A Data-Driven Methodology for Motivating a Set of Coherence Relations

Alistair Knott

Ph.D.
University of Edinburgh
1996

By *Consequence*, or TRAYNE of Thoughts, I understand that succession of one Thought to another, which is called (to distinguish it from Discourse in words) *Mentall Discourse*. When a man thinketh on any thing whatsoever, His next Thought after, is not altogether so casuall as it seems to be. Not every Thought to every Thought succeeds indifferently...

The generall use of Speech, is to transferre our Mentall Discourse, into Verbal; or the Trayne of our Thoughts into a Trayne of Words...

Thomas Hobbes: Leviathan *(1651)*

It is tempting to speculate that these coherence relations are instantiations in discourse comprehension of more general principles of coherence that we apply in attempting to make sense of the world we find ourselves in, principles that rest ultimately on some notion of cognitive economy.

Jerry Hobbs: On the Coherence and Structure of Discourse *(1985)*

Abstract

The notion that a text is coherent in virtue of the 'relations' that hold between its component spans currently forms the basis for an active research programme in discourse linguistics. **Coherence relations** feature prominently in many theories of discourse structure, and have recently been used with considerable success in text generation systems. However, while the concept of coherence relations is now common currency for discourse theorists, there remains much confusion about them, and no *standard* set of relations has yet emerged.

The aim of this thesis is to contribute towards the development of a standard set of relations. We begin from an explicitly empirical conception of relations: they are taken to model a collection of **psychological mechanisms** operative during the tasks of reading and writing. This conception is fleshed out with reference to psychological theories of skilled task performance, and to Rosch's notion of the **basic level** of categorisation.

A methodology for investigating these mechanisms is then presented, which takes as its starting point a study of **cue phrases**—the sentence/clause connectives by which they are signalled. Although it is conventional to investigate psychological mechanisms by studying human behaviour, it is argued here that evidence for the constructs modelled by relations can be sought in an analysis of the linguistic resources available for marking them explicitly in text.

The methodology is based on two simple linguistic tests: the **test for cue phrases** and the **test for substitutability**. Both tests are functional in inspiration: the former test identifies a heterogenous class of phrases used for linking one portion of text to another; and the latter test is used to discover when a writer is willing to substitute one of these phrases for another. The tests are designed to capture the judgements of ordinary readers and writers, rather than the theoretical intuitions of specialised discourse analysts.

The test for cue phrases is used to analyse around 200 pages of naturally occurring text, from which a corpus of over 200 cue phrases is assembled. The substitutability test is then used to organise this corpus into a hierarchical **taxonomy**, representing the substitutability relationship between every pair of phrases.

The taxonomy of cue phrases lends itself neatly to a model of relations as **feature-based constructs**. Many cue phrases can be interpreted as signalling just some features of relations, rather than whole relations. Small extracts from the taxonomy can be used systematically to determine the alternative values of single features; complex relation definitions can then be formed by combining the values of many features.

The thesis delivers results on two levels. Firstly, it sets out a methodology for motivating a set of relation definitions, which rests on a systematic analysis of concrete linguistic data, and demands a minimum of theoretical assumptions. Also provided are the relation definitions which result from applying the methodology. The new definitions give an interesting picture of the variation that exists amongst cue phrases, and offer a number of innovative insights into text coherence.

Acknowledgements

I would like to take this opportunity to express my deepest thanks to my supervisors, Chris Mellish, Robert Dale and Steve Isard, for their immense help and encouragement over the last four years. Robert first prompted me to think about the theoretical status of coherence relations; he led me to appreciate the need for a firm foundation for these constructs, as well as overseeing the assembly of the taxonomy of cue phrases. Chris was a huge help in guiding me through the theoretical interpretation of the taxonomy, in all its details. Steve provided a valuable perspective on the thesis as a whole.

I would also like to thank my examiners, Donia Scott, Graeme Ritchie and Alex Lascarides, for their detailed comments and suggestions on the manuscript, and for a very interesting discussion during the viva.

Many other people have been influential in the development of the ideas presented here. A big thank you to all of the following, whose feedback and suggestions were much appreciated: Jean Carletta, Ed Hovy, Elisabeth Maier, Kathy McCoy, Johanna Moore, Megan Moser, Jon Oberlander, Cécile Paris, Orin Percus, Ted Sanders, Karen Spärck Jones, Wilbert Spooren, Manfred Stede, and Keith Vander Linden. Special thanks to Ted, for long discussions in several countries, over many glasses of Belgian beer.

