by all researchers, speakers of Japanese appear to use *wa* to direct the listener's attention to the key referent in the utterance, its *topic* as some argue. Listeners who encounter *wa* automatically treat the noun phrase so marked as the key referent of the clause, integrating following information against this central referent. These decisions are largely automatic and non-consciously managed by speaker and listener.

Implicatures can be exemplified by something like the English conjunction *and*. In narrative discourse, when two clauses are conjoined by *and*, it is common for a listener to conclude that the event reported by the second clause occurred after the event reported by the first. However, the temporal order of events is *not* part of the meaning of *and*, since there are many uses of *and* which do not involve temporal order at all. The understanding of order derives from implicature, an understanding of event order arising from the interplay of the basic meaning of *and* with the larger context in which the conjunction is used (Levinson, 1983).

Planning and inference can be exemplified by simple inference across clauses. Suppose Mary has a new acquaintance, Sita, and has learned enough about her to know that Sita was raised in one of two national capitals, London or New Delhi. Upon hearing Sita utter (3a), Mary draws the inference that Sita is from London, permitting Mary to utter (3b):

(3a) I used to enjoy playing in the snow in our garden when I was a child and seeing all the trees covered in snow.

(3b) Did you ever see the Trafalgar Square Christmas Tree all covered in snow?

The inference Mary draws that Sita was raised in London derives not from any facts about language, but through Mary's employment of world knowledge while interpreting Sita's observation in (3a). Such inferences are important for knowledge integration, but the processes are strictly speaking non-linguistic.

It is important to see that all of these processes are involved in knowledge integration. But it is also important to see that these contributions are distinct. It is easy for the linguist to include facts about implicature and inference in descriptions of the function of linguistic forms in discourse. It is easy for the psychologist or psycholinguist to conclude that his or her observations about language use are facts about the linguistic system when they might be due to non-linguistic processes of planning or inference. Two excellent resources to help guard against such confusions can be found in Levinson (1983) or Leech (1983).

There are two useful models of knowledge integration in comprehension: Gernsbacher's (1990) structure building model and Kintsch's (1988) construction-integration model. In Gernsbacher's *structure building model*, the listener's goal is to build a coherent mental representation or 'structure' of the information in the discourse. In order to achieve this goal, the listener uses many general cognitive processes and mechanisms. Some processes and mechanisms are involved in 'laying a foundation' for the mental structure.

Once the foundation has been laid, the listener develops her mental structure by *mapping* incoming information onto the previous structure. However, the new information can only be mapped onto a current structure if it coheres with earlier information. If the new information is less coherent, listeners must *shift* to begin building a new structure.

In Gernsbacher's framework, these mental representations are built from memory cells which are activated by incoming information. If that information is coherent with previous information, it is more likely to activate similar memory cells. But if it is less coherent, it is not as likely to activate similar memory cells. Instead, it will activate different memory cells, and these newly activated cells become the foundation for a new structure or substructure. When memory cells are activated, they transmit processing signals which either enhance or suppress the activation of other cells. A group of memory cells is enhanced as long as the information they represent is needed in building the mental structure. When that information is no longer needed, they will be suppressed.

Kintsch's (1988) *construction-integration model* also deals well with knowledge integration. This model attempts to describe explicitly how knowledge is retrieved from memory and utilized in understanding utterances. Consider the sentence, *Mary baked a cake for Sally and burned her fingers.* In order to understand the utterance, the listener needs to know more than merely the words and phrases that were uttered by the speaker. She also needs to have a good deal of general knowledge about how the world works, in this case, that baking entails that the object will be very hot for a period of time. She must also know how language works, for example, that the verb *bake* requires an agent and that this role is filled by Mary in this sentence. Further, the listener needs to know specific information about the situation in which the words were uttered. It is not clear from the words alone whether Mary baked the cake as a gift for Sally, or whether Sally was obligated to make a cake and Mary did it in her place.

Kintsch's model of knowledge use in discourse comprehension has two stages. In the first stage, the words in the utterance are used as the raw material from which a mental representation of the meaning of the utterance is constructed. This mental representation is a network of linked propositions called a text base. In the second stage, the network is edited and integrated with other knowledge stored in memory. Each proposition in the utterance activates its closest neighbors in a general knowledge network. This process of spreading activation results in a text base which contains not only the propositions that were uttered by the speaker, but propositions retrieved from knowledge stores which are related to the propositions in the utterance. Thus, after the second stage of processing, the text base contains clusters of related propositions that combine the information in the utterance with the world knowledge and language knowledge stored in memory.

