

Chapter 5

Textual Resources

1. The Nature of Textual Systems

Textual semantics concerns the role of the text and its components as a *message*. While creating a text (whether a single utterance or a whole book), we have a certain amount of content we wish to encode. But there are various ways to encode this information, to present our message. The textual semantics represents the various strategies for structuring the message. The focus is on the text as an object of communication, and how the text is structured to effectively achieve communication.

While WAG is a system for sentence representation and processing, it is intended to function as part of multi-sentential processing systems. The micro-semantics thus needs to be responsive to multi-sentential concerns, so that it can offer the types of controls over the micro-representation that are needed by a higher-level processor. The textual semantics is important in this regard, since this meta-function is responsible for coherence and cohesion across sentences.

1.1 Textual Semantics from Above and Below

Matthiessen (1990) points out the difference between the semantics which results when we look at the types of distinctions necessary to motivate differences in form, and the semantics which results from projecting socio-functional needs onto language:

“(i) We can explore semantics from below, starting with the lexico-grammar ... -- what might be called a **decoding** or interpretive approach, since it works by decoding or interpreting lexico-grammar in semantic terms.

(ii) Alternatively, we can explore it from above, from outside the linguistic system -- what might be called an **encoding** approach since it looks at semantics as an encoding strategy and explores how contextual categories are encoded semantically.” (1990, p2)

The approaches differ in focus -- the decoding approach focuses on the lexico-grammatical resources. Their semantic function is secondary. The question to be asked, according to Matthiessen, is: *What is the interpretation of this word or grammatical category?* The semantics is built up as an explanation of differences in lexico-grammatical forms.

The encoding approach, on the other hand, views the text as a socio-culturally structured message:

“[the encoding approach] takes the context as its starting point and brings the communicative purpose into focus. The semantics constitutes the strategies for achieving communicative purposes in context.” (Matthiessen 1990, p7)

In relation to textual semantics, the decoding approach focuses on lexico-grammatical regions such as reference, substitution, ellipsis, conjunction and grammatical metaphor

(cf. Martin 1992; Halliday & Hasan 1976). Encoding approaches, on the other hand, tend to center on topics such as rhetorical structure, information structure, thematic structure, conversational implicature, etc.

Throughout this chapter, I will take an encoding approach, focusing on the contextual orientation of the textual resources. The textual resources will however be related to their lexico-grammatical realisations.

1.2 Macro- and Micro-Textual Resources

From the encoding view, we can deal with textual resources at two levels -- macro-textual (multi-sentential resources) and micro-textual (single-sentence resources). Textual Semantics primarily concerns macro-representation. However, this textual structure is realised through individual sentences, and we can look at how each sentence is organised to realise the macro-textual structure.

1.2.1 Macro-Textual Resources

I will describe briefly how multi-sentential text is organised as a message, focusing on three dimensions of the macro-textual structure:

1. **Thematic Structure:** how the text is structured to develop a Theme or number of Themes. The speaker/writer chooses a central Theme or Themes to organise the discourse, and each sentence expands upon one of these Themes.
2. **Rhetorical Structure:** how the text is organised to meet the rhetorical goals of the speaker. Each unit of text is serving some function towards these goals.
3. **Information Status:** the presentation of information as *recoverable* (already mentioned in the discourse, or part of the immediate non-verbal environment of the discourse), and *identifiable* (part of the shared knowledge of the speaker and listener).

1.2.2 Micro-Textual Resources

The text is composed of sentences, and it is largely through sentences that the macro-textual meanings are realised. The following discusses those aspects of macro-textual organisation which are visible in sentence-size units -- the micro-textual semantics:

1. **Themacity** (Thematic Structure): Is the information central to the speaker's thematic development strategy at this point of the text.
2. **Relevance** (Rhetorical Structure): Is the information relevant to the satisfaction of the discourse goals of the speaker? -- will the communication of the information aid in the developing rhetorical structure?
3. **Recoverability & Identifiability** (Information Structure): Is the information identifiable/recoverable from the context? The context here can be the prior text (i.e., information that has been introduced earlier in the text), or the shared knowledge of the participants (the speaker assumes that the information is known to the listener, and can be recovered easily in the context, e.g., "the Prime Minister").

These micro-textual patterns will be explored below.

2. Themacity

2.1 Macro-Thematic Structure

The patterning of Themes through a text is one means of structuring the text as message -- organising the text so that the listener/reader is aware of what the speaker/writer is trying to achieve. This thematic patterning is called *thematic structure* (Halliday 1985), *thematic progression* (Danes 1974), or *method of development* (Fries 1981; Martin 1992).