It seems to be easier to start a PhD than to finish one. This thesis would certainly not have been finished without the concerted efforts of a large and ruthless team of people, and I would like to say thank you to all of them for managing to force me into submission. Firstly to Chris, for the diverse rhetorical strategies he has employed to encourage me to finish, ranging from rational arguments and encouragements, through thinly veiled warnings, to direct threats and emotive appeals; and finally for eliciting from me a solemn oath to stop my research and start writing. Next, to the Engineering and Physical Sciences Research Council (EPSRC), for their generous support over three years; and also for their Faustian 'four year deadline', which I have watched approaching with increasing horror but which has finally—I think—had its intended effect on me. And lastly, to the assorted PhD students and denizens of the AI department, who have made my life on E floor most pleasant. Thanks are due to Alberto, Chris, Ed, Elena, Ian, Ion, Janet, Jean, Jeremy, Kim, Nam Seog, Steffan, and Wamberto, who have all been a great influence on me during the past few years. Needless to say, any defects that remain in my character are solely my own responsibility.

I would like to give special thanks to my family. Their constant support, particularly during the past year when I was running on empty, has been very important to me. Thank you Mum, Dad, Beck, Har, Jule. Thanks also to my wider family, in Britain, in New Zealand and in Tonga.

Above all I would like to thank Mele, who has had to put up with a lot of ranting and raving on the subject of discourse structure just recently. Thank you for your love, and cheerfulness, and steadfastness. And especially for the notice pinned above my bed, reminding me of the important fact that No-One Reads PhD Theses.

Declaration

I hereby declare that I composed this thesis entirely myself and that it describes my own research.

Alistair Knott
Edinburgh
December 3, 1996

Contents

List of Figures

Chapter 1

Introduction

1.1 Discourse Coherence: The Phenomenon Under Investigation

Theories of discourse coherence aim to investigate the rules which govern how clauses and sentences can be strung together into a text. While syntactic theories concentrate on the internal structure of sentences, theories of discourse look beyond single sentences to the organisation of larger units of language, such as conversations, books or newspaper articles.

The starting point for any such theory is that what we normally think of as 'connected discourse' is more than just a concatenation of random sentences. This fact is uncontentious, and a few examples will suffice to illustrate it. Consider, for instance, how a reader opening a newspaper might react to the two paragraphs given in Figure 1.1. The first is taken from the editorial of an issue of *The Economist*. The second contains exactly the same clauses as the first, but arranged in a different order. Whereas the former paragraph can easily be understood, the latter is at best odd, and at worst completely unintelligible. We can say, pending more precise definitions, that the former text is **coherent**, whilst the latter is **incoherent**.

1.1.1 Two Types of Incoherence

At least two separate problems can be identified in the second paragraph in Figure 1.1. One has to do with deciding about the referents of anaphoric expressions in the text. For instance, the phrase *these middle-aged rich people* cannot be interpreted—it occurs before any such group of people has been introduced. It is simply impossible for the reader to understand or evaluate any propositions in which such unresolved anaphora appear.

A second class of problems in the incoherent text relate to the reader's inability to understand why two portions of text are placed next to one another in the first place. In these cases, there are no problems with finding interpretations for the two portions of text; it is merely surprising that they have been juxtaposed. For instance, consider

The World in 1993

1993 will start with the world in a pessimistic frame of mind. That gloom should soon dispel itself. A clear economic recovery is under way. Though it will be hesitant at first, it will last the longer for being so. If you are sitting in one of the world's blackspots, this prediction will seem hopelessly optimistic. But next year's wealth won't return to yesteryear's winners; these middle-aged rich people need to look over their shoulders to the younger world that is closing in on them.

The World in 1993

1993 will start with the world in a pessimistic frame of mind. A clear economic recovery is under way. That gloom should soon dispel itself. These middle-aged rich people need to look over their shoulders to the younger world that is closing in on them. But next year's wealth won't return to yesteryear's winners; it will last the longer for being so if you are sitting in one of the world's blackspots. Though it will be hesitant at first, this prediction will seem hopelessly optimistic.

Figure 1.1: Coherent and Incoherent Texts

the first two sentences of the incoherent text:

(1.1) 1993 will start with the world in a pessimistic frame of mind. A clear economic recovery is under way.

Why should the world be pessimistic if an economic recovery is under way? The two statements appear to be contradictory, and we do not expect the writer to countenance both of them.

