In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

To appear in: Bunt, H & R.J. Beun, Selected Papers from the Proceedings of the Second International Conference on Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence, Berlin: Springer-Verlag.

Relating Imperatives to Action

Paul Piwek

ITRI – Information Technology Research Institute, Brighton
Email: Paul.Piwek@itri.bton.ac.uk

November 20, 2000

Abstract. The aim of this paper is to provide an analysis of the use of logically complex imperatives, in particular, imperatives of the form 'Do A_1 or A_2' and 'Do A, if B'. I argue for an analysis of imperatives in terms of Classical Logic which takes into account the role of background information in the use imperatives. I show that by doing so one can avoid some counter-intuitive results which have been associated with analyses of imperatives in terms of Classical Logic.

1 Introduction

Imperatives[1] are typically used to get an addressee to do something. As a result, a key notion in the analysis of imperatives is the notion of the *satisfaction of an imperative* (by its addressee). Understanding the meaning an imperative involves knowing under which circumstances it has been satisfied. Various authors, notably Stenius (1967) and Searle (1969), have proposed that we analyse this notion of *satisfaction* by assuming that an imperative conveys both a proposition and the instruction to make this proposition true. For instance, the imperative 'Get off the table' conveys the proposition *that the addressee gets off the table* and the instruction that *the addressee makes it true that s/he gets off the table*. Whether or not an imperative has been satisfied can now be understood in

[1] I will use the word 'imperative' as a short-hand for 'the utterance of an imperative clause'. Imperative clauses are traditionally characterised by the lack of a subject, the use of a base form of the verb, and the absence of modals as well as tense and aspect markers. Examples of naturally occurring imperatives are: 'Get off the table', 'Don't forget about the deposit' and 'Hold on, are we late?'. The definition and the examples are taken from Biber et al. (1999:219).

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

terms of whether its propositional content has been made true. Although this account of imperatives sounds plausible, it has been shown to give rise to some apparently counter-intuitive results when it is formalised by means of Classical (Predicate) Logic. In this paper, I discuss a number of these results and demonstrate how they can be avoided by embedding the aforementioned type of analysis in an Agent-based model which takes into account the influence of background information on the use of imperatives.

2 The Puzzle of Logically Complex Imperatives

Let me indicate how a formalisation of the aforementioned analysis of imperatives appears to give rise to counter-intuitive results. In such a formalisation, the propositional content of imperatives is represented by means of Predicate Logical sentences. More precisely, there is a function f_C which maps an imperative clause, given some utterance context C, to a Predicate Logical representation of its propositional content. In particular, the natural language connective 'or' is rendered by means of the logical disjunction '$\vee$' and the natural language construction '..., if ...' is represented by means of the material conditional '$\rightarrow$':

(1) $f_C(A_1 \textit{ or } A_2) = f_C(A_1) \vee f_C(A_2)$.

(2) $f_C(A, \textit{if } B) = f_C(B) \rightarrow f_C(A)$.

The notion of the satisfaction of an imperative can now be defined in terms of the truth of the propositional content of the imperative in a model M of the actual world:

(3) SATISFACTION
Imperative **I**, uttered in context C to addressee A, has been *satisfied* in a world represented by model M, if and only if A makes it the case that $f_C(\mathbf{I})$ is true in the world represented by M.

In words, an imperative has been satisfied if and only if its addressee made its content true. A first possible objection to this definition is that the mere fact that the addressee made the content true is to weak, especially, if we read 'satisfied' as 'complied with': consider a situation in which a person is told to get off a table and subsequently accidentally stumbles and falls off that table. In that situation, the propositional content of the imperative has been made true. Nevertheless, one would not want to claim that the person really complied with the imperative: his or her action was not influenced in the *appropriate way* by the utterance of the imperative clause.[2] The problem might be resolved by

[2] Suppose that the addressee of the imperative was a person named 'John' and the issuer a person named 'Mary'. In that situation, the following discourse would be, if not an infelicitous, then certainly a misleading description of what took place: 'Mary told John to get of the table and he complied'. Note that 'Mary told John to get of the table. He did not comply.' seems equally infelicitous/misleading, which suggests that '**A** complied with imperative **I**' presupposes (rather than that it asserts) that **A**'s action was influenced in the *appropriate way* by the utterance of **I**.

