Paper 1

Imperatives in Dialogue

Alex Lascarides
Division of Informatics, University of Edinburgh
alex@cogsci.ed.ac.uk

Abstract

In this paper, we offer a semantic analysis of imperatives. We explore the effects of context on their interpretation, particularly on the content of the action to be performed, and whether or not the imperative is commanded. We demonstrate that by utilising a dynamic, discourse semantics which features *rhetorical relations* such as *Narration*, *Elaboration* and *Correction*, we can capture the discourse effects as a byproduct of discourse update (i.e., the dynamic construction of logical forms). We argue that this has a number of advantages over static approaches and over plan-recognition techniques for interpreting imperatives.

1 Introduction

An adequate theory of dialogue interpretation requires a satisfactory account of imperatives. In this paper, we will address two inter-related questions. What is their compositional semantics? And how does the discourse context affect their content?

There are several puzzles which need to be addressed. The first concerns compositional semantics. Ross (1941) observed that imperatives aren't closed under logical consequence: *post the letter* does not entail *post or burn the letter*, even though the proposition that the letter is posted entails that it is posted or burned. This makes a straightforward analysis within modal logic problematic: if ! is a 'standard' modal operator and $!A$ means A *is commanded*, then $A \models B$ will incorrectly entail $!A \models !B$, regardless of the accessibility constraints on the !-worlds.[1] Segerberg (1990) bypasses the paradox via a modal logic of action. But the semantics is static and the base language is propositional. We will find that by making the semantics dynamic, the account can be significantly simplified.

The second puzzle concerns the interaction between context and imperatives. How does context—both linguistic and non-linguistic—affect the content of imperatives, particularly the content of the action, and whether or not the imperative is commanded? Consider, for example, the discourses (1.1), adapted from Webber *et al.* (1995):

(1.1) a. Go to Fred's office and get the red file folder.
b. Go to Fred's office and refile the red file folder.
c. John went to Fred's office. He got the red file folder.
d. John went to Fred's office. He refiled the red file folder.

[1] A similar problem holds for deontic statements: *You must post the letter* doesn't entail *You must post the letter or burn the letter.*

BI-DIALOG 2001 (P. Kühnlein, H. Rieser, & H. Zeevat (eds)).

Division of Informatics, University of Edinburgh
alex@cogsci.ed.ac.uk.

Discourses (1.1ab) both implicate that the actions should be performed in the order described and the second action is performed in Fred's office. (1.1a) implicates that the red file folder is in Fred's office whereas (1.1b) doesn't implicate this. Similar spatio-temporal implicatures hold of the indicatives versions (1.1cd).

Segmented Discourse Representation Theory (SDRT, Asher 1993, Lascarides and Asher 1993) accounts for the implicatures in (1.1cd) by stipulating within a dynamic semantic setting how one computes the *rhetorical relation* which connects the propositions (namely, *Narration* for (1.1cd)), and stipulating how such relations constrain the content of its arguments (e.g., the spatio-temporal content described above follows from the semantics of *Narration*). We aim to model imperatives in a similar manner. That is, we aim to account for their implicatures by identifying their rhetorical role. This involves specifying the semantics of the relations that take imperatives as arguments, and stipulating a precise default axiomatisation of how such rhetorical relations are computed on the basis of both linguistic and non-linguistic knowledge sources.

We will show that SDRT can provide an entirely uniform analysis of the imperative vs. indicative examples in (1.1), which is desirable given their similar implicatures. The uniformity rests on the fact that the SDRT axioms of interpretation that apply to these discourses are neutral with respect to sentence mood, instead relying on other compositional and lexical semantic features. In contrast, it would be hard to achieve such a uniform analysis with with plan recognition approaches (e.g., Grosz and Sidner 1986, 1990, Litman and Allen 1990, Lochbaum 1998), where interpreting the current utterance utilises only the *goals* of the prior utterances, rather than their compositional and lexical semantics directly. This is because the goals of indicatives (typically, that the interpreter believe the proposition) are radically different from imperatives (typically, that the interpreter perform the action). The similar interpretations of (1.1a/c) and (1.1b/d) suggest that the goal of the prior clause isn't primary in these cases.

This is not to deny the importance of beliefs and goals in interpretation, however. The fact that falling downstairs is undesirable whereas going to the hardware store is not underlies the difference between (1.2a) (where the imperative is not commanded) and (1.2b) (where it is):

(1.2) a. Go straight on and you'll fall down the stairs.
 b. Come home by 5pm and we can go to the hardware store before it closes.

(1.3) a. A: How does one make lasagne?
 b. B: Chop onions and fry with mince and tomatoes, boil the pasta, make a cheese sauce, assemble it, and bake in the oven for 30 minutes.

(1.4) a. A: Go straight on for 5cm.
 b. B: That takes me right into the crevasse.
 c. A: Go left then.

Similarly, the inference that the rhetorical role of (1.3b) is to provide sufficient information that A can compute an answer to his question (1.3a) is calculable from Gricean style principles of rationality and cooperativity (e.g., Cohen and Levesque 1990, Lascarides and Asher, 1999). And whether or not such responses to questions are commanded depends on the content of the question: (1.3b) is not commanded; but an imperative *is* commanded if it serves as a response to a question whose answers all implicate that the questioner is the agent of a deontic attitude (e.g., *Where should I go now?*). Finally, in (1.4), taken from the HCRC map task corpus, the undesirability of falling into the crevasse helps one infer that the request (1.4a) is 'cancelled' and replaced by (1.4c).

Our hypothesis is that for all these examples, the interplay between content, domain knowledge and cognitive states can be captured within the semantics of the rhetorical relations, and the axioms one uses to compute them during the construction of the discourse's logical form. We will test this by incorporating a semantic analysis of imperatives into SDRT. In an attempt to do justice to the complexity of interaction between the different information sources that contribute to interpretation—both conventional and non-conventional—many theories assume a radically unmodular framework, so that a single reasoning process can access the different kinds of information at any time (e.g., Hobbs *et al*, 1993). SDRT takes a different approach, assuming a high degree of modularity: reasoning with conventional clues about interpretation is kept separate from reasoning with non-conventional clues, but there are interactions between them.

2 The Compositional Semantics of Imperatives

Segerberg (1990) offers a semantics of imperatives which bypasses Ross' paradox. He augments a propositional language with two operators. First, the action operator δ takes formulae into *action terms*: e.g., if p is a propositional variable, then δp is an action term, corresponding to the action of seeing to it that p is true. Second, the command operator ! takes action terms into *practical formulae*; these are essentially the imperatives. So $!\delta p$ is a well-formed practical formula standing for "making p true is commanded"; $q \rightarrow !\delta p$ is also a practical formula standing for the conditional imperative "if q is true then making p true is commanded"; and $!p$ is ill-formed.