The two examples above appear incoherent in quite different ways. In fact, it is a common assumption amongst discourse linguists that coherence is not a unitary phenomenon at all, but rather the product of a number of different mechanisms. If this is the case, then it should be possible to undertake an explanation of some *aspects* of discourse coherence without providing a complete account of it. In the present study, this will be the objective.

1.1.2 The Role of Context in Decisions about Coherence

It should be noted immediately that when we rule the second paragraph in Figure 1.1 as incoherent, we are not making a judgement about an isolated piece of text; we are also making certain assumptions about its reader, its writer, and the situation in which it was produced. For instance, if the text was accompanied by a photograph of a group of well-heeled businessmen, then a plausible referent for the phrase *these middle-aged rich people* would be available. Again, if we imagine the text to be aimed at readers who believe that the world is full of anarchists who dread an economic recovery, then the sentences in Example 1.1 become intelligible.

Judgements about coherence, therefore, are not made about strings of sentences, but about complexes comprising several additional components, including a reader, a writer, the stock of world knowledge which they share, and what we might call a 'communicative situation' in which they find themselves. This point has often been made in the past: Halliday and Hasan (1976) were influential in distinguishing between a **text** and the **context** in which it appears, envisaging a strong role for the latter concept in an explanation of the former. While a text is relatively concrete, taking the form either of marks on a page or of a series of spoken sounds, its context is a much more intangible notion, and much harder to capture. Indeed, most of the difficulties which arise in simulating the tasks of text production and interpretation have to do with the adequate representation and use of the reader and writer's world knowledge, and of the situation in which a text is produced.

The importance of context in the study of discourse makes the presentation of example texts somewhat of a problem. It is impossible to specify the context for a text under discussion in all its detail; in what follows, the reader will often have to imagine plausible contexts for the texts being discussed. An effort will be made to keep this task as easy as possible.

It is interesting to discover that inventing texts which are incoherent in *any* context is actually quite difficult. This is particularly true as regards the kind of incoherence which results from juxtaposing apparently unrelated portions of text—if we allow ourselves to imagine unusual contexts, we can often give very bizarre discourses coherent interpretations. For instance, consider Text 1.2:

(1.2) Sally decided to take the history course. The ducks on the lake were not eating the bread.[1]

This text, although odd at first sight, can be perfectly well understood if we imagine Sally to have unusual superstitions about the ducks on the lake.

It seems as though readers often go out of their way to find a coherent interpretation of a text, even when this involves making some odd assumptions about it. In a way, this is to be expected: it cannot be easy for a reader to abandon the assumption that the text conforms at some level to Gricean Maxims (Grice (1975)); but this is often what is required in order to appreciate that a text is incoherent. In addition, readers are not often exposed to incoherent texts, and so it is natural for them to operate on the assumption of coherence.

1.1.3 Degrees of Coherence and Incoherence

A final observation about coherence is that it can be present in texts in varying degrees. Texts should not be thought of as either coherent or incoherent, rather as more or less coherent; the dividing line between coherent and incoherent texts is a hazy one.

Part of this haziness can be traced to the problem mentioned in the previous section, of the role of context in decisions about coherence, and of the tendency of readers to look for plausible contexts even for the most implausible texts. However, even if contexts are fully specified in advance, texts still admit of degrees of coherence.

[1] Linguistic examples are hand-crafted unless otherwise noted. One of the reasons for this is so that examples of incoherent texts can be provided as easily as examples of coherent ones.

There are two reasons for this. The first is that coherence is a *global* measure of text. Thus it is possible that a text is coherent in some places, but incoherent in others. It would of course be preferable to use a more local measure of coherence to avoid this problem. However, we are still working with a pre-theoretical concept at present, and without a theory of coherence, it seems premature to try tracking down problems to specific features of a text.

A second reason for the haziness of the concept of coherence is simply that we can understand some texts more easily than others. If a text is well written, the intended reader should understand it without much difficulty; however, even if it is badly written, a reader may still be able to piece together its meaning. Consider the following example, from 'The Brief English Handbook' (Dornan and Dawe (1984)):

(1.3) Alice Adams is a successful writer, and she lives in San Fransisco, and she has received grants from the Guggenheim foundation and the National Foundation for the Arts.

This text is stylistically awkward, but it still manages to get its message across. Ideally, however, we would like a way of distinguishing between such texts and 'well-written' ones.