Martin's (1992) discussion of method of development brings to the fore the fact that thematic structuring is oriented towards the effective communication of ideation:

"Method of development ... establishes an angle on field. ... [it] is the lens through which a field is constructed; of all the experiential meaning available in a given field, it will pick on just a few, and weave them through Theme time and again to ground the text -- to give interlocutors something to hang onto, something to come back to -- an orientation, a perspective, a point of view, a perch, a purchase." (p489)

We might find for a particular text that the Theme begins on *Health care*, moves on to a *virus* and how it spreads, and thence onto *prevention* of viral spread. The same information could alternatively be presented with a different method of development, thus leading us to read a different message from the text.

Danes (1974) proposes three typical patterns of thematic progression:

- 1) **Constant Theme:** Theme stays constant over sentences -- the speaker's rhetorical strategy is to develop more fully a single participant.

Mark gave the book to Paul. Then he came here.

- 2) **Simple Linear:** The Theme is drawn from the Rheme (the non-Theme component) of the prior sentence:

Mark gave the book to Paul. Paul will be here soon.

This strategy is often used to develop a chain of relevance: using the known to introduce the unknown. In the above case, the speaker used *Mark* as a starting point, to introduce *Paul*. The speaker could then use *Paul* as the centre of development.

- 3) **Derived Theme Progression:** The Themes of successive sentences are derived from a single over-riding Theme. For instance, each successive sentence develops some characteristic of one entity:

Mark is a fine bloke. His manner is good. His main problem is that he is too generous with his books.

Another pattern, not mentioned by Danes (1974), involves the new Theme drawing on both the Theme and Rheme of the prior sentence:¹

Theme_i+Rheme_i => Theme_{i+1}

Mark left Mary. They had been fighting.

Martin (1992) introduces a higher level of thematic organisation, based around what he calls a *hyper-Theme*, which is:

"an introductory sentence or group of sentences which is established to predict a particular pattern of interaction among strings, chains and Theme selection in following sentences." (p437).

¹Original source unknown.

A text usually intermixes these thematic progression patterns -- for instance, one may start with a simple linear pattern then move to a constant Theme for a few clauses. The Theme may pick up from the immediately prior sentence, or continue from a Theme several sentences prior.

Thematic progression is probably best viewed from the speaker's perspective (the generation perspective). At a local level, it is a concern about what will be the Theme of the next utterance. However, when used by a good speaker/writer, thematic progression is concerned with more than the choice of the one Theme. Sequences of thematic choices are planned to achieve some effect. Churchill's famous war rhetoric uses a long sequence of constant Theme to emphasise the collectivity and solidarity of the intended audience (the British public):

"We shall defend our island, whatever the cost may be, we shall fight them on the beaches, we shall fight them on the landing grounds, we shall fight in the fields and in the streets, we shall fight in the hills; we shall never surrender." (Churchill 1940).

In the analysis direction, we can examine the thematic progression as a means to recover the writer/speaker's intention in communicating. In Halliday's words:

"...by analysing the thematic structure of a text clause by clause, we can gain insight into its texture and understand how the writer made clear to us the nature of his underlying concerns." (1985, p67).

2.2 Micro-Thematic Structure

The Theme of a sentence is the point of departure for the message -- what the sentence is about. In *Mark has the virus*, we are talking about Mark, what has happened to him. But in *The virus has got to Mark*, we are centering on the virus, and how it is spreading.

In English, Theme is marked by initial position in the clause (Halliday 1985). However, in other languages, other means are used to mark Theme. Martin (1983b), for instance, mentions that Theme in Tagalog "is marked with the particle *ang* and typically appears in clause final position."

My implementation of Theme is restricted to what Halliday (1985) refers to as 'Topical Theme' (the themacity of participants and circumstances in the sentence). He also makes use of 'Interpersonal Themes' (the themacity of interactional information, such as modality -- "Probably..." -- or attitude -- "Unfortunately..."), and 'Textual Themes' (the themacity of markers of rhetorical structure, for instance, marking the role of a sentence in relation to the rest of the text, e.g., "In addition..."). These types of Themes will not be considered here.

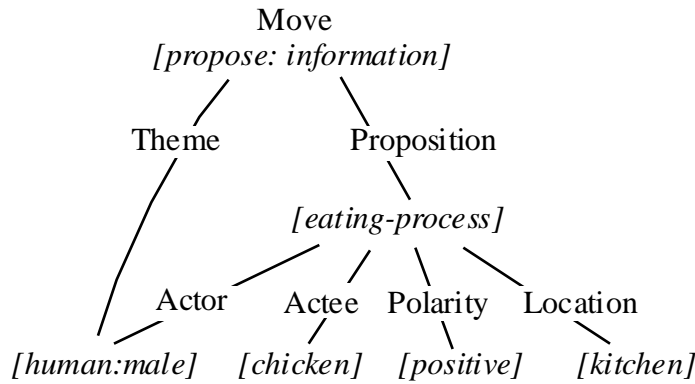


Figure 5.1: Showing Theme specification for
“He ate the chicken in the kitchen.”