An action term δp denotes a set of pairs of possible worlds. Intuitively, the first world of each pair corresponds to a possible state of affairs in which the action can be performed, and the second world describes a possible outcome of performing the action in that first world. Furthermore, for each action a there is a corresponding modal operator $[a]$: $[a]p$ is true in a model M at a world w just in case p is true at all worlds w' such that $\langle w, w' \rangle \in [\![a]\!]^M$ (as we'll see shortly, $[\![a]\!]^M$ is a rigid designator). In other words, p is a necessary postcondition of a. p' is a precondition if $\neg p' \rightarrow [a]\bot$.

Plans are also terms, constituting a *sequence* of actions $a_1; a_2; \ldots a_n$. These also denote sets of pairs of worlds: $\langle w, w' \rangle \in [\![a_1; \ldots a_n]\!]$ iff $\exists w_1, \ldots, w_{n-1}$ such that $\langle w, w_1 \rangle \in [\![a_1]\!], \ldots, \langle w_{n-1}, w' \rangle \in [\![a_n]\!]$. So the possible consequences of doing a_i must be compatible with the preconditions of a_{i+1}. Finally, one express *free choice*: $[\![a_1 + a_2]\!] = [\![a_1]\!] \cup [\![a_2]\!]$.

The formula δp receives its model-theoretic denotation via a function D in the model which takes propositions (i.e., a set of worlds) to actions; i.e., $[\![\delta p]\!]^M =_{def} D[\![p]\!]^M$. D satisfies the following constraint:

$$D[\![p]\!]^M \subseteq \{\langle w, w' \rangle : w' \in [\![p]\!]^M\}$$

This makes $[\delta p]p$ true at all worlds in all models; i.e., making p true guarantees that p is true.

Since the above constraint on D uses $\subseteq$ rather than $=$, the logical relationships among actions is almost entirely impotent, in the sense that it's *not* the case that if $A \models B$, then $[\![\delta A]\!] \subseteq [\![\delta B]\!]$ (i.e., all the actions for making A true aren't necessarily also actions for making B true). This is problematic, because reasoning with these action terms, and hence planning, becomes impractical because of their weak logic.

The semantics for theoretical formulas (i.e., formulas that contain no ! operator) is essentially Kripkean, and a logic of satisfaction $\models$ for theoretical formulas is defined in the usual way. Practical formulae have their own distinct but related logic: the logic $\models_r$ of requirement or 'commanding'. This logical consequence relation exploits the notion of a command system Γ *requiring* a formula; written $\Gamma \models_r^{M,w} A$, where A now is either theoretical or practical. A command system Γ is a semantic primitive, which stipulates which actions the authority commands; or more accurately, which action he commands in which situations.

More formally, a command system Γ is a set of *command sets*, one for each possible world in the model. And a command set Γ_w is a set of actions; intuitively, the actions that the authority commands in the world w (and any one of these actions may in fact only be describable by several imperatives). One can now define the logical consequence relation $\Gamma \models_r^{M,w} A$, of the command system Γ *requiring* a formula A at the world w in the model M. We present highlights here:

1. $\Gamma \models_r^{M,w} p$ iff $w \in [\![p]\!]^M$.

2. $\Gamma \models_r^{M,w} A \rightarrow B$ iff if $\Gamma \models_r^{M,w} A$, then $\Gamma \models_r^{M,w} B$.

3. $\Gamma \models_r^{M,w} [a]B$ iff for all w' such that $\langle w, w' \rangle \in [\![a]\!]^M$, $\Gamma \models_r^{M,w'} B$.

4. $\Gamma \models_r^{M,w} !a$ iff $[\![a]\!]^M \subseteq \Gamma_w$.

Note that for any theoretical formula A, $\Gamma \models_r^{M,w} A$ iff $M \models^w A$, reflecting the intuition that wishful thinking can't make things true. Furthermore, even if $A \models^M B$ in the logic of satisfaction and $[\![\delta A]\!]^M \in \Gamma_w$, it does *not* follow that $[\![\delta B]\!]^M \in \Gamma_w$. So $\Gamma \models_r^{M,w} !\delta p$ doesn't entail $\Gamma \models_r^{M,w} !\delta(p \vee q)$, thereby bypassing Ross' paradox. Unfortunately, it's also the case that $\Gamma \models_r^{M,w} !\delta(p \wedge q)$ doesn't entail $\Gamma \models_r^{M,w} !\delta p$, indicating that the logic of commanding, as well as of actions and plans, is perhaps weaker than it should be.

This account is at best incomplete. It cannot be used to analyse imperatives with quantifiers, since its base language is propositional. And it cannot be used to explore the interaction between content and anaphora since it's static. In fact, the static semantics would yield a highly complex translation from natural language imperatives in discourse into logical form. For note that the content of (1.5) is not adequately expressed by (1.5′), where A represents the proposition that you go to the traffic lights and B represents the proposition that there's a roundabout to your right.

(1.5) Go to the traffic lights. There's a roundabout to your right.

(1.5′) $!\delta A \wedge B$

(1.5″) $\delta A \wedge [\delta A]B$

This is because (1.5′) entails that the roundabout is to your right *now* (i.e., before the action is performed), rather than being conditional on the action being performed. In fact, its intuitive interpretation is captured in Segerberg's semantics by the formula (1.5″). But constructing such a formula on the syntax/semantics interface is impractical.

We'll shortly see that incorporating action terms into a dynamic discourse semantics simplifies this analysis. Not only will we achieve a uniform semantic construction procedure within the grammar (cf. example (1.5) above). But we can also abandon altogether Segerberg's command system and the logic of requirement. Whether or not an imperative is commanded will not be determined by a semantic primitive (i.e., the command system), but rather by the semantic consequences of its rhetorical connection to the context, which in turn is inferred from a wide variety of knowledge sources, both linguistic and non-linguistic.

Finally, as we mentioned before, avoiding Ross' paradox by sacrificing the capacity to reason about actions and commands is problematic. Hare (1967) takes a different view, arguing that one shouldn't avoid Ross' paradox at all. He suggests that $A \vdash A \vee B$ *is* in fact valid when A and B are requests. But it doesn't appear to be valid because of Gricean-style scalar implicatures. But this is unsatisfactory too, because no details are given of how scalar implicatures would have the desired effect. An alternative solution is to include some contextually determined formula ϕ within the postconditions of the action:

(1.6) $[\![\delta p]\!]^M = \{\langle w, w' \rangle : w' \in [\![p \wedge \phi]\!]\}$

The problem now is to compute the value of ϕ in different contexts. We'll argue that reasoning about the rhetorical role of the imperative goes some way towards systematically stipulating the content of ϕ in (1.6); for its rhetorical function will encode why it was uttered, and for what purpose.

3 Going Dynamic

The whole notion of meaning is reconstrued in dynamic semantics as a *relation* between an input context and an output context; this is known as the *context change potential* or CCP of a formula. These contexts can be characterized extensionally as assignment functions, which map the formula's variables to individuals in the model. However, to analyse imperatives and modal action operators, we need an intensional dimension. So we make contexts a world assignment pair (w, f). Thus the truth definition of a formula K will define exactly when K relates an input context (w, f) to an output context (w', g).

