

Semantic and pragmatic relations and their intended effects

Alistair Knott
Department of Computer Science
University of Otago, New Zealand

March 15, 1999

Abstract

According to many current theories of coherence relations, a distinction can be drawn between relations that hold between the content of the text spans they link (which can be termed SEMANTIC relations), and those that hold between the utterances of the text spans themselves or the beliefs which underlie them (which can be termed PRAGMATIC). However, this distinction is not always clearcut. One useful recent proposal by Sweetser effectively results in dividing the class of PRAGMATIC relations in two, distinguishing between EPISTEMIC relations (which hold between the beliefs which underlie a speaker's utterances) and SPEECH-ACT relations (which hold between the utterances themselves). In this paper I begin by reviewing the advantages and shortcomings of both these approaches. I then propose an alternative, intention-based definition of PRAGMATIC relations, which I argue retains the advantages of both approaches while avoiding at least some of their shortcomings.

1 Introduction

This paper takes as its starting point an assumption that has received much attention in both computational and psycholinguistic treatments of discourse: that the coherence of an extended text can be explained in terms of an account of the **relations** which hold between its component spans. Naturally, when it is stated as baldly as this, the assumption is almost vacuously true. However, it is useful in providing the beginnings of a vocabulary for discussing the issue of text coherence. It means that the empirical question of what makes a text coherent can be re-expressed as two more specific questions. Firstly, how can we define the set of relations which are permitted within coherent text? And secondly, what are the structural constraints on the configuration of these relations within coherent text? In the present paper, I will concentrate on the first of these questions.

Accounts of the semantics of coherence relations have often suggested that they should be thought of as composite entities, defined in terms of a number of different dimensions. This assumption is implicit in systemic approaches to conjunctive relations (e.g. Halliday and Hasan, 1976, Martin, 1992), in psycholinguistic/text-linguistic approaches (e.g. van Dijk, 1979, Redeker, 1990, Sanders, Spooren and Noordman, 1992), and in computational/semantic approaches (e.g. Hobbs, 1985, Elhadad and McKeown, 1990, Knott and Mellish, 1996, Oversteegen, 1997). In all of these approaches, specifying the semantics of a coherence relation involves specifying a number of different values, which position it within a multidimensional space. Furthermore, if we assume that intersentential/interclausal conjunctions can be used

to signal coherence relations in surface text, then the semantics of these conjunctions can be specified in a similar fashion. There is a large measure of variation in the above studies as to the closeness of the mapping between the set of coherence relations and the set of connectives. Likewise, the studies differ as regards the number of dimensions they propose, and the degree of interdependence they envisage between relations. However, the principle is the same in each case.

One of the fundamental dimensions to have been proposed distinguishes between relations that hold between the content of the text spans they link and those that hold between the utterances of the text spans themselves, or the beliefs which underlie them. This distinction is most clearly illustrated with reference to two kinds of ‘causal’ relations: those that simply describe a cause and effect occurring in the world (such as Example 1 below), and those that have an argumentative or rhetorical force (such as Example 2).

- (1) Bill was starving, so he had a sandwich.
- (2) Bill had five sandwiches, so he was/must have been starving.

In each case, *so* is taken to link a cause and effect. But while in Example 1 the cause and effect are taken to be the eventualities *described* by the respective clauses, in Example 2, the speaker’s *belief* in the eventuality described in the first clause is taken to cause her *conclusion* about the eventuality described in the second. Nearly all the theories which advocate a decomposition of relations include a dimension which reflects this difference. The main reason for this is the apparent productivity of the distinction, in particular in capturing alternative senses of sentence and clause connectives. There are many connectives, of several different types, which can be analysed as ambiguous with respect to this dimension. As well as causal/inferential connectives like *so* and *because*, there are conditional connectives such as *if*, contrastive connectives such as *but*, disjunctive connectives such as *or* and *otherwise*, and temporal connectives such as *then*. In each case, a similar ambiguity seems to be discernable; the challenge is to find a formulation of the distinction which covers all the cases.

Several different terms have been used for the two types of relation, and slightly different definitions have been given. Halliday and Hasan and Martin refer to EXTERNAL and INTERNAL relations; van Dijk refers to SEMANTIC and PRAGMATIC relations; and Redeker refers to IDEATIONAL and PRAGMATIC relations. I will adopt van Dijk’s formulation in the remainder of this paper.

In this paper, I begin in Section 2 by outlining the two types of relation in more detail. In Section 3 I then consider some problems with this account noted by Sweetser (1990)¹, which led to her proposing a new tripartite distinction in place of the bipartite one: in the new classification, the class of PRAGMATIC is divided in two, into the class of EPISTEMIC relations, which hold at the level of premises and conclusions about what is the case in the world, and the class of SPEECH-ACT relations, which hold between the utterances themselves. I also present Sweetser’s extension of the set of relations from causal to conditional, disjunctive and sequential relations. In Section 4 I consider the advantages and disadvantages of the proposals previously outlined, and also note some problems which apply equally to each proposal. In Section 5 I then propose an alternative, intention-based definition of PRAGMATIC relations which overcomes a number of these problems. An assessment of the new definition, noting

¹Sweetser in fact addresses two early versions of the bipartite distinction formulated by Ross (1967) and Davison (1973). However, her criticisms apply to the more recent versions too.

advantages relating to the scope of its applicability and to its explanatory force, as well as cases which it cannot handle, is given in Section 6. I conclude in Section 7 by considering some possible solutions to these problems, by introducing a new bipartite distinction, orthogonal to the SEMANTIC/PRAGMATIC distinction, between CAUSE-DRIVEN and RESULT-DRIVEN relations.

2 The bipartite SEMANTIC/PRAGMATIC distinction

Consider again the examples mentioned above.

- (1) Bill was starving, so he had a sandwich.
- (2) Bill had five sandwiches, so he was/must have been starving.

The distinction between SEMANTIC relations such as Example 1 and PRAGMATIC relations such as Example 2 has been captured by many researchers by interpreting the latter type of relation as containing an **implicit performative**. In the above example, the performative could be made explicit as follows:

- (3) Bill had five sandwiches, so *I conclude* he was starving.

