing]. [Direction]
 (d) Pro-verb phrase:
 Can you do *that* [pointing]? [Action]
 Can you do *this* [demonstrating]? [Action]
 (e) that . . . happen:
 If *that* [pointing] happens again, I'm leaving. [Event/Activity]
 (f) Pro-manner adverbial:
 You shuffle cards *thus/so/like this/this way*
 [demonstrating]. [Manner]
 I used to walk *that way/like that* [pointing to someone
 walking in a funny fashion]. [Manner]

- (g) Pro-measure expression:
 The fish that got away was *this* long [demonstrating]. [Distance]
 There were about *this* many people at Joe's party too.
 [gesturing toward the assembled throng]. [Amount]
 (h) Pro-time-PP:
 You may start . . . right . . . *now!* [clapping] [Time]

The deictic expressions here refer to entities in the conceptualized world that can be picked out with the aid of the accompanying gesture. But the entities referred to are not objects, they are sounds, sensations, locations, directions, and so on.

In order to accommodate these possibilities for reference, it is useful to introduce a kind of “ur-feature” that classifies the entity being referred to into an *ontological type*. Each of the ontological types – objects, sounds, actions, locations, and so forth – has its own characteristic conditions of identity and individuation. It is a task both for natural language semantics and for cognitive psychology/neuroscience to work out the logic and the characteristic perceptual manifestations of each type.

Are all the sorts of entities in (11) “in the world”? They are certainly not like refrigerators – you can't touch them or move them. In fact, it is odd to say they all *exist* in the way that refrigerators exist (“the length of the fish exists”?). Yet (11) shows that we can pick them out of the perceptual field and use them as the referents of deictic expressions. So we must accord them some dignified status in the perceived world – the “model” that serves as the basis for linguistic reference.

Now notice what has just happened. Up to a moment ago I was concerned with reference to objects, and I used perceptual theory to ground the theory of reference. Now all of a sudden I have turned the argument on its head: If this is the way reference relates to perception, perception must be providing a far richer range of entities than had previously been suspected. It is now a challenge for perceptual theory to describe how the perceptual systems accomplish this. In other words, examples like (11) open the door for fruitful cooperation between linguistic semantics and research in perception.

This repertoire of ontological types seems to me a good candidate for a very skeletal unlearned element of cognition. Again, because it is central not just to language but to perception and action, we need not call it part of Universal Grammar, the human specialization for learning language. But we do have to call it part of the innate human cognitive endowment.

So far, I've only talked about perceivable entities. What makes them perceivable is that they have features that connect to the perceptual interfaces in figure 13.1. Now suppose that conceptual structure contains other features that do not pertain to a perceptual modality, but which connect instead to the inferential system in figure 13.1. Such features would provide “bridges” to other concepts but no direct connection to perception. That is, they are used in *reasoning* rather than in identification. Let's call these *inferential* features by contrast with *perceptual* features.

What would be candidates for inferential features? Consider an object's value. This is certainly not perceivable, but it influences the way one reasons about the object, including one's goals and desires concerning the object. Value can be established by all sorts of means, mostly very indirect. Thus value seems like a prime candidate for

a type of inferential feature. Another candidate is high-level taxonomic category such as *natural kind* or *artifact*. The classification of a particular object (a fork, a refrigerator, a book, etc.) as an artifact predicts only that people typically use this object for something. But it does not predict what the object looks like, how it works, or what one does with it. Rather, beyond the ontological feature "physical object," which does have perceptual consequences, *artifact* is not on the whole a collection of perceptual features; rather it bridges to all sorts of inferences that can in turn potentially engage perceptual features.

Now suppose there were a concept that had only inferential features. This then would be an abstract concept such as those expressed by *the value of my watch* and *the meaning of this word*. From the point of view of language, abstract objects are not distinguished from concrete objects: all that is necessary for reference is an indexical feature. The interfaces to language can't "see" whether a concept's descriptive features are perceptual or inferential. So as far as language is concerned, reference proceeds in the same fashion whether or not there is a physical object out there in the world – or even a conceptualized physical object.

8 Satisfaction and Truth

An important piece has been left out of the story so far. Linguistic expressions used in isolation cannot refer: they can only *purport* to refer. For example, suppose I'm talking to you on the phone and say *Hey, will you look at THAT!* You understand that I intend to refer to something; but you can't establish the reference and therefore can't establish the contextualized meaning of the utterance.

A referential expression *succeeds* in referring for the hearer if it is *satisfied* by something that can serve as its referent. Remember: in realist semantics, satisfaction is a relation between a linguistic expression and an entity in the world; but in conceptualist semantics, the entity is in [the world as conceptualized by the language user]. It is this latter notion of satisfaction that we have to work out here.