As we have seen, knowledge integration involves the meshing of nonlinguistic or real world knowledge and knowledge about one's language

with the actual utterances in the discourse. The speaker must employ non-linguistic knowledge to observe and understand events in the world and their relevance to the listener. Then he must use this knowledge and knowledge about his language to choose particular linguistic structures which will be informative to the listener. The listener, for her part, must interpret these linguistic structures using her own linguistic and non-linguistic knowledge. This task is made easier if the speaker manages well his task of providing the listener with appropriate information. In the next four sections, we will explore how speakers manage the task of controlling information for their listeners.

The Rhetorical Management of Discourse

While a detailed discussion of rhetorical management falls outside the scope of the present chapter (but see Gill and Whedbee, Chapter 6 in this volume, for a rich discussion), most of the key concepts discussed below depend on their rhetorical setting for a full understanding. The process of speaking involves both information and action. The informational component includes the details of propositional content as well as pragmatic matters – emphasis, importance, presupposition – which guide how the semantic content should be interpreted. The action component includes the details of discourse planning – both global and local – which help direct pragmatic matters for the speaker and help constrain interpretation by the listener.

It is well known that the use of linguistic structures in discourse is related to linguistic actions taken by the speaker. At the sentence level, the linguistic action of, say, *issuing a command* can be carried out through a number of linguistic structures: an imperative (*Give me your money*; *Let me have your money*), an interrogative (*Could I have your money*), or a declarative (*I want your money*). These examples demonstrate that the form of an utterance is separable from the action, in this case the *speech act*, which the utterance carries out.

The key insight of speech act theory (Levinson, 1983; Searle, 1969; 1979) is that language is used to *do* things (Austin, 1962). Speech act analysis of discourse focuses on local matters affecting clause or sentence type. But language as action is reflected in higher level aspects of discourse organization as well. For example, Swales (1981) examined some 48 introductions to scientific and technical articles. He identified four crucial component actions within each introduction. These actions, which Swales called *moves*, capture critical kinds of information selected by the speaker from his conceptual representation of the subject matter. As shown in Figure 3.3, a typical introduction to a scientific article is composed of four moves: (1) establishing the field, (2) summarizing previous research, (3) preparing for present research, and (4) introducing present research.

There are numerous threads of research, many apparently unknown to each other, which pursue work in this area. Probably the best known is

Subject	HLA antigens in patients with scabies
Move 1 Establish the field	The cell-membrane molecules which are determined by the closely linked genes in the HLA chromosomal complex may be divided into two different classes (Thorsby, 1979): (a) The HLA-ABC molecules, which are determined by allelic genes at the A, B, and C loci, are present on probably all nucleated cells and are highly polymorphic. (b) The HLA-D/DR molecules which have a more restricted tissue distribution are present mainly on B lymphocytes and monocytes/macrophages. Typing for these antigens has become a tool of steadily increasing interest.
Move 2 Summarize previous research	Patients with certain diseases have an increasing frequency of particular HLA antigens compared to healthy individuals (Dausset and Svejgaard, 1997). This is also true for some dermatological diseases. The strongest appears to be the association between D/DR and dermatitis herpetiformis (Solheim et al., 1977), but discoid lupus erythematosus (Stenszky, Nagy, and Szerze, 1975), psoriasis (Williams et al., 1976), vitiligo (Retornaz et al., 1976) and lichen planus (Lowe, Cudworth, and Woodrow, 1976; Halevy et al., 1979) have been found to be associated with certain HLA antigens.
Move 3 Prepare for present research	The reason for these associations are unknown, but probably involve HLA gene control of T cell immune responses (Thorsby, 1978). Immunological mechanisms are also involved in patients with scabies (Mellanby, 1944; Falk, 1980; Falk and Bolle, 1980a, b).
Move 4 Introduce present research	In view of these observations we looked for an association between scabies and any of the HLA-ABC antigens.

Figure 3.3 *Analysis of article introductions (Swales, 1981)*

the British tradition of discourse analysis associated with Coulthard and Sinclair (Coulthard, 1977; Coulthard and Montgomery, 1981; Sinclair and Coulthard, 1975). Their classic work (Sinclair and Coulthard, 1975) examined the structure of classroom discourse in British schools. Mehan (1979) has developed similar lines of research within education.

Researchers in artificial intelligence have examined the goal oriented structure of discourse and its relation to the structure of the knowledge that the discourse is about. Grosz (1974) examines the language used in assembling a water pump, investigating connections between knowledge representations and referential form. Such task-based discourse work has been conducted by Cohen and Perrault (1979), Baggett (1982), McKeown (1985), Sidner (1983), and others.