Generalising from this example, we can state that SEMANTIC relations hold directly between the propositional content of the two related utterances, while PRAGMATIC relations hold between the utterances themselves, interpreted as speech acts. The clearest expression of this idea comes from van Dijk (1979):

Pragmatic connectives express relations between speech acts, whereas semantic connectives express relations between denoted facts. (van Dijk, 1979:449)

Many other theorists have advocated a similar distinction—although typically qualified in some way—including Halliday and Hasan (1976), Martin (1992), Redeker (1990), and Sanders *et al* (1992). (Different qualifications will be mentioned in Sections 3 and 4.2.)

Interpreting PRAGMATIC relations as holding between speech acts allows us to give a connective like *so* a single denotation, which applies equally to PRAGMATIC contexts and to SEMANTIC ones. Rather than assuming a lexical ambiguity in *so* (as signalling either a cause in the world, or an argumentative relation), we can assume that the relation signalled is simple causality in both cases, with the difference between the cases stemming from a systematic variation in which *propositions* are being linked by this relation. The effect of this analysis is essentially to reduce relations at the PRAGMATIC level to relations at the SEMANTIC level. SEMANTIC relations are ones in which information about a relation between two propositions in the world is conveyed from the speaker to the hearer. PRAGMATIC relations, according to the definition just proposed, are exactly the same; it is just that the information which the speaker is communicating to the hearer is (partly) about her own speech acts.

A point to note about the definition of PRAGMATIC relations is that it extends in a useful way to relations involving sentences which are not in the indicative mood. Consider the following two sentences, for instance:

- (4) Bill's starving. So why isn't he eating?
- (5) Bill's starving, so give him something to eat.

The point is that some transformation on the latter sentence in each example is necessary to render it directly comparable to sentences in the indicative mood. Moreover, we want the transformation to reflect the real differences between the different kinds of sentences. Moving to the level of the speech act which underlies the sentence is a way of doing this. The speech act behind a sentence can always be expressed as an indicative sentence, which gives us the possibility of making direct comparisons. Thus we can gloss Example 4 and Example 5 respectively as follows:

- (6) Bill's starving. So *I ask* why isn't he eating?
 (7) Bill's starving, so *I instruct you to* give him something to eat.

In summary, the distinction between SEMANTIC and PRAGMATIC relations can be seen as bearing on two separate problems: the problem of identifying the ambiguity of certain connectives, and the problem of finding a uniform representation for sentences in different moods. It is fair to say that it is the former problem which originally motivated the distinction, and that it is an advantage of the distinction that it goes some way towards addressing the latter problem. However, as we will see, it is far from presenting a complete solution to the latter problem. In fact, this should already be obvious from the fact that PRAGMATIC relations are to be analysed simply as providing information about the world; a complete treatment of imperatives and interrogatives is bound to require something more than this.

3 The tripartite CONTENT/EPISTEMIC/SPEECH-ACT distinction

Sweetser (1990) notes an important underspecification in a gloss such as that given in Example 3. Is the implicit act of 'concluding' to be thought of as a speech act? There are good reasons for maintaining that this is inappropriate. For Sweetser,

[t]here is a class of causal-conjunction uses in which the causality is that between premise and conclusion in the speaker's mind... , and there is another class of uses in which the causality actually involves the speech act itself.

Sweetser draws a sharp distinction between the act of drawing a conclusion and the act of stating it. A speaker is only obliged to *state* a conclusion she has reached 'inasmuch as the rules of conversation make it incumbent on us to say things we believe to be true'; it is therefore inappropriate to gloss an argumentative relation such as that in Example 2 in terms of speech acts. Essentially, what the speaker decides to *say* and what the speaker decides to *believe* are two quite different things, and it is not possible to accurately express one in terms of the other. Note that we can still analyse the text in Example 2 as containing an implicit performative—but the performative verb should be clearly understood as describing a theorem-proving act on the part of the speaker rather than a linguistic one, thereby nailing an important ambiguity in a word like *conclude*, which can be understood in both senses.²

²In fact, Sweetser is unwilling to commit to the view that explicit performatives should be understood as contributing to the semantics of a sentence (Sweetser, 1990:92). But the rationale for this is mainly concern for a proper demarcation the domains of 'semantics' and 'pragmatics'; the only relevant implication for our purposes is that the performative glosses are to be understood as falling in the latter domain rather than the former.

Sweetser proposes a division of the class of PRAGMATIC relations into two categories: SPEECH-ACT relations, which hold between the utterances themselves, and EPISTEMIC relations, which hold at the level of premises and conclusions about what is the case in the world. Her overall classification is tripartite; there is a third category of relations, termed CONTENT relations, which effectively cover the same ground as those termed SEMANTIC in the bipartite classification.

Note that the distinction between EPISTEMIC and SPEECH-ACT relations is to some extent independent of the question of whether a relation contains any non-indicative sentences. There can be SPEECH-ACT relations between two indicative sentences as well as EPISTEMIC ones. Example 8 is a case in point:

(8) The answer's on page 200, since you'll never find it for yourself.

Here we would certainly not want to suggest that the speaker is reaching a conclusion about where the answer is, on the basis of knowing that the hearer won't find it for himself. Glossing with an implicit speech-act performative (*I tell you that* the answer's on page 200) is much more appropriate.

However, there is still some degree of dependence on mood in Sweetser's definitions. As she says (p78),

if an utterance is imperative or interrogative in form, then it cannot reasonably be causally conjoined to another utterance except at the speech-act level.

In practice, it seems likely that many advocates of the bipartite position are aware of the variations which Sweetser points out within the class of PRAGMATIC relations. Definitions of PRAGMATIC relations tend in fact to hedge the issue to some extent, frequently containing disjunctive elements or appealing to somewhat ill-defined higher-level constructs. For instance, Sanders *et al* (1997:126) define PRAGMATIC relations as applying between the content of one span and the speaker's 'claim/advice/conclusion' about the content of the other, while for Redeker (1990:369) they hold between the 'beliefs and intentions' which underlie the two spans. Others appeal to somewhat ill-defined higher-level constructs; for instance, for Halliday and Hasan (1976: 240) the relationship in question is

not so much a relationship between speech acts (though it may take this form...) as a relationship between different stages in the unfolding of a speaker's COMMUNICATION ROLE (...) his choice of speech role and rhetorical channel, his attitudes, his judgements and the like.

