

## Discourse Semantics

Russell S. Tomlin, Linda Forrest,  
Ming Ming Pu and Myung Hee Kim

### The Problem of Meaning

It is remarkable how well, how routinely, we communicate with one another. Whether we relate the common and trivial events of our day-to-day lives or argue passionately for our ideologies; whether we read a simple newspaper account of a world event or an academic paper on the intricacies of human language; whether we compose a simple thank-you note or a legal brief, our problems remain the same: how can we ensure that our comprehender gets the message we intend, and how do we derive the message intended from what we hear or read? Our problems as linguists also remain the same: for the comprehender, whether listening or reading, we must describe and account for how meaning is derived from any of these multiple sources; and, for the producer, whether speaker or writer, we must account for how meaning is conveyed. This problem of meaning, this problem of *discourse semantics*, is a complex one, involving interplay among a wide array of linguistic and non-linguistic processes.

Let us consider a very simple example, the text fragment in (1).

#### (1) Text fragment

- 1 . . . puck knocked away by Dale McCourt,  
Ø picked up again by Steve Shutt/  
Now Shutt coming out, into the Detroit zone/  
He played it out in front/
- 5 There's Lemaire with a shot.  
and it was blocked by Reed Larson/

The fragment is a short segment transcribed from the concurrent and on-line description of an ice hockey game. The professional announcer rapidly produces his description as the events unfold before him. One can see that the announcer alternates between active and passive voice (lines 4 vs lines 1, 2, and 6), between nominal and pronominal form, and between existential or presentative clause type (line 5) and simple clauses (line 3). What is it about the message the announcer must convey that leads to the selection of these alternative structures? Our initial and intuitive guesses,

some grounded perhaps in our traditional schooling, is that the difference between the active and passive is related to which referent is given more emphasis or is somehow, at least at the moment, more important; that the selection of the proper names versus the pronoun is related to keeping clear which referent one targets at a given moment, and that the selection of the existential is tied to whether or not the referent has been recently 'on stage,' as it were, as the action unfolded.

This chapter provides an introduction to the concepts and processes underlying our intuitions about how matters of emphasis and importance and prior knowledge contribute to the meaning one derives from text and discourse during comprehension, and how they contribute to and shape decisions about the use of language structures in the service of larger meanings.

The remainder of this section provides a framework for this effort.

### *Metaphors of Discourse Interaction*

The way one thinks of discourse has a strong effect on the kind of theory or model of discourse semantics one creates. The most naive metaphor can be called the *conduit metaphor* of discourse. In this view, the speaker packages his intended meaning into a textual artifact. This artifact, in essence, *contains* the meaning intended by the speaker. It is conducted to the listener in either spoken or written form. The text is then unpacked and its meaning extracted from the text artifact by the listener.

Within the conduit metaphor one is invited to think of language and texts as *containing* meaning. It invites one to think of meanings as inherent to the text itself or to the clauses and words which make it up. Failure by the speaker to be clear derives from a failure to match well intended meaning with the details of text as produced. Failure by the listener to comprehend arises from the failure to extract proper semantic detail or nuance from the text itself. Language, in a sense, is viewed as a precision instrument, which is used to craft a precise meaning, fully embodied in the text. The text becomes both the target of analysis and the source, sometimes the only source, for explanatory information regarding its structure and meaning.

An alternative metaphor, the one we will embrace here, can be called the *blueprint metaphor* of discourse. In this view, the speaker holds a conceptual representation of events or ideas which he intends should be replicated in the mind of the listener. The listener is neither helpless nor passive in the endeavor but actively engages in constructing her own conceptual representation of the matters at hand. The speaker behaves as a sort of architect and his linguistic output, the text, can be viewed less as a fully fleshed out semantic construct than as a blueprint to aid the listener during the construction of a conceptual representation. Just as a true blueprint contains no actual building materials but depicts by convention how existing materials should be employed in constructing a given edifice,

so the text itself contains little or no meaning *per se* but serves by convention to guide the listener in constructing a conceptual edifice.

One can demonstrate the utility of the blueprint metaphor over the conduit metaphor with a brief demonstration or 'thought' experiment. Consider a simple task – like putting on your favorite sweater. Try to write a passage which instructs a naive reader – a very naive reader – on how to put on a sweater. No matter how much detail is provided in the text itself, the text will never alone, never without the employment of other sources of information, be adequate to get the sweater on. The conceptual model of sweater-wearing the speaker has in mind is more complex than can be captured in the text, and the text can only be helpful to the reader as it exploits and builds on conceptual models of sweaters and donning and whatever else can be brought to bear on the task.