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

making a clear distinction between the notion of *satisfaction* of an imperative, and the notion of an addressee *complying with* an imperative. For now, let us look at another problem which casts a shadow over definition (3). Take a rule-like imperative such as:

(4) Make a cross, if you encounter a vampire.

In line with (2), the propositional content of this imperative is rendered by the following Predicate Logical sentence, assuming that the addressee's name is John:

(5) $\mathit{encounter\_vampire(john)} \rightarrow \mathit{make\_cross(john)}$

In words, if the addressee encounters a vampire, then he makes a cross.[3] Note that in Classical Logic the formulae of the form $p \rightarrow q$ and $\neg p \vee q$ (i.e., it is not the case that p or it is the case that q) express the same proposition: whenever one of them is true/false, the other one is as well. Thus, the propositional content of our imperative can also be expressed by the formula in (6).

(6) $\neg \mathit{encounter\_vampire(john)} \vee \mathit{make\_cross(john)}$

By combining this with the definition (3) we obtain: the imperative in (4) has been satisfied by John if and only if John makes least one of the following two formulae true:

(7) a. $\neg \mathit{encounter\_vampire(john)}$;
b. $\mathit{make\_cross(john)}$.

According to this analysis, one possible way of satisfying the imperative would be for the addressee to simply avoid encountering a vampire (cf. Hamblin, 1987:84). In some sense, it is correct that in that case the imperative would be satisfied (and even complied with), as predicted by definition (3). Still, I am left with the feeling (and hopefully the reader to) that the definition provides us with a very incomplete (and potentially misleading) analysis of the relation between imperatives and the actions which they give rise to. Let me explain. Imagine that the addressee of the imperative is interviewed and asked why he was avoiding vampires. What could the addressee reply to this question? Would it, for instance, be right for him to say: 'I avoided vampires because I was told that if I see one, I should make a cross'? To me, this answer sounds odd; but why, and can we draw any conclusions from this? One important thing to note is that there do exist situations in which the response is felicitous: it makes sense if we assume that the addressee for some reason or other wants to avoid making a cross. An addressee who is already committed to not making a cross, would indeed need to avoid vampires in order to comply with the imperative (4).

Generally speaking, we may conclude that in order to understand how imperatives influence the behaviour of their addressees we need to take into account possible background information which also influences the actions of an

[3] For the moment, I am omitting an explicit representation of temporal reference.

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

addressee. In particular, we need to take into consideration further propositions to which the addressee is committed. This means that although definition (3) provides us with some idea of the meaning of imperatives (in terms of their conditions of satisfaction), it fails to give us a full picture of how imperatives are *used* and (are expected to) *affect* the behaviour of their addressees.

3 Imperatives in Context: An Agent-based Analysis

In order to examine the influence of background information on the relation between imperatives and the actions which they give rise to, I am going to construct a model of communicating agents and their environment. The influence of background information can then be formalised in terms of constraints on the behaviour of the agents in this model. These constraints will be normative in nature: they describe how the agents are *expected* to behave when they have agreed to comply with an imperative. Additionally, I will describe a concrete policy for action which allows an agent to fulfill these expectations.

At the basis of the model –a formally more explicit account can be found in Piwek (2000)– lies the notion of a time line. This line consists of discrete instants of time. For practical purposes, the time line is restricted to a finite number of instants of time.[4] There are events and states, such that an event always occurs *between* two subsequent instants of time. Henceforth, I call such a pair of instants a transition. A state always holds *at* one or more instants of time. For example, consider a situation in which an agent switches off the light in a room between the instants of time t_0 and t_1. This means that a switch event e took place at the transition from t_0 to t_1. The event token e of the type switch occurs uniquely at this transition. In general, the same event (token) can never occur at several different transitions, although several different events can be associated with one and the same transition. Similar to Parsons (1990), I assume that an event can involve a number of objects: in this case, its actor (the agent who switches off the light) and its patient (the light which is switched off). Furthermore, there are two states s_1 and s_2: one in which the light is on and one in which the light is off. s_1 holds at $\ldots, t_{-2}, t_{-1}$ and t_0 whereas s_2 holds at $t_1, t_2, \ldots$ In general states are assumed to span a connected series of instants of time. In other words, if a particular state holds at two different instants of time t_x and t_y, then it also holds at all the instants between t_x and t_y. Additionally, a state has the same type and involves the same objects at all the instants at which it occurs.