In Discourse Representation Theory (DRT, Kamp and Reyle 1993), a discourse is represented by a *discourse representation structure* or DRS, which is a pair consisting of a set of discourse referents (i.e., the individuals and events that the discourse is about) and a set of DRS-conditions (these convey properties and relations among the discourse referents). Since DRS-conditions can themselves include DRSs, DRSs are recursive. The syntax and semantics is as follows:

Syntax of DRT

Suppose $U \subseteq$ *Discourse-Referents.* Then the well-formed DRSs K and DRS conditions γ are defined recursively:

$$K \quad := \quad \langle U, 0\rangle \mid K^\frown\gamma$$

Let $R \in$ Predicates be an n-ary predicate and $x_1, \cdots, x_n$ be discourse referents.

$$\gamma \quad := \quad R(x_1, \cdots, x_n) \mid \neg K \mid K_1 \Rightarrow K_2 \mid K_1 \vee K_2.$$

The Semantics of DRSs

The truth definition involves embedding DRSs into a standard Tarskian model M; so $M = \langle A_M, W_M, I_M\rangle$, where A_M is a set of individuals; W_M is a set of worlds, and I_M is a function which assigns n-ary predicates at a world w a set of n-tuples of A_M. We define simultaneously the model theoretic transition P and the satisfaction of conditions V relative to the model M. From a dynamic logic perspective, P yields a change in the assignment function, extending the input over newly introduced discourse referents, while V treats the other DRS elements as tests.

The Truth Definition:

$$\begin{array}{rcl}
(w,f)P_M\langle U,\emptyset\rangle(w',g) & \textit{iff} & w = w' \wedge f \subseteq g \ \wedge \ \mathit{dom}(g) = \mathit{dom}(f) \cup U \\
(w,f) \in V_M(R(x_1,\cdots,x_n)) & \textit{iff} & (f(x_1),\cdots,f(x_n)) \in I_M(R)(w) \\
(w,f) \in V_M(\neg K) & \textit{iff} & \neg\exists g\,(w,f)P_M(K)(w,g) \\
(w,f) \in V_M(K \Rightarrow K') & \textit{iff} & \forall g\,((w,f)P_M(K)(w,g) \rightarrow \exists\, h\,(w,g)P_M(K')(w,h)) \\
(w,f) \in V_M(K \vee K') & \textit{iff} & \exists\, g\,(w,f)P_M(K)(w,g) \vee \exists h\,(w,f)P_M(K')(w,h) \\
(w,f)P_M(K^\frown\gamma)(w',g) & \textit{iff} & w = w' \wedge (w,f)P_M(K)(w,g) \wedge (w,g) \in V_M(\gamma)
\end{array}$$

As yet, we've not used the possible world component. But whereas extensional formulae transform the variable assignment function, action terms will transform the possible world (as well). And the dynamic semantics of DRS conditions of the form $[a]K$ invokes quantification over worlds.

More formally, we extend the language as follows:

1. If K is a DRS formula (e.g., K could be a DRS), then δK is an action term;
2. If a_1 and a_2 are action terms, then so are $a_1; a_2$ and $a_1 + a_2$.
3. If a is an action term and K is a DRS formula, then $[a]K$ is a DRS formula.
4. If K is a DRS formula and a is an action term, then $K \rightarrow a$ is a DRS condition.

The truth conditions of these new action terms and formulae are again defined in terms of a model theoretic transition P_M and satisfiability conditions V_M. The characteristic CCP of action terms is that they change the world parameter (see clause 1. below):

1. $(w,f)P_M(\delta K)(w',g)$ iff $(w',f)P_M(K)(w',g)$
2. $[\![a_1; a_2]\!]^M = [\![a_1]\!]^M \circ [\![a_2]\!]^M$
   (i.e., $(w,f)P_M(a_1;a_2)(w'',h)$ iff there is a pair (w',g) such that $(w,f)P_M(a_1)(w',g)$ and $(w',g)P_M(a_2)(w'',h)$).
   $[\![a_1 + a_2]\!]^M = [\![a_1]\!]^M \cup [\![a_2]\!]^M$
   (i.e,. $(w,f)P_M(a_1 + a_2)(w',g)$ iff $(w,f)P_M(a_1)(w',g)$ or $(w,f)P_M(a_2)(w',g)$).

3. $(w, f)P_M(K \to a)(w', g)$ iff either:

   (a) $\neg\exists h(w, f)P_M(K)(w, h) \wedge (w', g) = (w, f)$; or

   (b) $\exists h(w, f)P_M(K)(w, h) \wedge (w, h)P_M(a)(w', g)$

4. $(w, f) \in V_M[a]K'$ iff for every g and for every w' such that $(w, f)P_M(a)(w', g)$, $\exists h(w'g)(K)^M(w', h)$.

Note that, thanks to condition 2 above, the denotation of the complex action (1.7) is one where the individual who talks is the same as the individual who walks:

(1.7) $\delta\begin{array}{|l|}\hline \text{x} \\ \hline \text{walk(x)} \\ \hline \end{array};\delta\begin{array}{|l|}\hline \\ \hline \text{talk(x)} \\ \hline \end{array}$

And guarded actions (i.e., formulae of the form $K \to a$) can be the basis of conditional commands:

(1.8) If you want to get an A, study hard.

Overall, then, we will represent imperatives in DRT as action terms. For example, we assume that the grammar generates the action term (1.9′) for (1.9) (we've simplified slightly by ignoring temporal information):

(1.9) Walk!

(1.9′) $\delta\begin{array}{|l|}\hline \text{u,e} \\ \hline \text{walk(e,u)} \\ \hline \end{array}$

The discourse referent u is the addressee: we assume a conventional default within the grammar which generates u from the imperative sentence mood. This conventional default can be overridden when the subject is explicitly given (e.g., *Someone close the door!*) or in a sufficiently rich discourse context, as described in Lascarides and Copestake (1998).

Semantically, the defining characteristic of a discourse which includes a commanded imperative is that its CCP changes the input world into an output one where the action has been performed. (1.9′) changes the world this way, and thus it represents a discourse where the imperative (of walking) is commanded. The CCP of the DRS representing (1.8) also captures the command status of the imperative in the required way: the dynamic semantics of guarded actions means that the imperative is commanded if you want to get an A, and it's not commanded otherwise. This dynamic characterisation of an imperative makes introducing a semantic primitive for stipulating what's commanded and what's not redundant: we can rely instead on the signature CCP that the world is transformed into one where the action has been performed.[2] And so it can replace Segerberg's notion of a command system and the accompanying logic $\models_r^M$ of requirement, thereby considerably simplifying the semantics. We would need to bypass Ross' paradox, however, by including contextually specified information ϕ in the semantics of action terms (cf. the definition (1.6)). We return to this shortly.

A further advantage of the dynamic view is that, unlike Segerberg's static analysis, the semantics of the formula (1.5′) now captures the intuitive interpretation of (1.5) (as before, A stands for *you go to the traffic lights*, B stands for *there's a roundabout to your right*, and ; now stands for *dynamic and*):

(1.5) Go to the traffic lights. There's a roundabout to your right.