It needs to be shown how Theme is represented. Like other textual patterns, Theme is a dynamically evolving orientation towards the ideation-base: each successive sentence preserves or changes the Theme, manipulating the hearer’s focus on the ideation which is being constructed. In the WAG implementation, the Theme is a *role of the speech-act*. The role is basically a pointer to the ideational entity which is thematic. For example, see figure 5.1.

In this diagram, we have an informing move, which would be realised as “He ate the chicken in the kitchen”. The Theme role of the speech-act conflates with the Actor role of the Proposition. Figure 5.2 demonstrates a different Theme selection, where the Actee conflates with Theme.

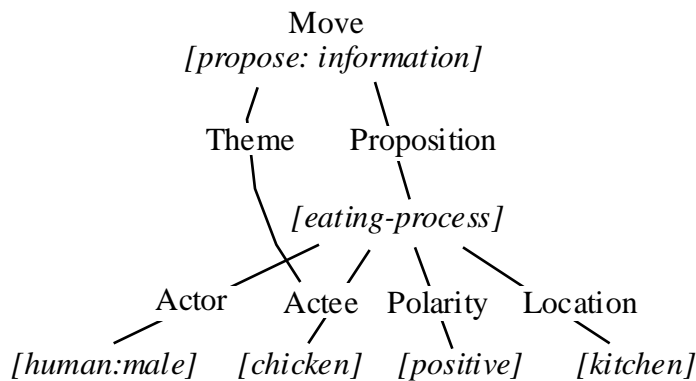


Figure 5.2: Showing Theme specification for
“The chicken was eaten by him in the kitchen.”

This approach to Theme differs from the Halliday’s, and the Penman, approach to Theme, where Theme is treated as an element of lexico-grammatical structure. In the WAG system, Theme is an element of semantic structure. The choice of semantic Theme is realised by influencing various ordering decisions in the lexico-grammar, e.g., active vs. passive, fronted complements, fronted circumstances, etc.

In some instances, the element of the proposition specified to be Theme is not at the top-level of the proposition. Take for instance the move specified in figure 5.3, where the Origin of the Actor is specified to be the Theme of the move. This could be realised as: “Sydney is where the man who ate the chicken came from.”² The sentence is awkward,

²This form of thematicisation has not yet been implemented within the WAG resource model.

but this is the result of selecting a Theme which is not a top-level element of the proposition.

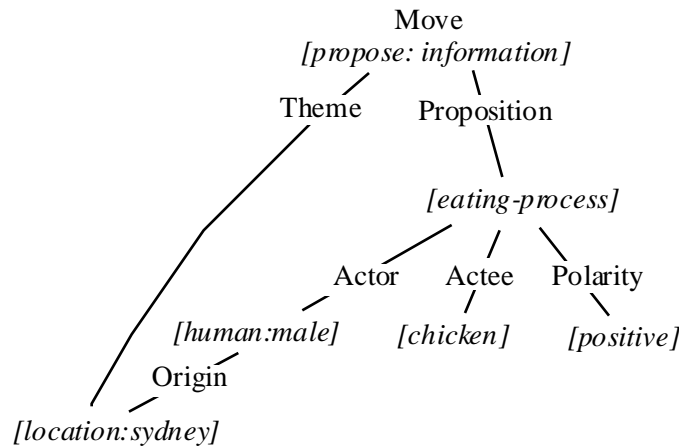


Figure 5.3: Move Structure for "Sydney was where the man who ate the chicken came from."

3. Rhetorical Relevance

Rhetorical Structure Theory (RST) (cf. Mann & Thompson 1987; Mann & Matthiessen 1987; Matthiessen 1988b; Mann, Matthiessen & Thompson 1992) views text structure in terms of dependency relations between units of text. An RST analysis segments a text, and relates these segments by showing the *rhetorical relations* between them. Typical RST relations include *evidence*, *motivation*, *concession*, *purpose*, *elaboration*, etc.

While RST is useful for discourse analysis by humans, it is not very useful for computational discourse analysis, since many of the rhetorical relations are not explicitly (or unambiguously) marked in the text, and are thus difficult to recover during the analysis of the text. RST is thus more useful in the text generation context, for the construction of coherent text.

3.1 Rhetorical Relevance

One of the main steps in the text generation process involves *content selection* -- the selection of information from the speaker's knowledge base for presentation. Such a process must decide what information is relevant at each point of the unfolding discourse. Relevance is here defined as working towards the achievement of the speaker's discourse goals.