Building or interpreting text blueprints requires dealing with two fundamental problems. The first is integrating the semantic information provided in each utterance into a coherent whole.<sup>1</sup> The speaker must select pertinent concepts and events from his experience and organize this in a way helpful to the listener. The listener must integrate utterances heard into a coherent representation which permits her to access or construct concepts and events virtually identical to those held by the speaker. We can call this the problem of *knowledge integration*.

The second fundamental problem is managing the flow of information between speaker and listener in dynamic, real time interaction. The speaker will help the listener succeed in knowledge integration in part by directing the listener's efforts to process the information provided through the text. For example, the speaker will help the listener by exploiting information held in common as a prelude to or anchor for information the speaker believes will be novel or unexpected for the listener. The coherence of the knowledge held by the listener will be affected by how cohesive the information is that the speaker offers. We call this the problem of *information management*.

Knowledge integration requires effective information management, but effective information management is not enough to account for knowledge integration.

### *What Is Linguistic and What Is Not*

The social and cognitive processes of discourse creation and comprehension involve a complex array of both linguistic and non-linguistic processes. It will prove immensely useful to keep in mind that both sorts of processes are implicated in our understanding of discourse and text and that the dividing line between linguistics and psychology cannot be neatly drawn. These processes include at least three different types:

- 1 *Morpho-syntactic codings*: grammatical codes employed automatically and generally non-consciously by speaker and listener to shape the flow and decoding of information during production and comprehension.

- 2 *Implicatures*: signals employed under routine conditions by the speaker and which result in common and rapidly drawn inferences by the listener.
- 3 *Planning and inference*: high level processes of planning on the part of the speaker and inference on the part of the listener which set general directions for less global processing by speaker and listener.

These issues will be dealt with more fully in the next main section where we discuss the problem of knowledge integration in discourse.

#### *Four Central Threads of Information Management*

The ease with which the listener can integrate knowledge into a coherent vision of the concepts and events the speaker has in mind is determined largely by how well the speaker controls the flow of information to the listener. Information flow (Chafe, 1979; 1980a; 1987; 1994) depends on the effective control of four distinct kinds of information:

- 1 *Rhetorical management*: participants must be clear about the goals and intentions of the discourse interaction as these constrain greatly the propositional content of production and the construal of what is heard.
- 2 *Referential management*: participants must keep track of the referents and propositions they have in common.
- 3 *Thematic management*: participants must keep track of the central elements around which the discourse is developed.
- 4 *Focus management*: the participants must keep track of which referents they are dealing with at any given moment and must take steps to ensure they are on the same one(s).

Each of these four threads will be pursued in separate sections in this chapter.

#### *The Granularity of Discourse Interaction: Three Key Levels*

Discourse is neither flat nor linear in its organization; it is hierarchical, with clauses forming higher order structures, paragraphs, which in turn combine to form larger episodes or sections of discourse. For our purposes, it is important to distinguish three key levels:

- 1 *Global coherence*: the participants develop a sense of what the overall narrative or procedure or conversation deals with.
- 2 *Episodic coherence*: the participants are sensitive to smaller scale units which contribute to global coherence but which display an internal gist of their own.
- 3 *Local coherence*: the participants make sense of the contribution of individual sentences or utterances.

These three levels of granularity will be important to keep in mind as we pursue each of the four threads of information management.

### *Historical Perspective*

The pragmatic notions discussed below emerged from research traditions seeking answers to questions of discourse and syntax (see Cumming and Ono, Chapter 4 in this volume), in particular to questions of typological differences between languages. For example, the precursor to Prague School work on *functional sentence perspective* was a volume by Weil (1887), originally published in 1844, which appealed to pragmatic and cognitive rationales to explain differences between the less constrained word orders of Latin and Ancient Greek and the more constrained word orders of modern European languages.

In the same vein, Prague School research was directed, as it is directed today, at explaining differences between languages like Czech and English. All of the pragmatic notions discussed below arose historically from attempts to explain otherwise opaque structural contrasts. Why, after all, should a language 'need' simple voice alternants like English passive, let alone complex systems typical of Philippine Austronesian languages like Tagalog or Cebuano, when the essential semantic content of passive and active is identical?