(1.5′) $\delta A; B$

[2] This can even apply to imperatives where one seemingly doesn't have to do anything to discharge the command; e.g., *Don't move!* and *Keep quiet!*

However, further investigation shows that this DRT-based analysis is flawed. As we saw earlier, not all imperatives are commanded, *even when* there are no linguistically explicit clues present to indicate this (e.g., (1.3)). In fact, maintaining the DRS language as it stands leads to one of two undesirable consequences. We either maintain a simple DRS-construction procedure for imperatives (i.e., use an action term, followed by ; if subsequent clauses are present), and therefore predict the wrong semantics of dialogues like (1.3). Similarly, this construction procedure would predict the wrong interpretation of (1.10a):

(1.10) a. Go to Fred's office. Take the file with you.

b. John went to Fred's office. He took the file with him.

The natural interpretation of (1.10a) is one where the actions are to be performed at the same time, rather than in sequence. The undesirable alternative is to generate several DRSs when updating the context with an imperative, one for each possible semantic contribution to the discourse. But proliferating ambiguity is undesirable.

These problems are in fact similar to the problems with DRT's analysis of temporal discourse. Kamp and Reyle (1993) note that their rules for DRS construction handle only those simple past-tensed discourses where event sentences move the time line forward (e.g., (1.1cd)). But not all discourses behave this way (e.g., (1.10b)). Here, we see that the DRT semantics of imperatives handles just those discourses where the imperative is commanded and the subsequent utterances should be interpreted with respect to a context where the action has been performed. But not all imperatives have this effect on content, as (1.2a), (1.3) and (1.10b) attest.

In view of these problems, we will maintain the analysis of imperatives as dynamic action terms, but incorporate it into SDRT. We'll then use SDRT's semantics of rhetorical relations to capture the various contributions imperatives can make to the overall content of the discourse.

4 Imperatives and Rhetorical Relations

Discourse is represented in SDRT as a recursive structure of labelled DRSs with rhetorical relations between the labels. For example, the logical forms of (1.1a) and (1.1c) are shown below (in slightly simplified form, since we have ignored presuppositions, tense and anaphora):

(1.1) a. Go to Fred's office and get the red file folder.

a′.
$$\begin{array}{|l|}\hline \pi_1, \pi_2 \\ \hline \pi_1 : \delta \begin{array}{|l|}\hline \text{f,u,o,}e_1 \\ \hline \text{fred(f), office(o),} \\ \text{own(f,o), go-to(}e_1\text{,u,o)} \\ \hline \end{array} \quad \pi_2 : \delta \begin{array}{|l|}\hline \text{r,}e_2 \\ \hline \text{red-file-folder(r), get(}e_2\text{, u,r)} \\ \hline \end{array} \\ Narration(\pi_1, \pi_2) \\ \hline \end{array}$$

c. John went to Fred's office. He got the red file folder.

c′.
$$\begin{array}{|l|}\hline \pi_1, \pi_2 \\ \hline \pi_1 : \begin{array}{|l|}\hline \text{f,j,o,}e_1 \\ \hline \text{john(j), fred(f), office(o),} \\ \text{own(f,o), go-to(}e_1\text{,u,o)} \\ \hline \end{array} \quad \pi_2 : \begin{array}{|l|}\hline \text{r,}e_2 \\ \hline \text{red-file-folder(r), get(}e_2\text{, u,r)} \\ \hline \end{array} \\ Narration(\pi_1, \pi_2) \\ \hline \end{array}$$

Note that the rhetorical relation *Narration* is used in both logical forms.

But how do these rhetorical relations affect the CCP of SDRSs? Well, unlike other non-logical predicates, rhetorical relations are assigned a truth definition with the capacity to *change* the world assignment pair. Moreover, we say that a relation $R(\alpha, \beta)$ is *veridical* iff $(w, f)P_M(R(\alpha, \beta))(w', g)$ entails $(w, f)P_M(K_\alpha)(w', g)$ and $(w, f)P_M(K_\beta)(w', g)$, where α labels the content K_α and β labels K_β. In other words, $R(\alpha, \beta)$ is veridical if it entails K_α and K_β. Similarly, R is left-veridical iff

$(w,f)P_M(R(\alpha,\beta))(w',g)$ entails $(w,f)P_M(K_\alpha)(w',g)$, and right-veridical if $(w,f)P_M(R(\alpha,\beta))(w',g)$ entails $(w,f)P_M(K_\beta)(w',g)$.

Now, *Narration* is veridical; in fact, its CCP satisfies the content of its arguments *in dynamic sequence*, as defined by the connective ';'. So its truth definition in SDRT is as follows:

- `Semantics of Narration`
  $(w,f)P_M(Narration(\pi_1,\pi_2))(w',g)$ iff:

  1. $\langle f(\pi_1), f(\pi_2)\rangle \in I_M(Narration)(w)$ and
  2. $(w,f)P_M(K_\alpha; K_\beta)(w',g)$

This semantics ensures that the imperatives in (1.1a′) are commanded, for thanks to clause 2., the CCP of (1.1a′) transforms the input world w into an output world w' where the actions have been performed (in sequence). Similarly, clause 2. also ensures that (1.1c′) is true only if the propositions expressed by the indicative clauses are true.

There are also *meaning postulates* on *Narration* which capture its spatio-temporal effects:[3]

- `Axiom on Narration`
  $Narration(\alpha,\beta) \wedge actor(x,e_\alpha) \wedge actor(x,e_\beta) \rightarrow$
  $\quad overlap(loc(x, prestate(e_\beta)), loc(x, poststate(e_\alpha)))$

In words, this stipulates that an actor x that is a participant in both events is in the same place, in space and time, at the end of the first event and at the beginning of the second. So, in (1.1ac), the agent is in the same place once he's finished going to the office as he is when he starts to get the file; i.e., he must start to get the file in Fred's office. And therefore, the file must be in Fred's office too, thanks to the lexical semantics of *get*.

One can think of such meaning postulates as constraining the admissible models in SDRT. More formally, consider (1.1a′). Because of the *Narration* relation, the CCP of this SDRS relates the world assignment pair (w,f) to (w',g) only if (w',g) verifies that (a) both actions have been performed, such that (b) $e_1 \prec e_2$ (i.e., e_1 preceded e_2), and (c) the red file folder is in Fred's office at the time when you get it. Similar constraints are imposed on the CCP of (1.1c′) by the very same axioms. These constraints on rhetorical relations thus account for implicatures.

This illustrates how in SDRT, implicatures are computed as a byproduct of computing discourse update: If one infers that a particular rhetorical relation must be used to connect the content of the current clause to the discourse context (we'll outline shortly how one does this), and if neither the content of that context nor the compositional semantics of the current clause verify the consequences of the rhetorical relation's meaning postulates (e.g., the spatio-temporal information of *Narration*), then this content is in essence accommodated, for it constrains the CCP of the updated SDRS. So SDRT predicts implicatures which are brought about by the demands of discourse coherence (i.e., the demand that we connect every bit of information in the discourse to some other bit of information with a rhetorical relation).