To work out a conceptualist notion of satisfaction, we invoke one component in figure 13.1 that I haven't yet discussed: the integration of concepts with the knowledge base. Suppose I say to you: *I talked to Henk Verkuyl today in the supermarket*. In isolation, the proper name *Henk Verkuyl* purports to refer; you assume that I intended it actually to refer. If you know Henk Verkuyl (i.e. he is part of your knowledge base), you can establish the intended reference in your own construal of the world. If you do not know him, the proper name is unsatisfied for you.

More generally, four different situations can arise.

- 12 (a) The purported referent of my utterance is present in your knowledge base or the readily available context, and it satisfies the referential expression.
- (b) The purported referent is not present in your knowledge base or the readily available context, and so you add a new character into your knowledge base to satisfy the referential expression. If I anticipate this situation for you, I'll use a specific indefinite expression like *a friend of*

mine or *this friend of mine*, *Henk Verkuyl* instead of a definite description or an unadorned proper name.

- (c) The purported referent is in conflict with your knowledge base, as in Russell's famous example *the present king of France*. Or it is in conflict with the readily available context, for example in a situation where I speak of *the apple on the table* and you see either no apples or two apples on the table. In this case you have to fall back on some repertoire of repair strategies: guessing what I intend, asking me for clarification, deciding to ignore me altogether, and so on.
- (d) The purported referent contains features that inherently conflict with each other, so that nothing can possibly satisfy the expression. In such a case, for instance *the square circle*, you'll find the expression anomalous, and you'll again have to fall back on repair strategies.

This is all clear with the reference of noun phrases. Next let's look at the reference of sentences. The standard custom in the formal semantics tradition, going back to Frege, is that the intended referent of a (declarative) sentence is a truth-value. I must confess I have never understood the argument for this position (e.g., Chierchia and McConnell-Ginet, 1990, ch. 2). I am going to explore instead the alternative position (proposed in Jackendoff, 1972) that the intended reference of a declarative sentence is a situation (an event or a state of affairs). Traditionalists should not worry: truth-values will get their due shortly.

The view that sentences refer to situations is motivated largely by linguistic parallels to referentiality in noun phrases. This is a kind of evidence not frequently cited in the literatures of philosophy of language and formal semantics – although Situation Semantics (Barwise and Perry, 1984) made a good deal of it. First notice how noun phrases and sentences are parallel in the way they can be used to accompany deictic reference:

- 13 (a) Will you look at that! A blimp!
- (b) Will you look at that! The baby's walking!

(13)(a) draws your attention to an object in the environment; (13)(b) to an event – not to a truth-value. Next, notice that discourse pronouns can co-refer with sentences as well as with noun phrases.

- 14 (a) A blimp appeared in the sky. It was huge.
- (b) Suddenly the baby started walking. It astounded her parents.

The antecedent of *it* in (14)(b) is the whole preceding sentence, so it presumably has the same referent. But *it* certainly does not refer to a truth-value: *it astounded her parents* does not assert that the parents are astounded by the truth of the proposition expressed by *the baby started walking*; they are astounded by the *event* of the baby walking. Hence this event must be the referent of the preceding sentence as well.

Next consider embedded *that*-clauses in a context where they alternate with noun phrases.

- 15 (a) The apple on the table astounded Max.
(b) That the Red Sox won today astounded Max.

What astounded Max wasn't a truth-value, it was an event.

Parallel to the four possible situations in (12) for satisfaction of a noun phrase's purported referent, we find four possible situations for the referent of *the Red Sox won today* in (15)(b).

- 16 (a) Your knowledge base includes the event of the Red Sox winning, and this satisfies the intended referent of the clause.
(b) Your knowledge base does not include the event of the Red Sox winning, so you add this to your knowledge base as the referent of the clause.
(c) Your knowledge base includes something in conflict with the purported event of the Red Sox winning (say, your take on the world is that the Red Sox didn't play). Then you have to engage in some repair strategy.
(d) The features of the purported event are inherently in conflict, so that there is no possible referent. In such a case, for instance *That the square is a circle astounded Max*, the clause is judged anomalous, and you again have to resort to repair.

There are of course, other cases of *that*-clauses that don't come out this way, notably in so-called intensional (with an *s*) contexts such as the complement of *believe*. However, noun phrases in this position are subject to the same distortions of referentiality:

- 17 (a) Max believes that there is a tooth fairy/the square is a circle.
(b) Max believes in the tooth fairy/in square circles.