In some systems, content selection is driven through the construction of the rhetorical structure of the text (e.g., Hovy *et al.* 1992). As we build a rhetorical structure tree, the ideation which is necessary for each rhetorical relation is selected. For instance, if we add an *evidence* relation to an existing RST tree, the ideation which functions as evidence is selected for expression. The rhetorical structure thus organises the ideational content to be expressed, selecting out those parts of the ideation which are *relevant* to the achievement of the discourse goals at each point of the text. I use the term *rhetorical relevance* to refer to this sort of relevance.

As an example of the interaction between relevance and content selection, consider an information service, where a customer is concerned with the location of the closest vehicle repair shop to their work. The name of the business, its street address, and

telephone number are probably relevant information. The fact that the shop next door is owned by my mum's best friend is perhaps interesting, but not relevant to the customer's goals.

At the sentence level, we use relevance to determine which roles of the proposition are to be expressed, e.g., which participants of the process are relevant (e.g., the Location of the process; the age, color, size or other quality of a participant, etc.). See Pattabhiraman & Cercone (1990) for a good computational treatment of relevance, and its relation to salience.

3.2 Representing Relevance

Rhetorical relevance is dynamic -- it changes as the text progresses. It represents a shifting *focus* on the ideation base (Halliday & Matthiessen, to appear, pp373-380). What is relevant changes as the text unfolds, as the rhetorical structure is realised. Relevance forms what Grosz (1977/86) calls a *focus space*.³ According to Grosz, focus is "that part of the knowledge base relevant at a given point of a dialog." (p353). However, Grosz's notion of relevance is different from that I use. Hers is based on the needs of a text understanding system -- which objects in the knowledge-base can be used to interpret the utterance. My sense of relevance is derived from relevance in generation -- what information has been *selected* as relevant to the speaker's unfolding discourse goals. She is dealing with a set of objects which may *potentially* appear in the text at this point, while I am dealing with the set of objects which most probably *do* appear in the text.⁴

I thus use the term *relevance space* to refer to the set of ideational processes, participants and circumstances which are relevant towards achieving the present discourse goal. The construction of an utterance is constrained by this relevance space, conditioning the decisions whether or not to include particular lexico-grammatical roles.

Figure 5.4 shows the dynamic evolution of relevance-space. The diagrams show a single macro-ideational structure, outlining a sequence of events. The left-hand column shows a text that one of the participant's in this event-sequence produced to describe the events. Each of the diagrams in the right-hand column represents the relevance space at a particular point in the unfolding of the discourse. The speaker developed a particular rhetorical structure to support the story, and this structure determines the relevance of information at each point in the discourse. As the rhetorical function of each unit changes, so too does the relevance status of each item of information.

The macro-ideational structure has been simplified here, since the focus is on textual, not ideational, representation. Note in particular the following simplifications:

- Temporal relations between processes are not shown, for instance the sequence relations between sentences 2, 3 and 4;
- the structure only shows ideation which is realised in the text. For some generation processes, this ideation base would present information not selected for realisation in the text;
- Polarity and Modality are not shown, unless marked (as in the third sentence).

³Various other linguists and computational linguists have also used the notion of 'spaces' to represent textual status, see for instance, Reichman (1978);Grimes (1982). Halliday & Matthiessen (to appear) extend Grosz's notion of focus space to include other types of *textual spaces*: thematic spaces, identifiability spaces, new spaces, etc. (p376). Each of these spaces can be thought of as a pattern stated over the ideation base (p373). This notion will be used later in this chapter.

⁴I say probably since even though information may be relevant to the speaker's current discourse goals, it might fail to appear in the text, for instance because the linguistic resources may not allow the expression of all the relevant information (the expressability problem). Alternatively, some relevant information is not expressed because it is totally recoverable from the context, and thus does not need to be stated.




Figure 5.4: Dynamically Evolving Relevance Space

3.2.2 Computational Representation

A relevance-space is basically a *set* of ideational entities. Since WAG doesn't support sets at present, a relevance-space is represented as a *list*. This list is held in a variable, named **relevant-entities**. For sentence generation, this list represents the entities to be included in the sentence being produced. The list is established before starting sentence generation, typically by including a field in the micro-semantic specification (the input to the generator -- see chapter 11 for detailed discussion), as in the following example:

```
(say utterance-1
  :is (:and initiate propose)
  :proposition arrival-event
  :speaker Caller
  :relevant-entities (arrival-event morning Mark) )
```

During the generation process, feature selection constraints (see next chapter) can query this list, to determine which of the elements of the ideation-base are to be expressed. The *:relevant* operator is used for this purpose. For instance, the following constraint tests whether genitive deixis should be used in a nominal group:

```
(:and (:exists Referent.Owner)
  (:relevant Referent.Owner) )
```

Interpretation: if the referent (ideational content) of this nominal-group has an Owner role, and the filler of that role is on the **relevant-entities** list, then the constraint is met.