Computing implicatures via discourse update has two desirable consequences. First, it means that we can go some way towards axiomatising inferences about the value of ϕ in the formula (1.6), which we suggested earlier as a basis for bypassing Ross' paradox. Inferences about ϕ will essentially be part and parcel of inferences about the rhetorical relations that hold and their semantic effects. For example in discourse (1.1a), the implicature that the file is in Fred's office would be part of the contextually specified postconditions ϕ of the action. Thus rhetorical relations provide a first step towards avoiding Ross' paradox without sacrificing logical relationships among actions in the way that Segerberg does.

The second desirable consequence is that using the same rhetorical relation in the logical forms of (1.1a) and (1.1c) helps to explain why they have similar (spatio-temporal) implicatures. The discourses (1.1b) and (1.1d) are also similar: `Axiom on Narration` entails that the file is refiled in Fred's office. The uniform analysis of these discourses actually goes further than this: the logical forms of (1.1a) and (1.1c) are *constructed* in the same way as well. To see this, consider the way

[3] In fact, this axiom is stated here in slightly simplified form, because it ignores the role of frame adverbials.

in which SDRSs are constructed in SDRT. This is done within a glue logic, which consists of default axioms for inferring which rhetorical relation one uses to attach the new information to the logical form of the discourse context that's been constructed so far (see Asher and Lascarides (1995) for details). These default axioms encapsulate how a variety of knowledge sources provide clues about which rhetorical relations holds. So the axioms feature a default connective: $A > B$ means *If A then normally B*. The general schema for the axioms is given in (1.11): $\langle\tau,\alpha,\beta\rangle$ means that β (which labels an (S)DRS) is to be attached to a label α with a rhetorical relation, where α is part of the SDRS τ which represents the discourse context so far; and $\mathit{Info}(\tau,\beta)$ is a gloss for formulae that tell us properties of τ and β.

(1.11) $\quad (\langle\tau,\alpha,\beta\rangle \wedge \mathit{Info}(\tau,\beta)) > R(\alpha,\beta)$

Lexical semantics, domain knowledge and maxims of conversation essentially instantiate rules like (1.11). But typically, the rule itself has an antecedent which contains information that's derivable from the SDRSs that τ, α and β label. In other words, even if the justification of the rule resides in, for example, the model of discourse participants as rational and cooperative agents, the rule itself may appeal only to *linguistic* information in the antecedent. We will see an example of such a rule in section 4.2.

The axiom for inferring *Narration* is treated as a 'basic' default in Asher and Lascarides (1995), and it captures aspects of Grice's (1975) Maxim of Manner (i.e., *be orderly*):

- `Narration`: $\langle\tau,\alpha,\beta\rangle > \mathit{Narration}(\alpha,\beta)$

This default axiom together with `Axiom on Narration` stipulates that by default, people describe things in the order in which they occur, or are to occur.

Note that this default rule is neutral with respect to sentence mood. In particular, it applies when attempting to construct the logical forms of both (1.1a) and (1.1c). And in both cases, the consequent of the rule is consistent with the monotonic information that's available. So the underlying logic for $>$ yields the inference that *Narration* holds. Hence not only do the SDRSs (1.1a$'$) and (1.1c$'$) capture the implicatures of (1.1a) and (1.1c) in a uniform way, but also, in spite of the different sentence moods, the *way* in which these logical forms are constructed is uniform.

This contrasts with the plan-recognition approach to discourse interpretation (e.g., Grosz and Sidner 1990, Litman and Allen 1990). These theories reason about the way new information updates the meaning of the discourse by reasoning about how the communicative intention of the current utterance relates to the communicative intentions of the prior utterances. The communicative intentions that are conventionally associated (by default) with indicatives vs. imperatives are quite different. And so it's unclear how these theories could use the same axioms and proofs to explain their interpretations.

This semantic uniformity of imperatives vs. indicatives extends to other rhetorical relations as well. For example, *if...then* is a monotonic linguistic clue that the clauses are connected with the rhetorical relation *Condition*. So the axiom that encapsulates this will introduce *Condition* into the logical forms of both (1.12a) and (1.12b):

(1.12) a. If Ewan's in his office, then tell Johan the meeting is at 2pm.

b. If Ewan was in his office, then John told Johan that the meeting was at 2pm.

Unlike *Narration*, *Condition* isn't veridical. Rather, the following holds:[4]

$$\mathit{Condition}(\alpha,\beta) \rightarrow (K_\alpha \rightarrow K_\beta)$$

So *Condition* correctly predicts that (1.12a) is a conditional command, and it also conveys the correct semantics of the indicative discourse (1.12b).

The rhetorical relation *Elaboration* can also account for the semantics of the imperatives in (1.10a) and the indicatives in (1.10b): i.e., its semantics ensures that the action of taking documents with you is *part of* the action of going to the meeting.

[4] In fact, *Condition* is the SDRS equivalent of $\Rightarrow$ in DRT.

(1.10) a. Go to Fred's office. Take the file with you.

b. John went to Fred's office. He took the file with him.

This is captured in `Axiom on Elaboration`:

- `Axiom on Elaboration`:
  (a) $\mathit{Elaboration}(\alpha, \beta) \rightarrow K_\alpha \sqcap K_\beta$
  (b) $\mathit{Elaboration}(\alpha, \beta) \rightarrow e_\alpha \subseteq e_\beta$

This stipulates that the CCP of $\mathit{Elaboration}(\alpha, \beta)$ includes the *intersection* of the CCPs of the constituents that α and β label (hence *Elaboration* is veridical, and imperatives connected with *Elaboration* are commanded). Furthermore, the events are in a part-of relation (and so $\mathit{Elaboration}(\alpha, \beta)$ and $\mathit{Narration}(\alpha, \beta)$ are mutually inconsistent). So representing (1.10ab) with *Elaboration* captures the desired implicatures, and makes the temporal properties distinct from (1.1).

One doesn't infer *Narration* (via `Narration`) for connecting the constituents in (1.10), because a more specific conflicting default axiom applies in the glue logic, namely `Elaboration`:

- `Elaboration:` $(\langle \tau, \alpha, \beta \rangle \wedge \mathit{part\text{-}of}_D(\alpha, \beta)) > \mathit{Elaboration}(\alpha, \beta)$

In words, `Elaboration` states that if you're connecting β to α, and there's evidence within the discourse that they're in a part-of relationship, then normally the rhetorical connection is *Elaboration*. Discourse evidence of a part-of relation is typically modelled via *monotonic* rules which feature *linguistic* information about the constituents in the antecedent; i.e., they're axioms of the form $\mathit{Info}(\alpha, \beta) \rightarrow \mathit{part\text{-}of}_D(\alpha, \beta)$ (see Asher and Lascarides (1995) for details). In (1.10), the monotonic rule which applies instantiates $\mathit{Info}(\alpha, \beta)$ with the information that α describes movement, and β also describes a causative movement performed by the same agent. Note that this is neutral with respect to sentence mood, and so the SDRSs for (1.10ab) are constructed via the same glue logic axioms.