In both of these, the speaker makes no commitment to the existence of the tooth fairy or square circles.

So far, then, I've tried to convince you that it makes sense to regard a clause as referentially satisfied by a conceptualized situation. What about truth-values, then? The judgment of a declarative sentence's truth-value follows from how it is referentially satisfied.

- 18 (a) In case (16)(a), where the intended referent of the sentence is present in your knowledge base (or can be deduced from it), you judge the sentence true.
(b) In case (16)(b), where there is no conflict with your knowledge base, you normally take the sentence to be informative and presumed true.
(c) In case (16)(c), you judge the sentence false.
(d) In case (16)(d), you judge it analytically false.

Thus truth is defined in terms of reference and satisfaction, just as proposed by Tarski (1936/1983).

In short, the parallelism in the reference of noun phrases and sentences lies in the parallelism between conceptualized objects and conceptualized situations. The notion of satisfaction applies identically to both. However, sentences have an additional layer of evaluation, in which they are characterized as true or false on the basis of how they are referentially satisfied.

In this approach, then, the problem of characterizing the conditions under which a sentence is judged true does not go away. It is just demoted from the paramount problem of semantic theory to one among many problems. What seems more basic here is the conditions of satisfaction for referential constituents and how they interact with the knowledge base. In fact, a great deal of research in "truth-conditional" semantics can easily be reconstrued as addressing this issue. For instance, the question of whether sentence S_1 entails sentence S_2 has nothing to do with their truth-values – sentences may describe thoroughly fictional or hypothetical situations. Rather, S_1 entails S_2 if adding the situation referred to by S_1 to an otherwise blank knowledge base enables the situation referred to by S_2 to be satisfied. The factors involved in such satisfaction, and the form of the rules of inference, may remain essentially unchanged from a truth-conditional account.

Above all, the conceptualist approach shifts the focus of semantics from the question "What makes sentences true?" to what I take to be the more ecologically sound question, "How do we humans understand language?" – where I mean "ecologically sound" in the sense that it permits us to integrate semantics with the other human sciences. I take this to be a positive step.

There are obviously many further problems in establishing a mentalistic semantics and in particular a mentalistic theory of reference and truth. *Foundations of Language* (Jackendoff, 2002) addresses some of these and not others. But here is where we are so far. We have not ended up with the rock-solid rigid notion of truth that the realists apparently want. Rather, I think we have begun to envision something that has the promise of explaining our human sense of truth and reference – and why philosophical and commonsensical disputes about truth and reference so often proceed the way they do. I have no illusions that this work is over, but it strikes me as a path well worth exploring.

References

- Barwise, J. and Perry, J. (1984). *Situations and Attitudes*. Cambridge, MA: MIT Press.
Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
Byrne, R. W. and Whiten, A. (eds.) (1988). *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes, and Humans*. Oxford: Clarendon Press.
Chierchia, G. and McConnell-Ginet, S. (1990). *Meaning and Grammar: An Introduction to Semantics*. Cambridge, MA: MIT Press.
Chomsky, N. (1995). *The Minimalist Program*. Cambridge, MA: MIT Press.
Dennett, D. (1991). *Consciousness Explained*. New York: Little, Brown.
Fiske, A. (1991). *Structures of Social Life*. New York: Free Press.
Fodor, J. A. (1975). *The Language of Thought*. Cambridge, MA: Harvard University Press.
— (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.

The Intentional Inexistence of Language – But Not Cars

Georges Rey

In the short space here, I want to address an issue about the reality of language and the ordinary external world that Jackendoff raises in his chapter of the present volume (chapter 13, LOCATING MEANING IN THE MIND, WHERE IT BELONGS), and that has been a persistent theme in his work for the last 20 years.

It would seem to be a commonplace that people, when they talk, produce tokens of such things as words, sentences, morphemes, phonemes, and phones – I'll call tokens of all such types, "standard linguistic entities" ("SLEs"). Part and parcel of this commonplace would be the presumption ("physical tokenism," hereafter abbreviated to "PT") that these entities can be identified with some sorts of *acoustic* phenomena, e.g., wave patterns in space and time. For instance, Devitt and Sterelny write:²

[PT] Tokens are datable, placable parts of the physical world . . . Inscription types and sound types are identifiable by their overt physical characteristics and so we might call them "physical types." (Devitt and Sterelny, 1987, p. 59)

Over the years, however, this latter presumption has been repeatedly challenged by linguists, such as Saussure (1916/1966), Sapir (1933/1963), Chomsky and Halle (1968), Jackendoff (1983, 1987), and Chomsky (2000). In his textbook on phonetics, for example, John Laver writes:

The stream of speech within a single utterance is a continuum. There are only a few points in this stream which constitute natural breaks, or which show an articulatory, auditorily or acoustically steady state being momentarily preserved, and which could therefore serve as the basis for analytical segmentation of the continuum into "real" phonetic units. . . . The view that such segmentation is mostly an imposed analysis, and not the outcome of discovering natural time-boundaries in the speech continuum, is a view that deserves the strongest insistence. (Laver, 1993, p. 101)

- (1990). *A Theory of Content and Other Essays*. Cambridge, MA: MIT Press.
- (1998). *Concepts: Where Cognitive Science Went Wrong*. Oxford: Oxford University Press.
- Frege, G. (1892). Über Sinn und Bedeutung. *Zeitschrift für Philosophie und Philosophische Kritik*, 100, 25–50. (English translation in P. Geach and M. Black (eds.), *Translations from the Philosophical Writings of Gottlob Frege*. Oxford: Blackwell, 1952.)
- Gigerenzer, G. (2000). *Adaptive Thinking: Rationality in the Real World*. New York: Oxford University Press.
- Goodall, J. van L. (1971). *In the Shadow of Man*. New York: Dell.
- Groenendijk, J., Stokhof, M., and Veltman, F. (1996). Coreference and Modality. In S. Lappin (ed.), *The Handbook of Contemporary Semantic Theory*. Oxford: Blackwell.
- Hauser, M. D. (2000). *Wild Minds: What Animals Really Think*. New York: Henry Holt.
- Hauser, M. D., Chomsky, N., and Fitch, W. T. (2002). The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298, 1569–79.
- Heim, I. (1989). *The Semantics of Definite and Indefinite Noun Phrases in English*. New York: Garland.
- Hubel, D. and Wiesel, T. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology* (London), 195, 215–43.
- Jackendoff, R. (1972). *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.
- (1976). Toward an explanatory semantic representation. *Linguistic Inquiry*, 7, 89–150.
- (1983). *Semantics and Cognition*. Cambridge, MA: MIT Press.
- (1987). *Consciousness and the Computational Mind*. Cambridge, MA: MIT Press.
- (1990). *Semantic Structures*. Cambridge, MA: MIT Press.
- (2002). *Foundations of Language*. Oxford: Oxford University Press.
- (forthcoming). *Language, Culture, Consciousness: Essays on Mental Structure*. Cambridge, MA: MIT Press.
- Kahneman, D., Slovic, P., and Tversky, A. (eds.) (1982). *Judgment under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.
- Kamp, H. and Reyle, U. (1993). *From Discourse to Logic*. Dordrecht: Kluwer.
- Katz, J. J. (1981). *Language and other Abstract Objects*. Totowa, NJ: Rowman & Littlefield.
- Koch, C. (2004). *The Quest for Consciousness: A Neurobiological Approach*. Englewood, CO: Roberts and Company.
- Koffka, K. (1935). *Principles of Gestalt Psychology*. New York: Harcourt, Brace & World.
- Köhler, W. (1927). *The Mentality of Apes*. London: Routledge & Kegan Paul.
- Kripke, S. (1972). Naming and Necessity. In D. Davidson and G. Harman (eds.), *Semantics for Natural Language*. Dordrecht: Reidel.
- Lakoff, G. (1987). *Women, Fire, and Dangerous Things*. Chicago: University of Chicago Press.
- Langacker, R. (1987). *Foundations of Cognitive Grammar*, vol. 1. Stanford, CA: Stanford University Press.
- Lewis, D. (1972). General Semantics. In D. Davidson and G. Harman (eds.), *Semantics for Natural Language*. Dordrecht: Reidel.
- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Montague, R. (1973). The proper treatment of quantification in ordinary English. In J. Hintikka, J. Moravcsik, and P. Suppes (eds.), *Approaches to Natural Language*. Dordrecht: Reidel.
- Pinker, S. and Jackendoff, R. (2005). The faculty of language: What's special about it? *Cognition*, 95, 201–36.
- Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.
- Tarski, A. (1936/1983). The establishment of scientific semantics. In A. Tarski, *Logic, Semantics, Metamathematics* (J. H. Woodger, trans; 2nd edn., ed. J. Corcoran). Indianapolis: Hackett.
- Tooby, J. and Cosmides, L. (1992). The psychological foundations of culture. In J. H. Barkow, Leda Cosmides, and John Tooby (eds.), *The Adapted Mind*. Oxford: Oxford University Press.