In the analysis direction, the list represents those entities which were expressed in the sentence being analysed. Constraints like the one above, rather than testing membership of the relevance space, would assert the entities membership onto this list.⁵

The representation of relevance-space using a list of ideational-entities is not a totally ideal representation. It is often the case that an entity plays multiple roles in the ideation base, and the entity is relevant in only one of these roles for a given utterance. For example, take a case where Mark owns both a dog and a house, and the dog destroys the house. If we nominate *Mark*, *dog* and *house* to be relevant, then the system as implemented would provide output like: *Mark's dog destroyed his house*. However, we might have decided that Mark's ownership of the dog was not relevant, thus desiring output more like: *A dog destroyed Mark's house*. The relevancy list thus needs to be re-represented as a list of the relevant roles of each entity.

3.3 Lexico-Grammatical Resources for Marking Relevance

The main way of marking a particular item as relevant is to include it in the text. If we don't include it, it is probably not relevant. We can avoid the expression of an entity using a number of strategies:

- the non-insertion of an optional grammatical role (e.g., circumstances);
- the ellipsis (deletion) of an element which is inserted;
- the use of a grammatical structure which doesn't express a role, e.g., using an intransitive verb to avoid expressing an Agent, such as in "Mark died."

The non-expression of information does not necessarily mean it is not relevant -- information that is recoverable from context need not be said. We thus need to see the

⁵Some constraint operators can switch between a closed-world interpretation (return true only if the condition is already true) and an open-world interpretation (will change the knowledge-base to reflect the constraint, unless this contradicts existing knowledge. A flag in the WAG Knowledge Representation System allows different processes to choose between open- and close-world processing.

decision to express an entity as an interaction between relevance and recoverability (see next section).

The *degree* of relevance may also be marked by the *linguistic prominence* of the item. There are various different strategies for making an item linguistically prominent: lexico-grammatically, information is presented more prominently at the end of the clause (see Halliday 1985, p275), or by using marked structures, e.g., *It is Mark who did it!*. Prominence can be lowered by *lexico-grammatical downgrading*, e.g., parenthesizing, rank-shifting or nominalisation. Prominence can also be marked phonologically (intonational stress), or graphologically (e.g., bold, underline).

3.4 Marking Rhetorical Function

The role the text-segment plays in the rhetorical structure -- its *rhetorical function* -- is often marked explicitly, for instance, via a clausal conjunction (e.g., *In summary*, *Consequently*, *Our first argument*, etc.). There is a wide variety of lexico-grammatical resources for encoding the rhetorical function, which I will not discuss here, but see Martin (1992; 1983a). I have not yet addressed the issues involved in encoding rhetorical function in lexico-grammatical structures.

4. Information Status

The participants in an interaction each possess a certain amount of information, some of which is shared, and some which is unshared. Figure 5.5 shows this state of affairs. The term *information status* is used to refer to the state of information as either shared or unshared.




Figure 5.5: Information Sharing Between Two Participants

I will focus on the information status of ideational entities -- and how their information status affects how they are referred to in discourse. Two factors will be considered -- identifiability and recoverability. This section is based loosely on Martin (1992) and Halliday & Hasan (1976). See also Dale (1988a, 1988b, 1992).

4.1 Identifiability

The information status of an entity affects the way we refer to that entity. When a speaker mentions an entity they believe to be shared, they use a form of expression which marks the shared nature, the fact that the speaker expects the entity to be one of the entities within the hearer's information-base. In other words, the speaker expects the hearer to already have access to the referent. In addition, the speaker is indicating that within the context of expression, the hearer should know exactly which referent is intended. For instance, the use of definite deixis in *the President* indicates that the hearer should be able to identify which president is intended. Two common forms are used to signal identifiability:

- **Definite Deixis**, e.g., *the President*;
- **Naming**, e.g., *Ronald Reagan*;

Non-identifiability is signaled through indefinite deixis, e.g., *a boy called John; Eggs; Some eggs*, etc. A speaker uses indefinite deixis to indicate that he believes the entity not to be known to the hearer. It is thus a strategy used to introduce unshared entities into the discourse. Once the entity is introduced, some form of definite reference is appropriate.

Identifiability is a quality of the expression, not the hearer's knowledge itself: it is possible to refer to a shared entity in such a way that the reference is not identifiable, perhaps because inadequate clues are given to uniquely identify the intended referent. For instance, *that boy on Smith street* may not identify the referent, even if I know the boy in question, perhaps because there are several boys on Smith street, or perhaps because, while I know the boy, I don't know where he lives (see Goodman 1986).