Linguists have also examined the organization of discourse in this fashion. Early work by Propp (1958) investigated the prototypical organization of Russian fairy tales. Grimes (1975) developed an inventory of

rhetorical predicates to capture the intentional structure of discourse. More recently, Levy (1979) examined the structure of informal interviews of students completing course schedules. Hinds (1979) looked at procedural discourse in Japanese. Mann and Thompson (1986) offer a rich system for describing the fine-grained details of rhetorical management in natural discourse.

All of these efforts propose an inventory of hierarchically organized actions of one kind or another. High level structures can be decomposed into constrained sets of lower order units; lower order units combine in constrained ways to form higher levels of discourse organization. So, Swales's introduction is decomposed into four moves; the four distinct move types combine to form a well formed introduction.

Understanding rhetorical management is important for discourse semantics for a number of reasons. One, the integration of information into or from the text is never merely a matter of processing individual utterances. The utterances are integrated with respect to higher order considerations, and these considerations are what is managed by the rhetorical component. Two, as this volume details elsewhere, there is an important role to be played by the syntax, what we have called morpho-syntactic coding, in signaling one or another information status as the discourse unfolds. The determination of which information is thematic or focused and so on is very much tied to the higher order rhetorical goals for which the discourse was initiated.

The Referential Management of Discourse

One of the characteristics of connected and coherent discourse is that entities, once introduced at a given point in text, are often referred to again at a later point. The problem of how reference is managed in discourse production and comprehension has been the focus of considerable research on discourse processing because it is fundamental in understanding the relationships among cognitive processes, knowledge integration, and information management.

The key insight within referential management is that certain concepts seem to be held in common or shared by both speaker and listener, while others are not. Information held in common forms part of the conceptual scaffolding on which speaker and listener depend for effective communication. The key questions are: (1) what does it mean to say that speaker and listener 'share' information, and (2) how is referential management related to higher level aspects of rhetorical and discourse structure?

Virtually every theory of discourse structure draws a distinction between *given information* and *new information* (also referred to as *old* vs *new*, *known* vs *unknown*, or *shared* vs *new*). Each clause or utterance is theorized to contain elements the speaker believes he holds in common with the listener and elements the speaker believes he does not. So, in the discourse

fragment in (4), the bold-faced NPs are generally taken to be given information and the italicized NPs new information.[3]

(4) Text fragment from popular novel *Sarum* by E. Rutherford (1988: 17)

1 The next day **he** discovered *the lake.*
It was *a small, low hill about five miles inland*
that first attracted **his attention.**
It looked like [*a place from*
5 *which he could spy out the land and*
where they could camp at least for the night].
When **he** reached **the place**, however,
he was surprised and delighted to find
that hidden below **it** and in his path
10 lay *a shallow lake about half a mile across.*
At **its eastern end**, *a small outlet* carried **its water** away towards **the sea.**
Tracing round **the lake**
he found that
it was fed from **the north** and **the west** by *two small rivers.*
15 On **its northern side** was *a flat, empty marsh.*
The water, sheltered by **the hill**, was very still;
there was *a sweet smell of fern, mud and water reed.*
Over **the surface of the lake**, *a heron* rose
and *seagulls* cried.
20 Protected from **the wind** it was warm.
It did not take **him** long to make *a small raft*
and cross **the little stretch of water.**

Overall, this paragraph describes a new location encountered by the main character in the novel. Several of the clauses exhibit straightforward cases of given information as well as straightforward cases of new information.

The bold-faced NPs in 1, 2, 7, 8, 12, 13, 21 represent given information because they have been mentioned before in the text. The italicized NPs in 1, 2, 10, 11, 15, 18, 19, 21 represent new information, for they have just been introduced. Other cases are a bit less clear. The NPs in 3, 11, 14, 15, 20 are marked bold, and their putative status as given must be related to knowledge shared by writer and reader about lakes and their environs. But then lakes also include native birds, like herons and gulls, so one wonders why the NPs in 18 and 19 cannot count also as given.

Such observations have led a number of researchers to propose more or less complex systems of given and new information.

Conceptual Foundations for Given and New Information

There are two basic ideas about given and new information: (1) given information represents a referent shared in some way by speaker and listener; or (2) given information is a cognitively activated referent.