Clearly, theorists are looking for a way of generalising across the various different types of PRAGMATIC relations. Sweetser, by dividing the class in two, has to some extent abandoned this goal, but by the same token her definitions are much more intelligible.

3.1 SPEECH-ACT and EPISTEMIC disjunctions

As mentioned in Section 1, one of the attractions of the SEMANTIC/PRAGMATIC distinction is its productivity: it seems to find application not only for causal relations, but also for other types of relations. A strong recommendation for Sweetser's finer-grained distinction is that it too seems productive across this range of relations. Some of the interesting distinctions between EPISTEMIC and SPEECH-ACT relations are outlined in the following three sections; I

will omit discussion of the corresponding CONTENT relations, which should in each case be quite clear.

Examples of EPISTEMIC and SPEECH-ACT disjunctive relations are given in Examples 9 and 10 respectively.

- (9) John is home, or somebody is picking up his newspapers.
- (10) Would you like to come round tonight? Or is your car still in the shop?

Sweetser analyses Example 9 as conveying that the alternative propositions presented are the only two possible conclusions that I can reach. On the other hand, she analyses Example 10 as containing a pair of alternative speech-acts: the speaker is asking to be understood by the hearer as performing either one or the other. The distinction between the two analyses seems useful here; it is inappropriate to gloss Example 9 in terms of speech acts, for the same reason as it is inappropriate to gloss Example 2 in such terms, as discussed at the beginning of Section 3.

3.2 SPEECH-ACT and EPISTEMIC conditionals

Examples of EPISTEMIC and SPEECH-ACT conditional relations are given in Examples 11 and 12 respectively.

- (11) If John went to that party, he was trying to infuriate Miriam.
- (12) How old are you, if it's not a cheeky question?

Example 11 is to be analysed as expressing an implication relation between the speaker's beliefs: the speaker is really informing the hearer that 'if I [*the speaker*] believe that John went to the party, I believe that he was trying to infuriate Miriam'. On the other hand, Example 12 is to be understood as the conditional performance of a speech act: the speaker only wants to ask about the hearer's age if it's not a cheeky question. Again, for the same reasons, it is not appropriate to gloss Example 11 in terms of speech acts, and thus there seems to be good reason for giving different analyses for the two cases.

3.3 SPEECH-ACT and EPISTEMIC temporal sequences

Finally, examples of EPISTEMIC and SPEECH-ACT temporal relations are given in Examples 13 and 14 respectively.

- (13) A: Why don't you want me to take basketweaving this Summer?
B: Well, Mary took basketweaving, and she joined a religious cult.
- (14) Go to bed now! And no more backtalk!

In each of these examples, according to Sweetser, *and* is to be interpreted sequentially. However, while the temporal sequence in Example 14 relates to the order in which the speech acts are performed, the sequence in Example 13 relates to the order of events in the epistemic world. The idea in the latter case is that the two propositions are both to be interpreted as premises in an argument that A should not take basketweaving, but that their ordering is significant. While this latter analysis is not completely clear to me, the fact that explicitly sequential conjunctions like *to begin with* and *next* can be used to link multiple premises

in an argument has often been noted, and provides an equally good rationale for temporal EPISTEMIC relations. For instance, to echo examples given by Halliday and Hasan:

- (15) John's unsuitable for the job. To begin with, he's too young. Next, he's too hot-headed. Finally...

Sweetser's objections against speech-act glosses of argumentative relations are as telling in this case as in the others; an EPISTEMIC interpretation of the sequential relation is again preferable.

4 Some problems with the existing distinctions

While the bipartite and tripartite classifications are both extremely useful in distinguishing between different types of coherence relation and uses of connective phrases, there remain some outstanding problems in each case. In this section, I will review some of these.

4.1 Generalisations across SPEECH-ACT and EPISTEMIC relations

One problem with the tripartite account is that EPISTEMIC and SPEECH-ACT relations do seem to share a lot in common. For instance, as Sweetser points out, many connectives (such as French *puisque*) are appropriate for both EPISTEMIC and SPEECH-ACT relations, but not for CONTENT relations. It is useful to be able to capture this by positing a single category of relations into which they both fall. Indeed, this is an important reason for thinking of PRAGMATIC relations as a single class.

On the other hand, as we have seen, the bipartite distinction hardly makes matters clearer. To my mind, those definitions of PRAGMATIC relations which are truly general (in the sense that they propose a single general characterisation of all relations in this class) are open to the objections which Sweetser has levelled at them, while those definitions that allow for the distinction which Sweetser notes are essentially disjunctions, and hence tripartite in every respect except perhaps a terminological one. And so the problem of defining the commonalities between SPEECH-ACT and EPISTEMIC relations is still outstanding for both types of account.

4.2 Problems in defining the level at which relations hold

A related problem with both proposed classifications concerns another type of disjunction which is present, more or less implicitly, in the definition of PRAGMATIC/EPISTEMIC/SPEECH-ACT relations. In each case, the general idea is to suggest that a basic relation, such as cause, disjunction or temporal sequence, applies between the two related spans, and that the different classes of relation are the result of this basic relation applying at different levels: between the propositions which are expressed by the spans, or between the speaker's beliefs in these propositions, or between the speaker's utterance of the propositions. To make this more concrete, we can imagine three functions that operate on an utterance U : one function $c(U)$ which returns its propositional content, one function $b(U)$ which returns the speaker's belief in its content (or some event relating to the genesis of this belief), and one function $u(U)$ which returns the utterance itself (i.e. the identity function). We could then state, generally, that the coherence relation between two utterances U_1 and U_2 is to be analysed as

$$R(f(U_1), f(U_2))^3$$

³Or perhaps more precisely, as stating that the proposition $R(f(U_1), f(U_2))$ is asserted by the speaker.

where R denotes a ‘basic’ relation (cause, sequence or whatever), and f denotes one of the above functions. The definitions of the three categories of relation would then be straightforward:

- for SEMANTIC (CONTENT) relations, $f = c$;
- for EPISTEMIC relations, $f = b$;
- for SPEECH-ACT relations, $f = u$;

However, while these definitions would be nice, they don’t always result in the analyses which have been suggested as suitable for the examples we have seen thus far. In fact, getting the analysis right frequently depends on us applying *different* functions to the two utterances in the relation. For instance, consider the following SEMANTIC, EPISTEMIC and SPEECH-ACT relations:

(16) Bill was starving, so he ate loads of food.