4.1 Defeasible Conditionals and Metatalk Relations

The rhetorical relations we've considered so far constrain the *contents* of the constituents they connect. Moore and Pollack (1992) argue convincingly that rhetorical relations can also reveal information about intentions and speech acts. They use (1.13) to observe that recognising a content-level relation is sometimes necessary for recognising the intentional one and *vice versa*:

(1.13) a. Come home by 5pm.

b. Then we can go to the hardware store before it closes.

c. That way, we can finish the shelves tonight.

At the content level, the clauses in (1.13) describe events which are in *consequence* relations: doing the action described in (1.13a) normally results in (1.13b) being true; and (1.13b) being true normally means (1.13c) is true too. Elsewhere we have used the non-veridical relation $\mathit{Def\text{-}Consequence}(\alpha, \beta)$ to mark this connection between propositions (e.g., Asher and Lascarides, 1998b):

- `Axiom on Def-Consequence`
  $\mathit{Def\text{-}Consequence}(\alpha, \beta) \rightarrow (K_\alpha > K_\beta)$

When the first constituent is a request, however, the corresponding action term cannot be an antecedent to $>$ *directly*, because it's not of the right semantic type. Rather, where α labels $\delta K'_\alpha$, the appropriate consequence relation is $[\delta K'_\alpha]\top > K_\beta$. Or, in words, any situation where the action described by the imperative α is performed is normally one where the proposition β is true as well. We encode this content-level relationship in the rhetorical relation $\mathit{Def\text{-}Consequence}_r$ (α :! means that α labels an imperative and β : . means that β labels an indicative):

- `Axiom on Def-Consequence`$_r$
  - (a) $\textit{Def-Consequence}_r(\alpha,\beta) \rightarrow (\alpha :! \wedge \beta : .)$
  - (b) $(\textit{Def-Consequence}_r(\alpha,\beta) \wedge \alpha :!\delta K'_\alpha) \rightarrow [\delta K'_\alpha](\top > K_\beta)$

Like *Narration*, *Def-Consequence*$_r$ encodes information about what results from doing the action described by α. But unlike *Narration*, *Def-Consequence*$_r$ is not veridical: the imperative is not commanded since the CCP of the SDRS will not have the characteristic transformation of the world index.

Def-Consequence$_r$ is part of the semantic representation of the discourses (1.2), (1.13) and (1.14):

(1.2) a. Go straight on and you'll fall down the stairs.

b. Come home by 5pm and we can go to the hardware store before it closes.

(1.14) Smoke 20 cigarettes a day and you will die before you're 30.

Discourse (1.14) doesn't implicate that the imperative *smoke 20 cigarettes a day* is commanded, largely because the consequent state (death) is undesirable. Similarly for (1.2a).

However, in contrast to (1.14) and (1.2a), the request is commanded in (1.13) and (1.2b). As Moore and Pollack (1992) point out informally, the consequence relations at the 'content-level' in (1.13), together with the background knowledge that (1.13c) is a desirable state, yield further inferences: these consequence relations *explain why* the speaker made the request. This is an example of what Polanyi (1985) calls a 'meta-talk' relation, for it connects the content of one utterance to the performance of uttering another. In words, the meta-talk relation *Explanation**(α,β) means that β explains why *Agent*(α) (i.e., the person who uttered α) performed the speech act α (e.g., *Explanation** is part of the representation of *Close the window. I'm cold*). Like *Narration* and *Elaboration*, it's a veridical relation. So (1.13) is represented as (1.13′) (for simplicity, we haven't stipulated the action terms and DRSs that represent the clauses):

(1.13′)
$$\begin{array}{|l|}\hline \pi_1, \pi' \\ \hline \pi_1\text{:[Come home by 5pm]} \\ \\ \pi' : \begin{array}{|l|}\hline \pi'_1, \pi_2, \pi_3 \\ \hline \pi'_1\text{:[Come home by 5pm]} \\ \pi_2\text{:[We can go to store]} \\ \pi_3\text{:[We can finish shelves]} \\ \textit{Def-Consequence}_r(\pi'_1, \pi_2) \\ \textit{Def-Consequence}(\pi_2, \pi_3) \\ \hline \end{array} \\ \\ \textit{Explanation*}(\pi_1, \pi') \\ \hline \end{array}$$

In words, (1.13′) stipulates that the following explains why the request (1.13a) is commanded: doing the action described by (1.13a) normally leads to being able to go to the hardware store before it closes, which in turn normally leads to being able to finish the bookshelves tonight. Note that (1.13′) represents the content of the imperative at two levels in the discourse structure: it's the first argument in an *Explanation** relation; and its also part of the representation of the second argument π' to this relation. This reflects the fact that rhetorically, the imperative plays a 'dual role': its content and its default consequences motivate its own command status. Since *Explanation** is veridical, the request expressed by (1.13a) is commanded. The SDRS labelled π' must be true as well; but because *Def-Consequence isn't* veridical, this doesn't mean that the propositions expressed by (1.13b) or (1.13c) are true.

The difference between (1.13) and (1.14) is (1.14) does not feature the veridical *Explanation** relation, but only the *Def-Consequence*$_r$ one. So the imperative in (1.14) isn't commanded. There's a similar difference between (1.2a) and (1.2b). But how can one infer these differences in the glue logic for constructing logical forms? Well, the SDRS (1.13′) can be inferred via monotonic axioms which take account of the cue phrases *that way*, *then* and punctuation. Similarly, when the cue phrase *and* connects an imperative to an indicative, it monotonically generates $Def\text{-}Consequence_r$; this applies to (1.2ab) and (1.14). Finally, *Explanation** would be inferred via a default axiom which states: if α is a request, $Def\text{-}Consequence_r(\alpha, \beta)$ holds, and β is a desirable state, then normally $Explanation^*(\alpha', \pi)$ holds, where α' labels a 'repeat' of the content of the request α, and π labels α's and β's content and the $Def\text{-}Consequence_r$ relation between them.

4.2 Imperative Answers

We suggested earlier that an imperative that's an (indirect) answer to a question isn't necessarily commanded. We follow the SDRT analysis of indirect answers from Asher and Lascarides (1998a), using the relation *IQAP* (standing for Indirect Question Answer Pair) to represent the connection between a question and its indirect answers. Semantically, $IQAP(\alpha, \beta)$ holds only if α is a question, and the speaker of α can infer a direct answer to his question (according to the compositional semantics of questions and answers) from the content that's labelled by β. This relation *IQAP* will feature in the logical form of (1.3), for it reflects the fact that an adverbial of manner, which constitutes the semantic type of direct answers to *how*-questions, can be computed from the contents of the imperatives.

(1.3) a. A: How does one make lasagne?

b. B: Chop onions and fry with mince and tomatoes, boil the pasta, make a cheese sauce, assemble it, and bake in the oven for 30 minutes.