Given Information as Shared Information Traditionally, referential management is taken to require that a given semantic argument also hold a pragmatic status like *old* or *given* or *known* information. Within the Prague School, Mathesius (1939) suggests that one portion of the utterance represents information that is assumed to be possessed by the listener from the preceding context or may be inferred by her from the context. Such information is *known* (*old, given*) information. It is contrasted with the portion of the utterance which the speaker presents as *new* (*unknown*) information and which is the content of the utterance. Mathesius examined how this status of information is signaled via strategies such as word order, intonation, and other constructions. These ideas were developed by other Prague School scholars, such as Daneš and Firbas.

Halliday (1967a; 1967b) is concerned with relating each unit of information in a given sentence to the preceding discourse. He draws a distinction between *given* information and new information. New information represents information the speaker treats as not known to the listener. Given information represents information the speaker treats as known to the listener. Halliday links the status of new information to focal sentence intonation. Unlike Prague School researchers, Halliday draws a further distinction between *known* and *unknown* information. For Halliday, information is known if the speaker assumes the listener can identify the referent and is unknown if the speaker assumes the listener cannot identify the referent. DuBois (1980) also considers the importance of identifiability in referential management.

Prince (1981) finds these intuitively appealing notions to be too simplistic. She proposes a multi-way distinction in types of information (types of statuses of referents). One, a referent is *new* when it is introduced into the discourse for the first time. New referents may be *brand-new*, that is newly created by the speaker, or simply *unused*, that is entities the listener is assumed to know about but which have not been mentioned previously in the discourse. Two, a referent is considered *evoked* if it is already part of the discourse. An evoked referent may be *textually evoked* if the listener had evoked it earlier on instructions from the speaker (as by the speaker's mention of the referent), or it may be *situationally evoked* if the listener knows to evoke it all by herself, such as 'you' referring to the listener. Three, a referent is *inferable* if the speaker assumes the listener could have inferred it, using knowledge and reasoning. A referent may be inferable either from the text or from the situation.

Given Information as Degree of Memorial Activation Chafe (1976; 1987; 1994) discusses information status in terms of what is activated (or not activated) in consciousness. He argues that the linguistic phenomena such as *given* and *new* information are manifestations of our basic cognitive activities. Our minds contain a very large amount of knowledge or information but only a very small amount of this information can be focused on, or be 'active', at any given moment. He proposes that a particular

concept may be in any one of the three different activation states at a particular time of discourse processing: *active* (corresponding to the 'given'), *semi-active* (accessible), or *inactive* (corresponding to the 'new').

> An active concept is one that is currently lit up, a concept in a person's focus of consciousness. A semi-active concept is one that is in a person's peripheral consciousness, a concept of which a person has a background awareness, but which is not being directly focused on. An inactive concept is one that is currently in a person's long-term memory, neither focally nor peripherally active. (Chafe, 1987: 25)

A speaker normally makes changes in the activation states of certain concepts which are partially reflected in their referential choice. If the speaker assumes, prior to uttering an intonation unit, that a concept is already active in the listener's mind, he will verbalize that concept in an attenuated manner, most probably pronominalizing it. If he assumes that a concept is not presently activated in the listener's consciousness, he will verbalize that concept in a less attenuated manner, most probably nominalizing it.

Clark and Haviland (1974) relate these notions to memorial processes in their discussion of the 'given–new strategy'. That is, each sentence produced by a speaker contains some information that is old or given, and some that is new. The old information serves as an indication of where, in the listener's memory, she will find information related to that conveyed by the present sentence, and thus 'an instruction specifying where the new information is to be integrated into the previous knowledge' (1974: 105). Consequently, pronouns and definite noun phrases (NPs) are more likely to refer to old or given entities and indefinite NPs to new information.

Givón (1983) also considers referential management in cognitive terms. He observes that the speaker estimates to what extent a given referent is mentally accessible to his listener. If accessibility is gauged to be high, the speaker will use an attenuated referential form to index the referent (ellipsis or pronominalization). If accessibility is judged to be lower, the speaker will use a longer form, perhaps a simple nominal NP or one with some modification. If accessibility is estimated to be very low, the speaker may introduce a referent into the conceptual representation through (at least in English) an indefinite NP or some other appropriate device.

Referential Management and Knowledge Integration

One important problem in reference management has been understanding how speaker and listener keep track of referents during discourse production and comprehension. Keeping track of referents involves three related problems: (1) introducing referents to the discourse, (2) sustaining reference once a referent has been introduced, and (3) reintroducing referents after a long hiatus. Virtually every approach at present employs some notion of managing a mental model or conceptual representation for this purpose. Speakers will use particular linguistic forms (see Cumming and Ono, Chapter 4 in this volume) to introduce referents to the discourse,