1.1.4 A Definition of Coherence

For the purposes of this thesis, an operational definition of coherence is required; preferably, one which does not pre-empt too many theoretical questions. In what follows, therefore, we will think of an incoherent text as one whose 'structure' a high school teacher would be inclined to question or correct if it appeared as part of a student's essay. We will not be concerned with any corrections a teacher might make arising from bad grammar or spelling mistakes, or from errors of fact: an incoherent text is one which avoids such errors, but is nonetheless hard or impossible to understand.

Admittedly, this is a very informal definition to start off with. But an initial description of 'what a theory is about' precisely *should* be informal, so as not to begin doing the work intended for the theory itself. Subsequent more precise conceptions of coherence should ultimately be judged according to how well they reflect an informal conception such as this one.

1.2 First and Second Order Tasks for Text Analysts

Theories of coherence call for texts to be analysed—that is, broken up into parts which are given descriptions in theoretical terms. The theoretical description of an individual text should contribute to an account of why it is or is not coherent.

Two quite separate questions confront a discourse analyst. Firstly, what are the appropriate theoretical terms to be used for describing texts, and how are they to be used? Secondly, what is the right analysis for any particular text—for instance the text on the left in Figure 1.1? The first question calls for a decision amongst competing theories of coherence. The second calls for a decision about how to represent a particular text in terms of one theory or another. We can refer to this latter task as the **first-order** task, and to the task of deciding between theories as the **second-order** task. Both tasks have their problems. Theorists can (and do) disagree about the primitives to be used by a theory of coherence; and also, once a theory has been agreed on, over individual analyses.

This thesis is directed primarily at the second-order task—that of deciding on the aspects of a text which an analysis should identify, rather than of determining *how* these aspects are to be identified for any given text, either by a human or by a machine. The second-order task seems logically primary; after all, it is not possible to analyse a text until a particular representation formalism has been chosen. The first-order task will also be addressed to some extent, however, as is detailed in Section 1.3.2.

1.3 Requirements for a Theory of Discourse Coherence

While the previous sections described *what* is to be investigated in this thesis, the present section examines *how* it is to be investigated. Two requirements for a theory of coherence are set out, one concerning the goals that the theory should have, and one concerning the kinds of intuitions which should be admissible as evidence.

1.3.1 Descriptive versus Explanatory Adequacy

To assess a theory of discourse coherence, we can make use of some Chomskyan terminology. One of the tasks of the theory is clearly to provide a means for distinguishing between coherent and incoherent texts. It should be possible to use the theory to decide whether any given text is coherent or incoherent; if the theory is good, then these decisions will correspond to our own judgements of coherence. A theory which satisfies this requirement can be termed **descriptively adequate**.

A second, more stringent requirement for a theory of coherence is that of **explanatory adequacy**. Chomsky's (1964) original formulation of this notion is with reference to syntactic theories. An explanatorily adequate theory is one which is able not only to distinguish between well-formed and ill-formed texts, but also to explain why it is that well-formed texts are well-formed; in other words, to explain why it is that we have the grammars that we do. The notion can apply just as easily to theories of coherence—in this case, it would demand that a theory is able not only to distinguish between coherent and incoherent texts, but also to explain why it is that readers and writers have these standards for coherence.

It is ambitious even to aim for descriptive adequacy in a theory of discourse. However, this is no reason to completely ignore the criterion of explanatory adequacy. In this thesis, particularly in Chapter 3, the notion of explanatory adequacy will figure quite prominently.

1.3.2 Pre-Theoretical and Post-Theoretical Intuitions

The raw data for a theory of discourse are 'intuitions of coherence'. As Section 1.1.4 specifies, we are examining the kind of intuitions that a high school teacher might have about a range of different texts—if corrections of a certain kind are felt to be needed, a text is classed as incoherent, otherwise, it is classed as coherent.

This type of judgement about a text is what a theory of coherence must explain; it can be termed **pre-theoretical**. It should be contrasted with another type of judgement about text, which we can call **post-theoretical**—such judgements are about 'the way a text should be described'. A post-theoretical judgement might be, for instance, that 'a discourse segment boundary falls between sentences s_1 and s_2', or that 'the focused entity in span S_5 is E_3'. The first-order task of text analysis is normally thought of as involving such judgements.

As readers and writers of text, we might well have our own ideas about what makes a text hang together well, and these might even involve notions of text segments, focus and so on. If asked what the topic of a given sentence is, or how a given discourse should be segmented, we would probably be able to make some suggestions. However, these intuitions should be differentiated sharply from our intuitions of coherence. They cannot be relied upon in the same way: while intuitions of coherence have the status of irrefutable facts to be explained, post-theoretical intuitions are intended as *part* of an explanation of those facts, and are subject to the same standards of assessment as any other explanation.