In this model, agents are just another sort of objects. What is still missing is a counterpart to the fact that (human) agents can carry information; informa-

[4] I assume that the time line corresponds with the life span of a particular agent. This allows us to study the behaviour of that agent. A time line which is infinite in both directions might be more realistic, but it would also force us to explicitly model the life span of the agents in our models, which makes the analysis less concise without the benefit of any additional insights. The choice for a discrete as opposed to a continuous time line is another assumption which I make to keep the model as transparent as possible.

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

tion which guides their actions.[5] To remedy this I assume that at each instant of time an agent is associated with a commitment state. This commitment state consists of all the information which an agent is personally committed to (cf. Hamblin, 1971). What this precisely means is fleshed out in the remainder of this paper.[6] Basically, the commitment state puts constraints on the behaviour of the agent. When we look at the model from the outside, we can speak of the satisfaction conditions of the commitment state. I assume that as an external observer we have access to a function Int which maps commitment states (of individual agents) onto Predicate Logical sentences. The conditions of satisfaction of a commitment state correspond with the truth-conditions of the corresponding Predicate Logical sentence in the model (which includes the agent and its commitment state).[7] The reader might feel somewhat uncomfortable with the assumption that an observer has a function which returns the conditions of satisfactions of an agent's commitment state. It seems that we can only guess at such a function based on the behaviour of the agent, assuming that an agent strives for the satisfaction of his or her commitments. This is certainly correct when we consider human agents. However, the nice thing about building models is that we can simply define such a function for the agents with which we populate the model and then compare different policies for action with respect to this function. Thus we can determine how successful different policies are with respect to satisfying the agents' commitments. It will turn out that this perspective is very helpful for understanding the role of imperatives. Now that we have a model in place, let us examine how we can use it to analyse the use of an imperative. Take the imperative (4), here repeated as (8).

(8) Make a cross, if you encounter a vampire.

The utterance event u of this imperative is modelled as occurring between two instants of time, say t_0 and t_1.[8] This utterance event has certain physical prop-

[5] Cf. Beun, (1992), Bunt & Katwijk (1979), the collection of papers in Bunt and Black (2000), Bunt et al. (this volume), Lewis (1979) and Power (1979).

[6] But let me briefly indicate by means of an example what a commitment state amounts to. We can compare it to a shopping list. Such a list records what goods the owner of the list is committed to buying. The list might be carried about on a piece of paper, but alternatively, the person in question might have memorised it. Memorisation, of course, brings with it some problems, such as the fact that we can forget things. In this paper, I make the simplifying assumption that the commitment state behaves like a written record: the information in it is reliable. Of course, most of the commitments of human agents are not recorded in this way. Human agents mainly rely on their memory. It is, however, beyond the scope of this paper to take into account the complications which this entails.

[7] Technical details with respect to Int can be found in Piwek (2000). There, a partialized truth-conditional semantics is employed to avoid paradoxes such as 'The Liar'.

[8] For the sake of the argument I assume that an utterance event consists of just one single transition. In a more sophisticated analysis we might want to model an utterance as consisting of a sequence of events.

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

erties and some of these properties will be perceived and processed by the addressee of the utterance. In other words, in parallel with the utterance event, i.e., also between t_0 and t_1, there will be an observation event (by the addressee) of the utterance event. This observation event causes a change to the commitment state of the addressee. It is beyond the scope of this paper to provide a detailed analysis of how the perception affects the information state of the addressee (but see, for instance, Beun, 1994). Let us assume that the addressee perceives and understands the imperative and believes that the utterance was felicitous (e.g., the speaker was sincere. See Austin, 1962). The commitment state of the addressee at t_1 will then contain the information of the addressee's commitment state at t_0 extended with the information that it is the goal of the speaker that the addressee makes a cross if s/he encounters a vampire.[9] It is now up to the addressee to signal whether s/he is committed to satisfying the speaker's goal. Note that signalling commitment to an action is stronger than signalling that one intends to perform the action. Consider the following dialogue fragment:

(9) A: Finish the paper before tomorrow.
B: Ok, but I might fail to do it.