(17) Bill’s starving, so he’ll want loads of food.

(18) Bill’s starving. So why isn’t he eating?

In each case, the basic relation can be taken to be cause. In Examples 16 and 17, the functions seem to work out: the relation applies between the propositional contents of the two spans in the former case, and between the beliefs in these propositions in the latter case. But in Example 18 we would not want to say that the relation applies between the two utterances. The cause of my asking the question *Why isn’t Bill eating?* is simply the fact that Bill is starving, or possibly my belief in this fact; it is certainly nothing to do with my utterance of it.

In fact, the case of SPEECH-ACT relations is even more problematic. While some SPEECH-ACT relations only involve one of the utterances, others involve both. The following relation is a case in point:

(19) Would you like to come round tonight? Or are you busy?

In this case, as we have seen, we must hold that the basic relation (here, disjunction) applies between the two interrogative speech acts to get the desired analysis.

To sum up: in the accounts we have seen so far, it does not seem possible to state in a general way ‘which variable changes’ when we move from one level to another. Again, we have to rely on disjunction to get things right. For instance, Sanders *et al* define the class of PRAGMATIC relations as those where the segments are related ‘because of the illocutionary meaning of one or both of the segments’. What we really want is a formula with a free variable in it, and different value for this variable for CONTENT, EPISTEMIC and SPEECH-ACT relations; but it is far from clear how this can be achieved.

4.3 Problems with imperative sequences

A general problem we have already mentioned is that the classifications of relations considered so far have a certain amount to say about relations involving non-indicative sentences, but seem unlikely as the basis of a complete account, either of these sentences or of the relations between them. To take a specific example of this problem, consider the case of a temporal sequence of imperatives. For instance:

(20) Peel the onions. Then chop them.

Interestingly, there are no examples in Sweetser of *and* or *then* being used to relate two imperatives in a temporal sequence. This is because the account of SPEECH-ACT relations is clearly not suitable in such cases. We certainly don't want to interpret the temporal sequence here as holding between the speaker's utterances: the important thing, of course, is that the actions themselves need to be performed in the right sequence. The problem, then, is to come up with an account which does justice to cases such as Example 20 while still giving the kind of analysis of sequences of premises or speech acts that we have already seen in Section 3.3.

It should be quite clear that any complete account of non-indicative sentences will involve a departure from the current definitions of PRAGMATIC/SPEECH-ACT relations, in the sense that these definitions are all expressed in terms of the *information* which relations convey to a hearer, over and above any information which is contained in the individual utterances. It would obviously be wrong to analyse a simple imperative proposition in this way. There is clearly more to the imperative *do X* than the information *I [the speaker] tell you to do X*; we have to make reference in the definition to the conditions under which the imperative has the desired effect on the hearer. The same seems uncontroversially true for the sequence of imperatives given above: the relation of sequence has to be seen as part of the desired effect of the utterance. What is needed is a way of unifying this analysis with the analyses we have seen so far of imperatives as they feature in other types of relation.

4.4 Problems of explanatory adequacy

The problem just mentioned for SPEECH-ACT relations has a parallel in the definitions proposed so far for EPISTEMIC or argumentative relations. Again, this stems from the fact that the content of these relations is essentially taken to be informative. Consider the following example:

(21) Bill's starving, so he'll want loads of food.

According to the analyses seen so far, the relation in this text basically provides the hearer with information about the causes of one of her beliefs. This seems plausible enough in some contexts. However, a common—even prototypical—function of argumentative relations is to achieve a rhetorical effect on their hearers, causing them to believe a proposition in circumstances where simply stating it would not. (Indeed, the category of PRAGMATIC relations has long been associated with the effect of increasing some desire or positive regard in the hearer; see for instance Mann and Thompson, 1988; Bateman and Rondhuis, 1997.) Analysing an argumentative relation as providing information about why the *speaker* believes something does not by itself explain how it can have this effect, at least not directly. It requires in addition the stipulation that when the hearer learns about the causal processes amongst the speaker's own beliefs, this will prompt similar processes amongst his own beliefs. This stipulation is not part of the relation definitions we have seen so far. Moreover, the idea that this kind of analogy is responsible for the effect of an argumentative relation is not particularly plausible: an argument which is compelling for one person is not always compelling for someone else.

4.5 Problems with 'conditional speech acts'

A final problem for the definitions we have seen so far relates to their application in conditional or disjunctive contexts. The problem is specific to the class of SPEECH-ACT relations.

Consider, for instance, Sweetser's analysis of the following texts:

(22) Would you like to come round tonight? Or is your car still in the shop?

(23) How old are you, if it's not a cheeky question?

As we have seen, Example 22 is analysed as a case of SPEECH-ACT disjunction: the disjunction is between whether the speaker is taken to perform the interrogative speech act associated with the first utterance, or that associated with the second. Example 23 is a SPEECH-ACT conditional; that is to say, the question *how old are you?* is only to be understood as being asked if it's not a cheeky one.

Both of these analyses rely on the notion of conditional speech acts. But this notion is not unproblematic, as Sweetser recognises. Is it really up to the hearer to decide whether the speaker has performed a speech act? Clearly, conversation involves a great deal of co-operation between speakers, and the assumption of certain shared goals. But is the issue of whether a speaker has performed a speech act really a matter for negotiation between the two participants? It seems that to claim this is at least to diverge from what is normally meant by a speech-act or utterance.

Moreover, there are some cases where a speech-act analysis would require a fairly radical divergence from what is normally meant. Consider the following example:

(24) If you run out of money, come to me for a loan.