In fact, our intuitions about text analysis do not seem to be especially reliable. As already noted, disagreements between analysts are very common. And although some measure of consensus among analysts is often claimed for a particular theory, differences between analysts from different theoretical backgrounds are often quite significant. The problem is that it is hard to see how such differences can be resolved. Why should one theorist's intuitions be any more reliable than another's? The theorists are competing for an explanation of the same text; if we are looking for a single theory of text, then there is no question of both theorists' intuitions being equally good.

A second, more fundamental problem with a reliance on post-theoretical intuitions can also be noted. There is evidence in the psychological literature that people's intuitions about how they perform complex and highly learned tasks differ considerably from the way they are actually performed. For instance, Berry and Broadbent (1984) have found that subjects' verbalisations about how they perform a task do not change markedly as their skill increases, although qualitative changes in performance can be observed. The tasks of text generation and interpretation are highly skilled, and thus it could be that our intuitions about them (however much we agree or disagree) are inaccurate.

How then are we to go about analysing texts? In this thesis, a method will be suggested that relies less heavily on post-theoretical intuitions.[2] It is proposed that some of the constructs in a theory of discourse coherence can be linked a priori to a *second* class of pre-theoretical data, independent from straightforward judgments of coherence. The

[2] Of course, the development of the theory still requires post-theoretical intuitions. The problem of using *only* pre-theoretical intuitions is highlighted by Scott and Paris (1995).

new class of data concerns the judgements of ordinary writers about when two phrases of a certain type can be substituted for one another in a text, and will be described in detail in the following chapters. The important thing to note now is that the judgements are ones which a normal writer is frequently called upon to make; they do not require the analyst to think explicitly about theoretical concepts like 'focus' or *text segments* and so diverge from tasks present in the normal writing process. If the theory of coherence is successful, then, it will identify a correlation between the two classes of pre-theoretical intuitions.

This approach to theorising about discourse is much closer to the traditional pattern of empirical theories; it consists in making predictions about the relationship between two independent sets of data to be explained. For another example of this approach, we can again refer to theories of syntax. In a syntactic theory, the original pre-theoretical intuitions to be accounted for are judgements of well-formedness. These are explained in terms of a theory whose central construct is that of a 'constituent'. This construct is in turn grounded in *other* pre-theoretical intuitions, enshrined in the so-called 'tests for constituency'. For instance, it is specified that constituents can typically be moved from one part of a sentence to another, or replaced by an anaphoric expression, without affecting well-formedness. The crucial point is that our intuitions about these manipulations are different from our intuitions of well-formedness themselves. And moreover, they cannot be questioned; they are just another phenomenon to be explained. Essentially, what we have is a theory which makes predictions about a relationship between two independent classes of data. If these predictions are borne out, we have an empirical result.

Of course, the development of the theory itself still requires post-theoretical intuitions. Definitions for theoretical constructs must be proposed, and we must justify the decision to link these constructs to pre-theoretical judgements.

1.4 An Outline of the Thesis

We began this chapter by introducing and expanding on the notion of discourse coherence. We then distinguished the task of creating a theory of discourse coherence from the task of analysing particular texts in terms of such a theory—only the former task will be attempted in this thesis. Finally, we set out the requirements for a theory of coherence; namely that it seeks for an explanatory account of the phenomenon, and that it draws only on the pre-theoretical intuitions of ordinary readers and writers.

A summary of the rest of the thesis will now be given.

Chapter 2: Coherence Relations

The thesis targets a particular class of discourse theories; those which attribute the coherence of a text (at least in part) to the **relations** which hold, at different levels of hierarchy, between its various sub-parts. Many such theories can be found in the literature; in the last few years, the notion of coherence relations has become increasingly popular as they have been successfully adapted for use in computational applications such as text generation systems.

This chapter reviews the many different theories of coherence relations, emphasising the differences between them, and their lack of agreement on a standard set of relations. The proliferation of alternative sets of relations is the central problem to be addressed in the thesis. What is needed is a clear conception of the role of relations, and based on this, a method for justifying one particular set of relations over the others.