Now, we must encode within the truth definition of $IQAP(\alpha, \beta)$ that imperative answers aren't always commanded: i.e., $IQAP(\alpha, \beta) \rightarrow K_\beta$ isn't valid when K_β is a request. However, this is in contrast to *IQAP* when it relates propositions, for these are right-veridical. I.e., $IQAP(\alpha, \beta) \rightarrow K_\beta$ *is* valid when K_β is a proposition (see Asher and Lascarides, 1998a).

In fact, whether or not $IQAP(\alpha, \beta)$ makes the imperative β commanded depends on the compositional semantics of the question. The imperatives in (1.3) aren't commanded. But if the question is about what plans should be on the questioner's agenda, the imperative answer does seem to be commanded (e.g., (1.15) and (1.16)):[5]

(1.15) a. A: Where do I go now?

b. B: Go to platform 1.

(1.16) a. A: What should I do now?

b. B: Own up to the police.

We need to reflect this in the semantics for *IQAP*. First, we must have a monotonic axiom which makes $IQAP(\alpha, \beta)$ right-veridical when β is a proposition:

- `Veridicality of Propositional Answers:`
  $(IQAP(\alpha, \beta) \wedge \beta : .) \rightarrow K_\beta$

Second, when β is an imperative, β is commanded only if it is an indirect answer to a question whose *direct* answers implicate a deontic modality on the questioner. We can represent such implications as a $>$-statement. And so the axiom below captures the required information (where *QAP* stands for Question Answer Pair, and $QAP(\alpha, \beta)$ means that β is a true direct answer to the question α, according to the compositional semantics of questions and answers):

- `Veridicality of Imperative Answers: Deontic`
  $(IQAP(\alpha, \beta) \wedge \beta :! \wedge (QAP(\alpha, \gamma) > [deontic]_{Agent(\alpha)}(\phi))) \rightarrow K_\beta$

[5] Actually, it's not clear whether or not (1.16) is a fragment answer where its message type isn't imperative at all, and thus not commanded.

So, the logical form of (1.3) is (1.3′):[6]

$$
(1.3')\qquad
\begin{array}{|l|}
\hline
\pi_1, \pi \\
\hline
\pi_1\text{:[How make lasagne?]} \\
\pi : \begin{array}{|l|}
\hline
\pi_2, \pi_3, \pi_4, \pi_5, \pi_6 \\
\hline
\pi_2\text{:[Chop onions]},\ \pi_3\text{:[fry with mince]} \\
\pi_4\text{:[boil pasta]},\ \pi_5\text{:[make cheese sauce]} \\
\pi_6\text{:[assemble it]},\ \pi_7\text{:[bake]} \\
Narration(\pi_2, \pi_3),\ Narration(\pi_3, \pi_4),\ Narration(\pi_4, \pi_5), \\
Narration(\pi_5, \pi_6),\ Narration(\pi_6, \pi_7) \\
\hline
\end{array} \\
IQAP(\pi_1, \pi) \\
\hline
\end{array}
$$

Since *IQAP* isn't right-veridical, the imperatives in (1.3b) aren't commanded. However, the *Narration* relations ensure that the complex action described in π has the desired temporal properties; e.g., the individual actions would be performed in the order they were uttered.

Note that all direct answers to (1.16a) are propositions that the the questioner should ϕ, for some value of ϕ. So the semantics of *IQAP* correctly predicts that the imperative answer (1.16b) is commanded. Direct answers to the question (1.15a) don't *semantically* entail a deontic proposition, but they do implicate one. Assuming that this implicature is captured in SDRT, the imperative (1.15b) is commanded according to the above semantics of *IQAP*. *IQAP* will also deal adequately with the command-status of an imperative to a conditional question:

(1.17) a. A: If the exam is tomorrow, then what should I do?

b. B: Revise your notes tonight!

That is, it correctly predicts that B's imperative is a conditional command; conditional on whether the exam is tomorrow.

The glue logic axiom for inferring *IQAP* exploits Morgan's (1975) notion of *short-circuiting* calculable implicatures. In Lascarides and Asher (1999), we demonstrate that a logical model of discourse participants as rational and cooperative agents validates the following: if β is a response to a question α, then normally $IQAP(\alpha, \beta)$ holds; i.e., β supplies sufficient information that the questioner can infer a direct answer to α from it. This is represented in the glue logic via the following axiom:

- `IQAP`: $(\langle \tau, \alpha, \beta \rangle \wedge \alpha :?) > IQAP(\alpha, \beta)$

Note that although the *justification* for this axiom involves inferences that are derived from reasoning about the beliefs and intentions of the dialogue participants, the rule `IQAP` itself has an antecedent which includes only information about the sentence moods; something that is given by the grammar. So in essence, `IQAP` short-circuits the calculable inferences about when the speech act of providing an indirect answer is performed, because it allows the interpreter to entirely bypass reasoning with cognitive states, using just the sentence mood of α instead. This axiom plays a central role in constructing the SDRSs for (1.3) and (1.15–1.17).

4.3 Corrections

Consider (1.4), where intuitively A's second imperative 'replaces' the first one:

(1.4) a. A: Go straight for 5cm.

b. B: That will take me straight into the crevasse

c. A: Go left then.

[6]The axiom `Narration` is used to connect the imperatives with *Narration*, and the axiom `IQAP` which we'll specify shortly is used to connect the segment of imperatives to the question. We forego giving details here, however, of how one chooses which labels are connected to which other labels (but see Asher and Lascarides (1998a, 1998b)).

We must model how rhetorical relations can yield such non-incremental interpretations: whereas the logical form for the discourse context entails that an imperative is commanded, the logical form of the updated discourse context cancels this entailment.

Commanding the imperative (1.4a) is *not* incompatible with its (undesirable) outcome (1.4b). And yet (1.4b) functions as a corrective move, since it brings into dispute that the imperative is commanded (or, more accurately, that it should be commanded). Now, in earlier work (e.g., Asher and Lascarides, 2001), we have used the relation *Plan-Correction* to model this: *Plan-Correction*(α, β) holds if β indicates that the goal which lay behind uttering α is incompatible with *Agent*(β)'s goals. This is analogous to Searle's speech act of rejection: it features in the analysis of (1.4) and (1.18) (taken from Searle, 1969).

(1.18) a. A: Let's go to the movies tonight.

b. B: I have to study for an exam.

Plan-Correction(α, β) is right-veridical but not left-veridical, thereby providing the non-incremental interpretation we require: if α labels an imperative, then an SDRS that contains *Plan-Correction*(α, β) does *not* have a CCP with the characteristic transformation of the world index, indicating that this SDRS is one where the imperative is not commanded. This then leaves A free to issue a further command in response to the *Plan-Correction*. In fact, A requests (1.4c) as a *result* of B's utterance (1.4b), as indicated by the cue phrase *then*. But this result relation is at the speech-act level (i.e., it's a metatalk relation, connecting the content of B's utterance to A's speech act of uttering the request (1.4c)). And so (1.4) is represented as (1.4′), where *Result** encodes the appropriate metatalk relation:

(1.4′)

π_1, π_2, π_3
π_1:[Go straight for 5cm] π_2:[π_1 takes me into crevasse] *Plan-Correction*(π_1, π_2) π_3:[Go left then] *Result**(π_2, π_3)

Like *Explanation**, *Result** is veridical. Therefore, (1.4′) entails that the imperative (1.4c) is commanded, but it does not entail that (1.4a) is commanded.