I should point out that Example 24 is not the kind of text which Sweetser uses to exemplify SPEECH-ACT conjunctions. However, as already stressed, it would be useful to have an account which was general enough to cover this case as well. The alternative, of having distinct treatments for the two kinds of conditional imperative, seems to involve a measure of redundancy. In any case, if we were to attempt a SPEECH-ACT reading of this example ('If you ever run out of money, *I tell you to* come to me for a loan'), we would certainly encounter problems: we would have to envisage the speech act as occurring some way into the future, at some as-yet unknown moment. Moreover, if we interpret the text as expressing a habitual conditional (where we could substitute *if* with *whenever*), the speech act must also be allowed to occur arbitrarily often. Clearly, in either case we are stretching the notion of 'speech act' beyond its original meaning.

5 A new bipartite definition

5.1 Intention-based definitions of coherence relations

The problems mentioned in Sections 4.4 and 4.3 both point towards the need to include reference to the intentions which underlie an utterance in the definitions of SPEECH-ACT and EPISTEMIC relations. In fact, reference to intentions in relations has been long been advocated by those interested in discourse relations from the point of view of computational linguistics, and in particular from text generation. For Mann and Thompson (1988), a specification of the effect achieved on a reader by juxtaposing the two spans in a relation is a central component of its definition. This idea in turn gave rise to a procedural conception of relations, as planning operators in discourse structuring systems; for a review, see Hovy (1993). More recently, Moore and Paris (1993) and Moore and Pollack (1992) have also emphasised the need for a representation of intentions in discourse, though in their model it is proposed that intentional

relations should be represented alongside informational relations, rather than as an alternative to them.

While this work provides us with many of the concepts needed to think of relations in terms of intentions, none of it directly addresses the issue we have been concerned with, namely the problems with and tensions between the definitions of PRAGMATIC, EPISTEMIC and SPEECH-ACT relations. An intention-based approach to the problems is proposed in the next section.

5.2 Intention-based definitions for SEMANTIC and PRAGMATIC relations

The idea I am proposing is quite simple. Let us assume we are characterising the relation between two utterances U_1 and U_2 . Again, we assume that we can identify a basic relation R (cause, disjunction, or whatever) that is to be used in this characterisation, and that the task is to identify what this relation applies between. I am proposing that two classes of relation are sufficient to capture much of the data so far considered; retaining the bipartite terminology, these will be termed SEMANTIC and PRAGMATIC. In each case, a relation between U_1 and U_2 is expressed in terms of the intended effect of the complex utterance $U_1 + U_2$. We also make use of the functional notation introduced in Section 4.2. Two functions are needed: one function $c(U)$ returns the propositional content of an utterance U as before; the other function $ie(U)$ returns the intended effect of U . We can then define $ie(U_1 + U_2)$ in terms of R , U_1 , U_2 , and these two functions:

RELATION-TYPE	Definition
SEMANTIC	$ie(U_1 + U_2) = believes(hearer, R(c(U_1), c(U_2)))$.
PRAGMATIC	$ie(U_1 + U_2) = R(ie(U_1), ie(U_2))$.

In English: the intended effect of a SEMANTIC relation is that the hearer *believes* that the basic relation R holds between the *propositional contents* of the two related utterances; while the intended effect of a PRAGMATIC relation is that the basic relation *actually* holds between the *intended effects* of the two related utterances. In each case, the relation definition evaluates to a proposition, which can be thought of as the goal state which the complex utterance is expressed to achieve. In the case of SEMANTIC relations, this proposition is that the hearer believes a second, embedded proposition, namely that the relation applies. For PRAGMATIC relations, the proposition that the relation applies is directly asserted, rather than scoped within a hearer belief. The other difference between SEMANTIC and PRAGMATIC relations concerns which function is applied to U_1 and U_2 to obtain the propositions between which the relation is asserted to apply.

Before we proceed to some examples of these definitions in action, some axioms need to be specified for determining the intended effect $ie(U)$ of an atomic utterance U . While the notion of the intentions behind an utterance is notoriously ill-defined, I will be using it in the following technical, and fairly concrete, sense:

- The intended effect of an indicative sentence is that the hearer believes its propositional content.
- The intended effect of an imperative sentence is that the hearer performs the action in question.
- The intended effect of an interrogative sentence is that the hearer answers the question.

Note that in each case, the intended effect evaluates to a proposition.

We can now illustrate how the definitions work in a range of cases.

5.3 Causal relations

The following two texts are SEMANTIC and PRAGMATIC causal relations respectively:

(25) Bill was starving, so he ate five sandwiches.

(26) Bill ate five sandwiches, so he was [must have been] starving.

For Example 25, the propositional content of the first utterance is that Bill was starving; the propositional content of the second utterance is that Bill ate five sandwiches; and the intended effect of the complex utterance is that the reader believes that Bill's being starving caused Bill's eating of five sandwiches.

For Example 26, the intended effect of the first utterance is that the hearer believes that Bill ate five sandwiches; the intended effect of the second utterance is that the hearer believes that Bill was (must have been) starving; and the intended effect of the text as a whole is that the hearer's belief that Bill ate five sandwiches actually causes the hearer's belief that Bill was starving. This analysis seems to provide a good account of an argumentative relation. In describing the cause in such a case as something which the speaker intends to occur between two hearer beliefs, it provides the basis for an explanatory account of how argumentative relations achieve their persuasive impact; this is something which in Section 4.4 we argued was independently needed anyway.

Example 26, being argumentative, would be classified by Sweetser as EPISTEMIC. Now note that the new definition of PRAGMATIC relations applies equally well for SPEECH-ACT causal relations. For instance:

(27) We're late, so hurry up.

In this case, the intended effect of the first utterance is that the hearer believes that we are late; the intended effect of the second utterance—an imperative—is that the hearer hurries up; and the intended effect of the whole text is that the hearer's belief that we are late actually causes the hearer to hurry up. Again, this seems a plausible account of the text in question. In fact, it seems more direct than an account in terms of speech-act glosses: if we assume the speech act gloss ('so *I tell you to* hurry up') there is still a step to take between the hearer's understanding the cause of the speaker's statement and his actually hurrying up.

In comparing Examples 26 and 27 the main thing to note is that the work which in Sweetser's account is done by differences between the definitions of EPISTEMIC and SPEECH-ACT relations is done in the new account simply by the differences between the intended effects of an imperative and an indicative sentence. As we have argued in Section 4.3, these are differences which is something which any theory of sentence semantics needs to represent anyway. Moreover, in framing a single definition which covers EPISTEMIC and SPEECH-ACT relations, we are addressing the generalisations between the two classes of relation which, as noted in Section 4.1, are problematic for the previous accounts.