Chapter 3: Motivating a Set of Relations

In this chapter two existing approaches for justifying relations are described and compared. One approach associates relations with linguistic conjunctions or **cue phrases**. This permits a subtle classification of relations, but seems to lack explanatory adequacy; it is not clear why relations thus defined should provide a particularly revealing account of text. A second approach is to think of relations as modelling a set of psychological constructs used by readers and writers when they process text. This promises a more explanatory account; but there are problems with the experimental methodologies designed to investigate these constructs, because we have no reliable experimental window on 'the relations people use'.

The main proposal in the thesis is that these two methods can be combined, in such a way as to maximise the advantages of both, while avoiding their main problems. It is argued that cue phrases can be taken as evidence for relations precisely if they are thought of as modelling psychological constructs.

Chapter 4: A Data-Driven Methodology for Determining a Set of Relations

On the basis of the argument in Chapter 3, a step-by-step methodology is proposed for motivating a set of relations, making use of the pre-theoretical intuitions of readers and writers. The first step is to gather a corpus of cue phrases. A pre-theoretical **test for cue phrases** is described, which is used to gather a corpus of some 200 phrases from several hundred pages of naturally occurring text. These phrases are then organised into a taxonomy, using a second pre-theoretical **test for substitutability**, which taps writer's intuitions about whether one phrase can replace another in a given context. The corpus of cue phrases is given in Appendix A, and the taxonomy of cue phrases in Appendix B.

Chapter 5: Preliminaries for Defining a Set of Relations

This chapter outlines how the taxonomy of cue phrases can be used systematically to motivate a set of relation definitions. It is argued that the taxonomy lends itself very naturally to a conception of relations as **feature-based constructs**. Some general criteria for the individual features to be motivated are then discussed.

Chapter 6: Using the Taxonomy to Create Relation Definitions

In this chapter, a preliminary set of relation definitions is put forward. To begin with, the taxonomy of cue phrases is used systematically to motivate a number of orthogonal features: each feature is justified individually on the basis of appropriate extracts from the taxonomy. After this, the issue of how these features can be assembled to create complex relation definitions is discussed. The complete set of feature definitions is given in Appendix C, and a preliminary set of relation definitions is given in Appendix D.

The task of motivating features from the taxonomy is an arduous one, and there is plenty of room for improvement and addition to the sets of features and relation definitions reached in this chapter. However, a number of interesting new ideas do emerge as the theoretical interpretation of the taxonomy is developed. The fact that substitution methodology yields these fruitful results is an additional argument in its favour.

Chapter 7: An Evaluation of the Substitution Methodology

In this chapter, a number of problems with the substitution methodology are taken up. These have to do partly with the operation of the substitutability test itself, and partly with the set of relations which it eventually sanctions. None of these problems appears fatal to the proposed methodology, although they all point to further interesting avenues of research.

Chapter 8: Conclusions

The objectives of the thesis are quite straightforward. It attempts to remedy a current problem for theories of discourse, namely the confusing proliferation of coherence relations. It does so by giving firstly a concrete proposal about what relations should be thought of as modelling; secondly a clear methodology for investigating these entities; and thirdly, as a result of applying this methodology, an embryonic set of relation definitions.

The contribution of the thesis is twofold. The most tangible contribution is the set of relation definitions itself—these have several features not found in other sets of relations, and promise to be useful both in text analysis and in computational applications. However, it really requires more justification than this for bringing yet another set of relations into the world. The primary aim of the thesis is to establish and argue for a systematic methodology for *determining* a set of relations, based on a battery of fairly replicable linguistic tests. While it is unrealistic to suppose that every researcher who uses it will emerge with the same set of relations, it would at least be preferable if the differences between them could be traced to disagreements over the interpretation of concrete linguistic data, rather than being expressed in terms of intangible first principles.

1.5 The Scope of the Project, and Some Terminology

Before moving on, a word is in order about the type of texts which are within the scope of this study. I will only be looking at *monologue* in this thesis; and I shall concentrate (but not exclusively) on written monologue. Texts which incorporate devices from dialogue such as direct speech or question-and-answer patterns ('How do we know this? By looking at the evidence') will not be considered. I will also concentrate on *English* texts—although some cross-linguistic comparisons will be made in Chapters 3 and 7.

Finally, some terminology. I will use the words **text** and **discourse** interchangeably. Unlike Halliday and Hasan, for whom 'a text' is by definition coherent, I shall think of a text or discourse as any sequence of sentences produced by one writer in a particular context—texts can thus be ruled as coherent or incoherent.

I will refer throughout to the **readers** and **writers** of a text; these terms are intended cover hearers and speakers too, where applicable. To circumvent the biases of English, writers will be referred to generically as 'she', and readers as 'he'.