Identifiability also doesn't depend on the referent being known to the hearer, as long as the referred entity is somehow related to a known entity, and the relationship is stated in the expression. For instance, if I know John then *John's mother* is an identifiable reference.

The discourse context is important in determining the identifiability of a reference, for instance, *the President* may not be identifiable in the middle of a cooking class, but while discussing US or French politics, it will be identifiable. Expressions are just sign-posts, and will identify different entities in different contexts.

In summary, the identifiability of a reference depends on three factors:

- a) The hearer's information-base;
- b) The form of the expression used;
- c) The discourse context.

We also need to distinguish the *identifiability* of a reference from the speaker's *indication of identifiability*: it may not be possible for a speaker to identify a reference, even though a definite reference form was used. The definite form indicates that the speaker expects the hearer to be able to identify the reference, it doesn't indicate that the hearer actually will identify the referent. For several reasons they may fail, possibly because the speaker was wrong in their assessment of the hearer's information state, and may use definite or indefinite reference inappropriately. Sometimes, also, the speaker may falsely indicate the identifiability of a referent, perhaps for rhetorical purposes, or perhaps as a joke.

4.2 Recoverability

Shared information can be divided between that which is part of the *immediate discourse context*, and that which is not. In relation to information sharing, the immediate discourse context includes:

- information introduced earlier in the discourse;
- the discourse participants (speaker, listener, or speaker+listener);
- the immediate physical context of the discourse -- those entities which the participants can point at, for instance, a nearby table, or some person.

Entities which are part of the immediate discourse context are more prominent in the awareness of the participants, and so less sign-posting is required for a speaker to refer to one of these entities -- less specific information needs to be included in the reference. Such entities are termed *recoverable*, since the hearer can recover an entity from just a non-specific pointer towards it. The most common forms of reference for recoverable entities include:

- pronominalisation, e.g., *she, them, it, this*, etc.
- substitution, e.g., *I saw one*;
- ellipsis (the non-mention of an entity), e.g., *I am going to the shop*.

These forms are reserved for recoverable reference, and cannot be used for more general identifiable reference. A common feature of these forms is that they provide insufficient information to identify any entity not in the immediate discourse context. For instance, a pronoun will usually only indicate some combination of gender, number, proximity, or speech-role. By including only minimal information, the speaker signals to the hearer that the referent *must* be part of the immediate discourse context. The hearer must then try to locate (recover) the most salient entity in the discourse context which matches the semantic types of the expression. For instance, when a speaker uses the term “she”, the hearer looks for the most salient female in the discourse context, perhaps the most recently mentioned female, or perhaps someone present in the immediate physical context.

Other forms of identifiable reference (e.g., naming, definite deixis) can also be used for recoverable entities, but these forms are not restricted to recoverable reference, since they can be used to introduce entities which are shared, but not yet introduced to the discourse. However, when these forms *are* used to refer to recoverable entities, the reference form tends to be less specific, since the intended referent is contextually salient. For instance, if definite-deixis is used, usually only a super-ordinate of the term is necessary, e.g., *the boy*; *the vehicle*, etc. Super-ordinate reference is generally not appropriate to introduce an entity which is not part of the immediate discourse context.

Recoverability can be complicated by the following two phenomena, neither of which is appropriately handled in the WAG system as yet:

- **Reference Shadow:** A *reference shadow* occurs when the referential device the speaker intends to use to refer to an entity applies also to some entity which has been mentioned more recently. In such cases, a more specific type of reference is required, which more uniquely refers to the intended entity. There is also a problem with reference ambiguity, where two entities who were last referred to in the same utterance are picked out by the same referential form.
- **Recoverability Decay:** As a discourse continues, the ‘presence’ of entities in the discourse-context decays. After a certain time, references based on the supposed recoverability of an item may fail. For instance, ten minutes after the last mention of someone “he” will probably not adequately refer. I call this phenomenon *recoverability decay*.

4.3 Representing Recoverability and Identifiability

Recoverability and identifiability can be represented in a similar way as relevance: using textual spaces. The WAG system allows the user to specify which entities are candidates for identifiable reference (part of the shared information-base), and which are recoverable (part of the discourse-context). These reference spaces evolve dynamically, each sentence introducing new information, raising the information status from unshared to shared, or non-recoverable to recoverable.⁶

I use a *shared-entity* space, rather than an *identifiability* space, because, as stated above, identifiability is a characteristic of expressions, not of entities. Shared entities are *candidates* for identifiable reference. Other candidates for identifiable reference include the entities which are related to shared entities, if the relations are explicated in the expression, e.g., *John’s mum*.⁷

The diagrams of figure 5.5 demonstrate the evolving shared-entity and recoverability spaces of the various discourse entities. These spaces are marked as follows:

⁶At the same time, a process of recoverability decay moves some entities out of the recoverability space. These entities, however, remain in the shared-entity space -- a definite reference to the entity will bring it back into the recoverable space.