5.4 Disjunctive relations

Now consider some disjunctive relations.

- (28) The milk is in the fridge, or it is on the sideboard.
- (29) John is home, or somebody is picking up his newspapers.
- (30) Would you like to come round tonight? Or is your car still in the shop?

Example 28 is a SEMANTIC relation. The analysis here is straightforward: the propositional content of the first utterance is that the milk is in the fridge; the propositional content of the second utterance is that the milk is on the sideboard; and the intended effect of the complex utterance is that the hearer believes that either the milk is in the fridge or it is on the sideboard.

Example 29, which is EPISTEMIC in Sweetser's terms, is PRAGMATIC in ours. Accordingly, the intended effect of the first utterance is that the hearer believes that John is at home; the intended effect of the second utterance is that the hearer believes that someone is picking up John's newspapers; and the intended effect of the complex utterance is that either the hearer believes John is home, or the hearer believes someone is picking up his newspapers. This analysis seems at least as good as that given by the definition of EPISTEMIC relations.

Example 30 is also PRAGMATIC for the present account. Thus, the intended effect of the first utterance is that the hearer answers the question (i.e. says whether he would like to come round tonight); the intended effect of the second utterance is also that the hearer answers the question (i.e. says whether his car is still in the shop); and the intended effect of the complex utterance is that either the hearer answers one question, or answers the other. Again, this analysis seems to be perfectly good. In fact, it improves on the speech-act account: note that we do not have to introduce the problematic notion of 'conditional' speech acts.

5.5 Conditional relations

Now consider how conditional relations fare under the new definition:

- (31) If John goes to a party, he gets drunk.
- (32) If John went to the party, he was trying to infuriate Miriam.
- (33) How old are you, if it's not a cheeky question?

Example 31 is a SEMANTIC relation. The propositional content of the first utterance is that John goes to a party; the propositional content of the second utterance is that John gets drunk; and the intended effect of the complex utterance is that the hearer believes that if John goes to a party he gets drunk.

Example 32, which is EPISTEMIC in Sweetser's terms, is PRAGMATIC in ours. Accordingly, the intended effect of the first utterance is that the hearer believes that John went to the party; the intended effect of the second utterance is that the hearer believes that John was trying to infuriate Miriam; and the intended effect of the complex utterance is that if the hearer believes that John went to the party, the hearer believes that John was trying to infuriate Miriam. Again, this seems as good as the EPISTEMIC account.

Example 33, which is SPEECH-ACT in Sweetser's terms, is also PRAGMATIC for the new definition. Thus, the intended effect of the first utterance is that the hearer answers the question (i.e. says how old he is); the intended effect of the second utterance is that the hearer believes that the question is not cheeky; and the intended effect of the complex span is that if the hearer believes the question is not cheeky, he answers it. This seems to me to be

an improvement on the SPEECH-ACT account. For one thing we avoid the problematic notion of conditional speech acts, as before. Moreover, the analysis specifies that it is the *hearer's* judgement as to whether or not the question is cheeky on which something is conditional—the SPEECH-ACT account does not say anything about the perspective from which this is to be judged.

Finally, consider another type of conditional imperative, the kind which Sweetser's account does not treat:

(34) If you run out of money, come to me for a loan.

The new definition of PRAGMATIC works just as well for this example as for the others: the intended effect of the complex span here is that if the hearer believes he has run out of money, the hearer comes to the speaker for a loan.⁴ Note that the account works equally well if the conditional is interpreted as a habitual: there are none of the problems associated with conditional, future or repeated speech acts that would arise if we extended the SPEECH-ACT account.

5.6 Sequential relations

Finally, consider some sequential relations.

(35) John peeled the onions. Next he chopped them.

(36) John's unsuitable. To begin with, he's too young. Next, he's hotheaded...

(37) Peel the onions. Next, chop them.

Example 35 is a SEMANTIC relation: the intended effect of the complex utterance is that the hearer believes that John peeled the onions and then chopped them.

Example 36 is a PRAGMATIC relation, and would presumably be taken as EPISTEMIC by Sweetser. On the new analysis, the intended effect of the complex utterance is that the reader first believes that John is too young, and then believes that John is hotheaded. Again, this analysis seems as good as the EPISTEMIC analysis; we are still talking about the ordering of events in the epistemic world, but it is the hearer's world rather than the speaker's.

Example 37 is also PRAGMATIC in our terms. It is not the kind of example of a SPEECH-ACT sequence that gets discussed in the literature, because as we saw in Section 4.3, it would be wrongly analysed in this case. However, by the new definition it is unproblematic: the intended effect of the first utterance is that the hearer peels the onions; the intended effect of the second utterance is that the hearer chops them, and the intended effect of the complex utterance is that the hearer first peels the onions and then chops them.

⁴Note that the protasis has to be evaluated from the perspective of the hearer in this case. This reading certainly seems possible, but it might be thought that it could also be evaluated from the speaker's perspective (i.e. the hearer comes to the speaker for a loan if *the hearer believes* the speaker believes he has run out of money). If this is considered a problem, it could be borne in mind that the intended effect of an indicative utterance with content *p* can also be taken to be that the hearer believes that the speaker believes that *p*. With this alternative interpretation, both of the required perspectives are available.

6 An assessment of the new definitions

6.1 Advantages of the new definitions

I will begin by briefly summarising the advantages of the new definitions of SEMANTIC and PRAGMATIC.

Firstly, the definitions permit some useful generalisations across EPISTEMIC and SPEECH-ACT relations. Because the intended effect of an imperative speech act is different from that of an assertive, expressing the definition of PRAGMATIC relations in terms of intended effects allows a single characterisation of these two kinds of relation.