5 Some Concluding Remarks

We have examined the content of imperatives in dialogue, paying particular attention to their compositional semantics and to the ways in which the discourse context affects their interpretation. We argued that dynamic semantics provides an elegant account of their compositional semantics based on action terms: the defining characteristic of a discourse which features a commanded imperative is that its context change potential (CCP) transforms the world index into one where the action has been performed.

We observed that context can affect whether imperative is commanded, and also the content of the action term—for example, the time at which the action is to be carried out. We argued that these contextual effects are best explained through capturing the *rhetorical role* of the imperative. Indeed, representing the content of discourse in terms of the rhetorical connections between the propositions and requests partly contributes to the simplicity of the compositional semantics of the imperatives, in that it makes a semantic primitive for stipulating what's commanded unnecessary, in contrast to Segerberg's analysis. The command status of an imperative is instead represented via the veridicality of its rhetorical connection to the rest of the dialogue. And since this rhetorical relation is inferred on the basis of both linguistic and non-linguistic information, SDRT provides a framework in which the information flow between the content of an imperative, domain knowledge and goals can be modelled.

We also observed similarities in the implicatures of imperatives and indicatives. SDRT is distinct from plan-recognition approaches to discourse interpretation, in that the rhetorical relations are inferred on the basis of axioms which have direct access to the linguistic form of the context. We argued that this allows for a uniform analysis of these implicatures which would be hard to achieve through plan recognition.

There are many outstanding issues. For instance, we need to examine more closely the semantic relationships between imperatives and adverbials of manner; compare (1.19a) and the semantically similar (1.19b):

(1.19) a. Go to the kitchen and make a cup of coffee.

b. Go to the kitchen to make a cup of coffee.

The interaction between imperatives, presuppositions and anaphora also deserves closer attention, as does the link between interpreting imperatives and planning (see, for example, Stone (in press)). We will examine these issues in future work.

Acknowledgements

This paper could not have been written if it weren't for extensive discussions with Nicholas Asher, Johan Bos and Ewan Klein. Nicholas Asher pointed out to me that a dynamic semantics makes Segerberg's command system redundant. Much of the work presented here will also appear in Asher and Lascarides (forthcoming). Alex Lascarides is supported by an ESRC (UK) research fellowship.

References

Asher, N. [1993] *Reference to Abstract Objects in Discourse*, Kluwer Academic Publishers.

Asher, N. and A. Lascarides [1995] Lexical Disambiguation in a Discourse Context, *Journal of Semantics*, **12.1**, pp69–108, Oxford University Press.

Asher, N. and A. Lascarides [1998a] Questions in Dialogue, *Linguistics and Philosophy*, **23.3**, pp237–309, Kluwer Academic Publishers.

Asher, N. and A. Lascarides [1998b] The Semantics and Pragmatics of Presupposition, *Journal of Semantics*, **15**, pp239–299, Oxford University Press.

Asher, N. and A. Lascarides [2001] Indirect Speech Acts, *Synthese*, **128.1**.

Asher, N. and A. Lascarides [forthcoming] *The Logic of Conversation*, Cambridge University Press.

Cohen, P. and Levesque, H. [1990] Rational Interaction as the Basis for Communication, in in P. R. Cohen, J. Morgan and M. Pollack (eds.) *Intentions in Communication*, pp.221–255, MIT Press.

Grice, H. P [1975] Logic and Conversation, in Cole, Peter and Morgan, Jerry L. (eds.) *Syntax and Semantics Volume 3: Speech Acts*, pp41-58. Academic Press.

Grosz, B. and Sidner, C. [1986] Attention, Intentions and the Structure of Discourse. *Computational Linguistics*, **12**, 175–204.

Grosz, B. and Sidner, C. [1990] Plans for Discourse, in P. R. Cohen, J. Morgan and M. Pollack (eds.) *Intentions in Communication*, pp.417–444, MIT Press.

Hare, R. M. [1967] Some Alleged Differences between Imperatives and Indicatives, *Mind.*

Hobbs, J. R., M. Stickel, D. Appelt and P. Martin [1993] Interpretation as Abduction, *Artificial Intelligence*, **63 (1–2)**, pp.69–142.

Kamp, H. and Reyle, U. [1993] *From Discourse to Logic*, Kluwer Academic Publishers.

Lascarides, A. and N. Asher [1993] Temporal Interpretation, Discourse Relations and Commonsense Entailment, in *Linguistics and Philosophy*, **16**, pp.437–493.

Lascarides, A. and N. Asher [1999] Cognitive States, Discourse Structure and the Content of Dialogue, in *Proceedings from Amstelogue 1999*, pp1–12, Amsterdam, May 4th–7th 1999.

Lascarides, A. and A. Copestake [1998] Pragmatics and Word Meaning, *Journal of Linguistics*, volume **34.2**, pp387–414, Cambridge University Press.

Litman, D. and J. Allen [1990] Discourse Processing and Commonsense Plans, in P. R. Cohen, J. Morgan and M. Pollack (eds.) *Intentions in Communication*, pp.365–388, MIT Press.

Lochbaum, K. [1998] A Collaborative Planning Model of Intentional Structure, *Computational Linguistics*, **24.4**, pp525–572.

Mann, W. and S. Thompson [1987] Rhetorical Structure Theory: A Framework for the Analysis of Texts. In *IPRA Papers in Pragmatics*, **1**, 79–105.

Moore, J. D. and M. Pollack [1992] A Problem for RST: The Need for Multi-Level Discourse Analysis *Computational Linguistics*, **18** 4, pp537–544. Morgan, J. L. [1975] Some Interactions of Syntax and Pragmatics, in Peter Cole and Jerry L. Morgan (eds.), *Syntax and Semantics Volume 3: Speech Acts*, Academic Press, pp289–303.

Polanyi, L. [1985] A Theory of Discourse Structure and Discourse Coherence, in Eilfor, W. H., Kroeber, P. D., and Peterson, K. L., (eds), *Papers from the General Session a the Twenty-First Regional Meeting of the Chicago Linguistics Society*, Chicago, April 25–27, 1985.

Ross, A. [1941] Imperatives and Logic, *Theoria*, vol 1941, pp53–71.

Searle, J. [1969] *Speech Acts*, Cambridge University Press.

Segerberg, K [1990] Validity and Satisfaction in Imperative Logic, *Notre Dame Journal of Formal Logic*, **31**, pp203–221.

Stone, M. [in press] Towards a Computational Account of Knowledge, Action and Inference in Instructions, *Language and Computation*.

Webber, B., N. Badler, B. DiEugenio, C. Geib, L. Levinson, M. Moore [1995] Instruction, Intentions and Expectations, *Artificial Intelligence*, **73**, pp253–269.