⁷This feature has not yet been added to WAG.

- Circled elements are part of the shared-entity space, and are thus candidates for definite reference.
- The circled elements with white background are part of the shared-entity space, but not the recoverability space, i.e., they are part of the shared knowledge of the participants, but have not yet been introduced to the discourse. Definite reference can be used -- the participant does not need to be introduced using indefinite reference. However, at this level, entities cannot be pronominalised.
- The circled elements with gray background are part of the discourse-context, and thus in the recoverability space (and thus also the shared-entity space. They are thus candidates for pronominalisation, substitution, presumption, or superordinate reference, in addition to the other forms of identifiable reference.

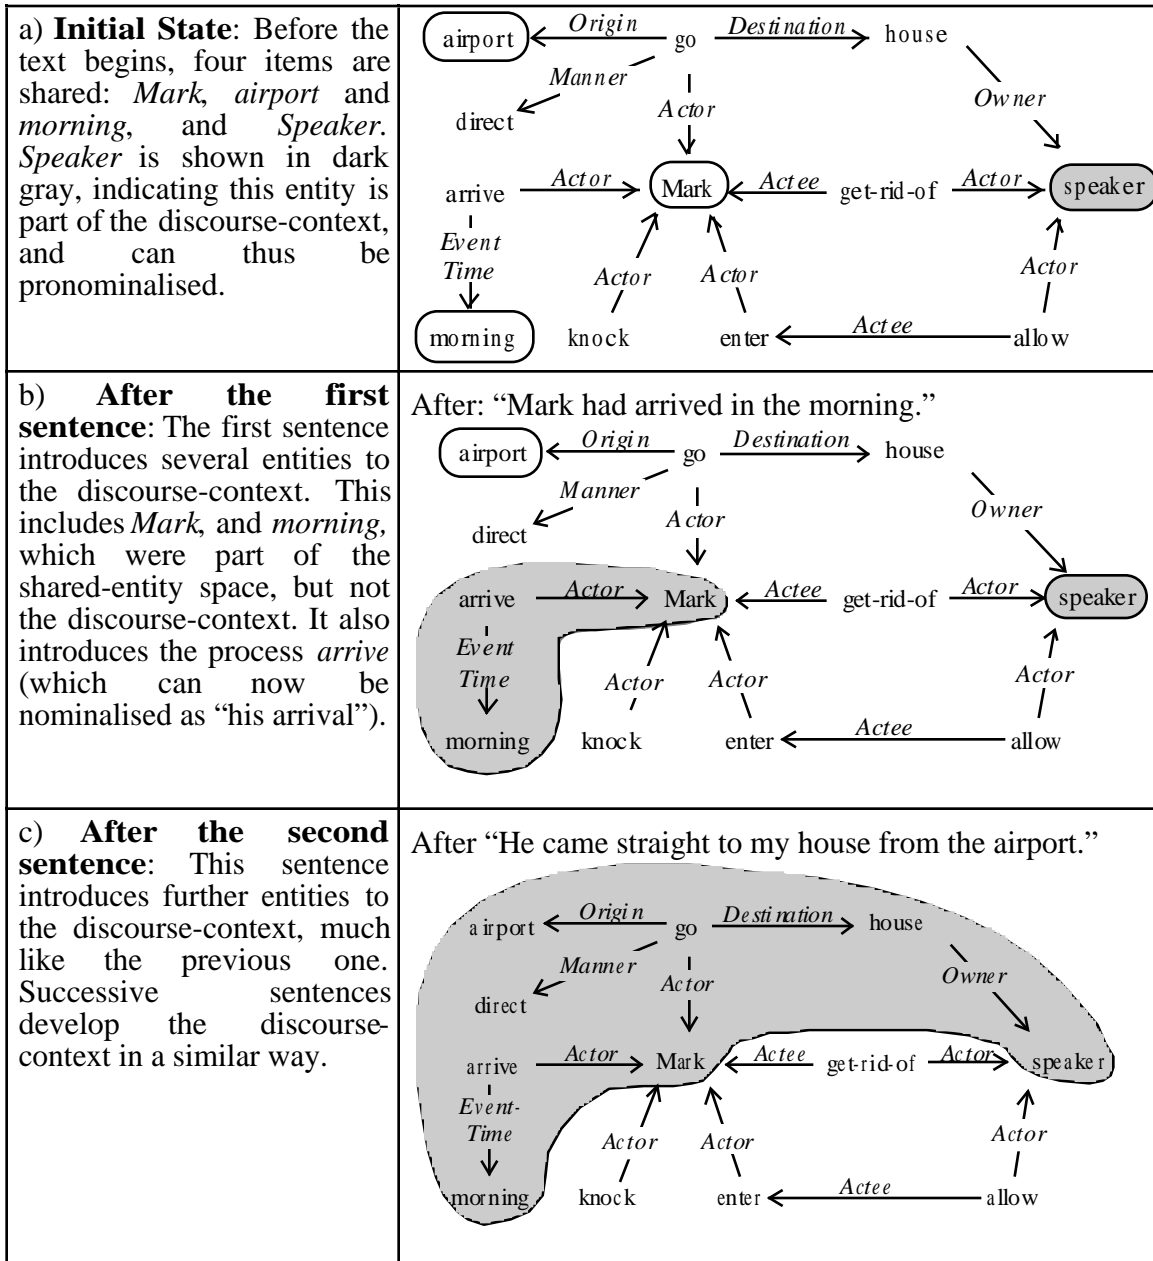


Figure 5.5: Evolving Recoverability and Identifiability Spaces

Computationally, recoverability and sharedness of information are represented as lists of recoverable and shared entities, stored in global variables. For instance, the textual spaces shown in figure 5.5 (b) would be represented as follows:

identifiable-entities : '(Mark house speaker airport morning arrive go)

recoverable-entities : '(Mark house speaker morning arrive go)

The two lists are mostly identical since most of the identifiable entities are also part of the discourse context, except for *airport*, which is a shared entity, but not yet mentioned.

The system needs to be extended to handle reference shadowing and recoverability decay. One modification would be to maintain a reference stack -- a list of mentioned entities with the most recently mentioned on top. Before using a reference-form to refer to an entity, the program would need to check that this reference doesn't fit any entity which is higher on the stack. If it does, then a more specific reference-form should be used. The reference stack can also be used in analysis to resolve the referent of a mention. To find the intended referent, we try to unify the semantics of the reference with each entity on the stack, working from the top, until one unifies. For instance, the semantics of "he" is [thing: male: human: singular-thing]. We look for an entity on the stack which shares this ideational specification.

For sentence generation, WAG allows the user to include a specification of the recoverability and shared-entity spaces directly in the micro-semantic specification (see chapter 11 for detailed discussion). These specifications are placed into the global variables mentioned above. For instance, the micro-semantic specification for "Mark had arrived in the morning":

```
(say-example