Secondly, the new definitions provide a clearer picture of compositionality. They are compositional in two senses. Firstly, they allow a consistent account of the variable that changes when we move from SEMANTIC to PRAGMATIC relations—namely, the identity of the function ($c(U)$ or $ie(U)$) that is applied to the two related utterances. Moreover, it is compositional in a more conventional sense. When we think of compositional semantics, we say that the meaning of a sentence can be derived from the meaning of its constituents, plus their manner of combination. The new definition allows us to make a similar statement about intended effects, for PRAGMATIC relations: we can say that the intended effect of the complex text span formed by a PRAGMATIC relation is derivable from the intended effects of its two constituent spans, together with the relation which links them.

Thirdly, the new definitions provide an explanatory story for argumentative relations. The new definition of PRAGMATIC relations, as applied to an argumentative text, does not simply represent it as a description of the speaker's own thought processes, but specifies how it is that it can, if successful, have a rhetorical impact on the hearer in order to be successful.

Fourthly, the new definitions provide for a better treatment of conditional speech-acts. As we have seen, by making the conditional events the intended effects of utterances, rather than the utterances themselves, we avoid a number of problems in analysing texts where *if* and *or* are used to conjoin imperative or interrogative sentences.

Finally, the new definitions appear to provide the basis for a more general treatment of imperatives. We have seen that the new definition of PRAGMATIC relations allows them to apply more widely, to sequences of imperatives and conditional instructions, as well as to the set of cases dealt with under the rubric of SPEECH-ACT relations.

6.2 Some problems for the new definitions

While there are many advantages with the new definitions, I should also note two cases where they do not apply straightforwardly, and where some modification of the account appears necessary. Both of these cases involve relations which Sweetser has analysed as SPEECH-ACT; and in each case there are aspects of this definition which the new definition of PRAGMATIC does not apparently capture.

6.2.1 Problematic SPEECH-ACT causes

Consider the following case of a causal SPEECH-ACT relation.

(38) Since you asked nicely, I'm 78.

In this text, the definition of SPEECH-ACT relations seems more natural: it does indeed seem as though the primary effect here is informational, with the hearer being told why the speaker

is making a certain utterance. If we used the new conception of PRAGMATIC relations, we get a strange analysis, where the intended effect of the complex utterance is that the hearer's belief that he asked nicely should cause the hearer to *believe* that the speaker is 78. This seems off the mark.

There are other cases where the new PRAGMATIC definition is possible, but not particularly attractive. For instance:

(39) The rules can't be broken, so 'no'.

(40) Since we're on the subject, when was George Washington born?

In each case, the original definition of SPEECH-ACT relations is a lot more natural. In summary, it must be acknowledged that there are cases where the new definition of PRAGMATIC relations doesn't work. So we haven't eliminated the need for a class of SPEECH-ACT relations altogether; we have just made it smaller, and possibly more homogeneous.

The question therefore remains as to how to define the class of SPEECH-ACT relations. It might be thought that, now that the class is smaller, we could frame a definition for it on the same terms as those given for SEMANTIC and PRAGMATIC relations. We could imagine re-introducing the function $u(U)$, which for an utterance U returns the utterance itself, and suggesting that for SPEECH-ACT relations, the intended effect of the complex utterance $ie(U_1 + U_2)$ given some basic relation R is something like $believes(hearer, R(u(U_1), u(U_2)))$. But this doesn't work. The problem is the same as before: we want to apply different functions to U_1 and U_2 . (It is not because the speaker *says* that we're on the subject of George Washington that he asks about when he was born, but because he *is* on the subject, or believes he is.) The attractive feature of the new definitions for SEMANTIC and PRAGMATIC is that the same function is applied to both utterances.

6.2.2 Problematic SPEECH-ACT conditionals

Another class of problematic cases includes a group of conditionals of the following form:

(41) If you're hungry, there's some soup in the fridge.

Sweetser would analyse this as a SPEECH-ACT conditional: 'if you're hungry, then [let us consider that] I tell you that there's soup in the fridge'—the idea being that the assertive speech act is not understood to have been made unless the hearer is hungry. This certainly seems strange, for the reasons already discussed, but in this case it is not clear that the new definition of PRAGMATIC is any easier to understand. On the PRAGMATIC account, the intended effect of the complex utterance is that if the hearer believes that he is hungry, then he also believes that there's soup in the fridge. But surely, given the text in question, the hearer is going to believe there's soup in the fridge regardless of whether he believes he is hungry? The notion of conditional hearer actions (the obeying of orders, the answering of questions) is perfectly intelligible, but the notion of conditional hearer beliefs in this context seems at least as strange as that of conditional speech acts.

7 Future directions: generalising over the intentions of participants and protagonists

Before concluding, I will sketch an extension (still very speculative) to the current definitions of SEMANTIC and PRAGMATIC relations, which points to further possible generalisations using

the concept of intentions. This involves the introduction of a second bipartite classification of relations, into CAUSE-DRIVEN and RESULT-DRIVEN relations, which I have discussed elsewhere; see e.g. Knott and Mellish (1996).

7.1 CAUSE-DRIVEN and RESULT-DRIVEN relations

The new distinction can be introduced by considering the following pair of sentences.

(42) John had been up all night, but he looked fresh as a daisy.

(43) John was hungry, but he couldn't find anything to eat.

But is typically analysed as signalling a violated expectation of some kind. For instance, in Example 42, the fact that John has been up all night leads to the expectation that he will look tired, which is then not forthcoming. The analysis is often formalised in terms of presupposed defeasible rules; see e.g. Knott and Mellish (1996), Lagerwerf (1998). In Example 42, for instance, the presupposed rule would include 'X has been up all night' on its left-hand side, and 'X looks tired' on its right: the *but* indicates that this rule is defeated in the situation being described.