The central concepts employed in developing theories of discourse semantics have a long history. This chapter will respect that history, linking current conceptualizations with their historical antecedents.

### *Methodological Issues and Dilemmas*

While it is not particularly fashionable to be concerned with methodology, it is nevertheless important to examine a number of issues tied to the methods employed in the analysis of discourse semantics. Of particular interest will be the contributions made to discourse studies by introspection, by text counting procedures, and by experimental methods. Methodological issues are taken up later in the chapter.

### *Cognitive and Discourse Semantics*

There has always been an interest in the relationship between discourse and text structure and their processing and storage in the mind. As early as the nineteenth century, scholars such as von der Gabelenz appealed to primitive ideas of attention in discussing 'psychological subject' (related to notions of topic and theme). Similarly, early psychologists like Stout also explained linguistic behavior by appeal to the activity of the mind (cf. Lambrecht, 1994 or Gundel, 1974). More recently, many linguists, within both discourse studies and more general linguistics, have begun to appeal to cognitive notions in their theoretical efforts, in particular to ideas of activation (memory) and attention. The penultimate section examines the promise this line of research holds for the future of discourse semantics.

### Knowledge Integration in Discourse

To appreciate the richness and complexity of discourse semantics in discourse production and comprehension, let us consider two examples. First, from the point of view of the speaker, consider the drawing in Figure 3.1. Describe the scene as you examine it for the first time, perhaps even recording your description, for a listener who cannot see the drawing.

There are a number of important matters one must deal with in developing a theory of discourse semantics. First, the speaker must have in mind some sort of *conceptual representation*<sup>2</sup> of the subject matter under discussion. It is this representation which is the fundamental 'meaning' the speaker works with in constructing a discourse. Most of the time as linguists we think of the conceptual representation as a set of propositions, sometimes referred to as a text representation or text base (Kintsch, 1974; van Dijk and Kintsch, 1983), but a conceptual representation need not be simply or only propositional. In the case of the Figure 3.1 the speaker will form some sort of *visual* representation for the picture, which he will access when speaking. For other kinds of experience, there may well be other kinds of representation. For example, a conceptual representation for, say, tasting coffee may well be different from one for listening to music, which in turn may be distinct from that for shifting gears in a car. And, we must remember that we can create conceptual representations as well as access them from memory or perceptual experience.

In addition, one should keep in mind that conceptual representations are *dynamic*. While the picture in Figure 3.1 is itself static, our viewing of it is not: our eyes move from one element to another, pausing here and there, skipping here and there, forming smaller scale images and impressions. In the same vein, most conceptual representations are dynamic: we do not view pictures so much as we do scenes and other sorts of unfolding events. Even our abstract efforts are dynamic, as we leap from idea to idea or slog through a difficult argument. Conceptual representations are thus foundational for understanding discourse production.

The conceptual representation alone, though, is not the meaning of the discourse. Somehow the speaker must select information from the overall conceptual representation and build a text from that selected information. The speaker will select information he believes the listener needs, and that information will be presented to help the listener make good sense of what is heard. How the speaker *manages* information plays a critical role in discourse production and comprehension. Discourse management involves (at least) four independent threads.

One, the speaker never *merely* outputs a conceptual representation via language. The speaker always has some *purpose* or *goal*, perhaps even several, in mind. Such goals constrain how conceptual representations are searched, and they constrain the selection of information to convey. The picture in Figure 3.1 will yield different descriptions depending on whether the speaker is asked to 'describe what seems to be happening' or 'evaluate




Figure 3.1 *Scene example (artwork by Jennifer Jones, 1995)*

the quality of the drawing'. In each case the speaker will access a visual representation of the picture, but the information selected and ultimately incorporated into a verbal description will be different. We can think of these different goals or purposes in speaking as distinct *rhetorical goals* for discourse production and the associated selection processes as defining an area we will call *rhetorical management*.

Two, constrained by rhetorical goals, the speaker accesses the conceptual representation, dynamically selecting referents and propositions for the listener. As he does so, the speaker makes real time decisions about which referents and propositions are more central or important to the developing discourse. Such *starting points* ultimately assist the listener in building her own conceptual representation. We can think of important referents or starting points as specifying the *thematic organization* of discourse and the

associated pragmatic notions and processes as defining an area we will call *thematic management*.