  :speech-function (:and propose information)
  :proposition (Arrive :is (:and motion-process motion-termination)
                    :Actor (Mark :is male)
                    :Event-Time (Morning :is temporal-duration))
  :theme Mark
  :shared-entities (Mark Morning Airport Speaker)
  :recoverable-entities (Speaker))
```

In the constraint language, sharedness and recoverability can be tested using two operators: *:shared-entity* and *:recoverable*. The following constraint tests whether pronominalisation is appropriate:

```
(:recoverable Referent)
```

The following constraint tests whether a proper noun is appropriate (if the entity has a name, and is part of the identifiability space):

```
(:and (:exists Referent.Name)
      (:shared-entity Referent))
```

When an entity is being realised lexico-grammatically, the semantic constraints on the grammatical features will refer to the informational status of the entity: can I use a definite noun phrase? yes if it is on the shared-entity list. Can I use a pronoun? -- yes if it is on the recoverability list (and the item is not in a reference shadow). A more detailed example of how these textual spaces constrain the nominal group structure is provided in chapter 6 on inter-stratal mapping.

5. Summary of Textual Resources

This chapter has offered a summary of the textual resources as implemented in the WAG system. These resources are still at an early stage of development. However, they do offer limited means to provide textual control of sentence generation, and in analysis, allows textual spaces to be recovered.

This chapter made the following points:

- 1) **Two views on the textual meta-function:** Following Matthiessen, I distinguish textual semantics construed to explain grammatical differences, and textual semantics construed to encode context. I have chosen the latter approach, yet relate such a textual semantics to the lexico-grammatical realisations of these resources.
- 2) **Textual Representation:** Following Grosz (1977/86), and Halliday & Matthiessen (to appear), I use the notion of textual spaces, partitionings of the ideation base, which shift dynamically as the discourse unfolds. I have outlined:
 - a) a relevance space: the information which is rhetorically relevant at the present point of the discourse;
 - b) a shared-entity space: the information which is part of the shared knowledge of the speaker and listener.
 - c) a recoverability space: the information which has entered the discourse context, including the entities which have been mentioned up to this point in the discourse. Information in the recoverability space can be presumed, or pronominalised.

Theme is represented as a role of the speech-act, not using a textual space.