While this analysis works for Example 42, it does not work for Example 43, as is noted by Spooren (1989) and Knott and Mellish (1996). To conform to the proposed representation, we would have to assume a rule of inference whose premise is that X is hungry and whose conclusion is that X finds something to eat. More generally, to deal with cases like Example 43, we would need a rule of inference whose premise is the fact that an agent has a goal, and whose conclusion is that the goal is achieved. However, this rule of inference is clearly not always a valid one: agents in the world frequently have goals which cannot be achieved. Of course, if we also know that the conditions in the world are such that the goal is achievable, then the inference is perfectly possible to make. But that is precisely what is not known in an example like 43—until the clause introduced by *but* is processed.⁵ Knott and Mellish propose that instead of attempting to reduce cases like Example 43 to a violated expectation, we should distinguish between two kinds of *but*: one, termed CAUSE-DRIVEN, which signals a violated expectation, and one, termed RESULT-DRIVEN, which signals a frustrated plan. The notion of a protagonist's plan is thus accepted as a primitive in the account. A generalisation between the two types of *but* is then proposed, which draws on the fact that a defeasible rule can be used in two ways: **deductively**, to generate an expectation of its right-hand side based on knowledge of its left-hand side, or **abductively**, to plan an action which achieves a state of affairs which features on its left-hand side in pursuance of a goal to achieve the state of affairs on its right-hand side. See Knott and Mellish for details of this distinction. The crucial point to note about the distinction for present purposes is that the definition of RESULT-DRIVEN relations makes explicit reference to the *goals* of a protagonist.

⁵Spooren apparently suggests that a violated expectation *does* arise in a case such as Example 43, as the necessary additional premise (in this case, that there is food to be found) is an implicated by the utterance expressing the protagonist's goal (i.e. *John was hungry*). While we must certainly implicate the fact that John has a goal if this is not explicitly stated, there seems no good reason to suggest that the speaker's informing the hearer about this goal should lead to an implicature that the goal is satisfiable. If anything, hearers are used to hearing about situations in which a protagonist's goals are *not* satisfied: this is certainly the predominant case in narratives.

7.2 Speaker and protagonist goals

The presence of goal-based primitives in the definitions of PRAGMATIC and RESULT-BASED relations opens up interesting possibilities for further generalisations. Specifically, there seem to be some interesting similarities between the goals a speaker pursues by making utterances, and the goals a protagonist being described in a discourse pursues by taking actions. Compare these two texts, for instance:

(44) Bill told Jim to go to bed, but Jim wasn't tired.

(45) Bill [to Jim]: Go to bed!
Jim [to Bill]: But I'm not tired!

In each case, we can talk about Bill seeking to satisfy a goal that Jim goes to bed by telling him to do so, and thereby discovering a circumstance that indicates that this action on its own will not be enough. In each case we can also use the conjunction *but*. But the two texts are nevertheless quite different: in Example 44, the intention is that of the protagonist in a narrative monologue, while in Example 45 it is that of the first speaker in a conversational exchange; moreover, Example 44 contains only indicative clauses, while Example 45 contains an imperative clause too. It seems quite likely that the definition of PRAGMATIC which is required to account for the imperatives in Example 45 will also suffice to introduce the goals necessary for the analysis of the *but* in this example as RESULT-DRIVEN. This possibility is something which will be pursued in further work.

8 Summary

This paper began by noting some difficulties with existing notions of PRAGMATIC, EPISTEMIC and SPEECH-ACT relations. A solution to some of these problems was then proposed, in the form of a new definition of PRAGMATIC relations framed in terms of the intended effects of a speaker's utterances. This new definition has a number of advantages over its predecessors, although there are still some remaining problems with it. Finally, a possible extension to the account was presented, by introducing a new category of RESULT-DRIVEN relations: there appear to be some interesting generalisations in prospect relating to speaker and protagonist goals. Whether these can be exploited in improved definitions of SEMANTIC and PRAGMATIC relations is a matter for future research.

References

- Bateman, J. and Rondhuis, K. J. (1997). "coherence relations": Towards a general specification. *Discourse Processes*, **24**(1).
- Davison, A. (1973). *Performative verbs, adverbs and felicity conditions: an inquiry into the nature of performative verbs*. Ph.D. thesis, University of Chicago.
- Elhadad, M. and McKeown, K. R. (1990). Generating connectives. In *COLING-90*, pages 97–101.
- Halliday, M. and Hasan, R. (1976). *Cohesion in English*. Longman.

- Hobbs, J. R. (1985). On the coherence and structure of discourse. Technical Report CSLI-85-37, Center for the Study of Language and Information, Stanford University.
- Hovy, E. (1993). Automated discourse generation using discourse structure relations. *Artificial Intelligence*, **63**, 341–385.
- Knott, A. and Mellish, C. (1996). A feature-based account of the relations signalled by sentence and clause connectives. *Language and Speech*, **39**(2–3), 143–183.
- Lagerwerf, L. (1998). *Causal connectives have presuppositions*. Ph.D. thesis, Katholieke Universiteit Brabant, The Hague.
- Mann, W. C. and Thompson, S. A. (1988). Rhetorical structure theory: A theory of text organization. *Text*, **8**(3), 243–281.
- Martin, J. (1992). *English Text: System and Structure*. Benjamin, Amsterdam.
- Moore, J. and Pollack, M. (1992). A problem for RST: The need for multi-level discourse analysis. *Computational Linguistics*, **18**, 537–544.
- Moore, J. D. and Paris, C. L. (1993). Planning text for advisory dialogues: Capturing intentional and rhetorical information. *Computational Linguistics*, **19**, 651–694.
- Oversteegen, L. (1997). On the pragmatic nature of causal and contrastive connectives. *Discourse Processes*, **24**(1), 51–86.
- Redeker, G. (1990). Ideational and pragmatic markers of discourse structure. *Journal of Pragmatics*, **14**, 367–381.
- Ross, J. R. (1967). *Constraints on variables in syntax*. Ph.D. thesis, MIT. Now reprinted as *Infinite Syntax*. Norwood, NJ: Ablex, 1986.
- Sanders, T. (1997). Coherence relations in context; on the categorization of positive causal relations. *Discourse Processes*, **24**(1), 119–147.
- Sanders, T. J. M., Spooren, W. P. M., and Noordman, L. G. M. (1992). Towards a taxonomy of coherence relations. *Discourse Processes*, **15**, 1–35.
- Spooren, W. (1989). *Some Aspects of the Form and Interpretation of Global Contrastive Coherence Relations*. Ph.D. thesis, Catholic University of Nijmegen, the Netherlands.
- Sweetser, E. (1990). *From Etymology to Pragmatics: Metaphorical and Cultural Aspects of Semantic Structure*. Cambridge University Press.
- van Dijk, T. A. (1979). Pragmatic connectives. *Journal of Pragmatics*, **3**, 447–456.