Three, in parallel with thematic management, the speaker will also monitor dynamically which referents and propositions seem to be available to the listener already and which require introduction or reintroduction. Referents and propositions fall coarsely into two classes: those that the speaker believes the listener can readily access, the information in common; and those that the speaker believes the listener will need explicit assistance to access, the targeted information. We can think of the interlocking pattern of referents and propositions held in common or not as defining the *referential organization* of discourse, and the associated pragmatic notions and cognitive processes as defining an area we will call *referential management*.

Finally, the speaker will monitor dynamically which referents and propositions he wishes to ensure are brought to the attention of the listener. From time to time, the speaker may not be certain that the listener has the right referent in mind at the right time, or he may wish to somehow *highlight* or *emphasize* a particular referent or proposition. We can think of efforts to highlight or emphasize or steer the listener's attention to particular referents and propositions as defining the *focus organization* of discourse, and the associated pragmatic notions and cognitive processes as defining an area we will call *focus management*.

These four arenas of discourse management define the central problems of information management in discourse. While there is an extensive literature dealing with each of these areas, there is at present no single, comprehensive model of how they fit together to account for knowledge integration (but cf. Vallduvi, 1992 or Lambrecht, 1994). Still, one useful model for understanding what the speaker does is Levelt's (1989) blueprint for the speaker. In Levelt's model there are three principal components. There is a *conceptualizer*, which is responsible for the creation of a pre-verbal message, the central proposition to be encoded, annotated in some way for pertinent pragmatic statuses like *topic*, *focus*, or *given*. The pre-verbal message is then processed by an *utterance formulator*, which (in concert with the lexicon) assigns first a grammatical representation to the message and then a phonological representation. The phonological representation is then processed by an *articulator*, which is responsible for the generation of overt speech.

A model of discourse semantics will deal with the creation of Levelt's pre-verbal messages with its pragmatically annotated propositions. Figure 3.2 (based loosely on Levelt, 1989) presents a simple model of production, setting out the contribution of the four arenas of discourse management to the overall discourse creation task.

Briefly, the conceptualizer operates on a conceptual representation to produce something like a pre-verbal message or annotated proposition. The conceptual representation may come from several sources – current perception, long term memory, or some creative production. The






Figure 3.2 *A blueprint for the speaker*

conceptualizer employs rhetorical management to select a particular event or other conceptualization for communication consistent with the speaker's global and local plans for the discourse. This selected event is, in turn, and in parallel, annotated for thematic, referential, and focus statuses (all discussed below). The output of this process is a pre-verbal message, a proposition annotated for fundamental pragmatic statuses and ready for processing by the utterance formulator.

Turning to comprehension, one can see that the listener's problem parallels the speaker's, but it is not simply its reversal. The listener will form a conceptual representation as discourse is encountered. She will use information extracted from the text in concert with information already at hand to construct a conceptual representation. Consider the discourse sample in (2) and imagine the scene created in your mind as you read it.

- (2) A boy with a pail on his head is trying to catch a frog who is falling off a tree stump and the boy has accidentally caught his dog with the net.

The listener encounters utterances one at a time in real time. She will attempt to create a conceptual representation of the events the speaker blueprints via the text, and she will do this by employing all of the information at her disposal. Most researchers see the listener as creating some sort of *text representation*. This text representation is something like the gist of the text, a cohesive and complete general picture of the events or other subject matter discussed. The text representation is generally taken to be a set of propositions linked to one another to form a coherent whole. Perhaps the best known such model is that of Kintsch and van Dijk (1978). A text representation for them is a set of propositions, linked together through common referents and other features of events like time or place, and connected to a higher order discourse theme or *macroproposition*. Other models (Gernsbacher, 1990; Kintsch, 1988; Reinhart, 1981; Vallduví, 1992) follow this general pattern. None of these models grapples with the larger problem of how the text representation – the understanding of the message – is tied to developing the final conceptual representation which the text blueprints, for instance the (possible) mental image resulting from comprehending (2).

In dealing with the text blueprint as it arrives in dynamic time, the listener employs at least three distinct kinds of knowledge to decode the text and integrate its information into a text representation. The listener exploits the *morpho-syntactic details* of her grammar as they code the flow of information in the text and discourse. She entertains *implicatures* derived from the pragmatic interpretation of the sentences in the text against the wider discourse context. And, she employs general processes of *inference* to help make wider connections between the current text and her larger store of knowledge.

Morpho-syntactic coding can be exemplified by something like the particle *wa* in Japanese (Hinds et al., 1987). While the details are not agreed to