Here, B's utterance appears to be incoherent. B cannot first signal compliance with the imperative and then suggest that s/he might fail to comply. Now let us substitute 'I intend to' for 'ok' and reconsider the dialogue fragment:

(10) A: Finish the paper before tomorrow.
B: (Well) I intend to, but I might fail.

Now, there seems to be nothing wrong with B's reply. From this we can draw the conclusion that the meaning of 'ok' cannot be identical to 'I intend to'. Imperatives give rise to commitment, which is stronger than intention: when one is committed to carrying out a certain action, one cannot at the same time be committed to the idea that one might fail to carry out that action. Now let us return to (8). The propositional content of this imperative is added to the commitment state of the addressee. Formally, we have that if $Int(c_{t_1}) = \Phi$ then $Int(c_{t_2}) = \Phi \wedge f_C(8)$.[10,11] I assume that the propositional content $f_C(8)$ corresponds with the following formula, where $occur(e, t)$ means that event e took place between t and the instant of time which immediately succeeds t:

---

[9] In case the imperative is intended as advice, the speaker will have expressed the proposition that it is beneficial to the addressee to adhere to the propositional content of the imperative. In other words, the speaker is suggesting a goal to the addressee. In this paper, the focus is on imperatives which are used as commands: that is, imperatives which express a goal of the speaker or somebody else on whose behalf the speaker issues the imperative.

[10] Where c_{t_x} stands for the commitments of the addressee at t_x.

[11] This representation does not allow for anaphoric reference to indefinite antecedents. For that we would need a dynamic system such as Discourse Representation Theory, e.g., Kamp & Reyle (1993).

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

(11) $(\forall v\,t\,t'\,e.(vampire(v) \wedge encounter(john, v, e) \wedge occur(e, t) \wedge successor(t', t) \wedge after(t, t_{utt})) \rightarrow (\exists e'.make\_cross(john, e') \wedge occur(e', t')))$

In words, if the addressee –John– encounters a vampire at a time t after the utterance time of the imperative t_{utt}, then he makes a cross at the instant of time which immediately succeeds t. Now let us look at two possible constraints governing the role of imperatives. Firstly, one might consider an individual imperative and demand that its content is true in the world.

(12) CONSTRAINT: TRUTH-CONDITIONAL IMPERATIVE SATISFACTION
If $f_C(\mathbf{I})$ is the content of an imperative **I** uttered at some point in time, then $f_C(\mathbf{I})$ is true in the world.

Note that we omit the condition that this content should be made true by its addressee. We return to this issue in a moment. Secondly, we can look at the commitment states of agents. We have seen that the content of an imperative is added to the commitment state of the addressee of an imperative. Thus we can define an alternative notion of satisfaction which pertains to entire commitment states:

(13) CONSTRAINT: TRUTH-CONDITIONAL COMMITMENT STATE SATISFACTION
If c is the commitment state of an agent at some point in time, then $Int(c)$ is true in the world.

The addressee of an imperative is supposed to actively make sure that the constraint (12) is satisfied. An agent who manages to live up to this expectation is said to have *complied with* the imperative. However, an agent who tries to comply with an imperative cannot address that imperative in isolation from his or her other commitments. Consider, for instance, our agent John. Suppose he is committed to encountering a vampire, because he replied 'ok' to the imperative 'Find a vampire'. This commitment prevents John from satisfying (8) by avoiding vampires. The question then is how an agent can make sure that all his or her commitments are true in line with (13), if it is not possible to consider them one at a time. *Is there a policy which allows an agent to make sure that all his or her commitments are/become true?*

Let us first try to formulate a more precise version of this question. At the basis of this section lies the idea of a model of reality. The models which we have described up till now are total in nature: such a model determines completely which states and events take place at any given instant of time/transition and who and what are involved. With respect to such a model we can check whether constraints such as (12) and (13) are satisfied.[12] Now consider a model which

[12] In Piwek (2000) we also allow for partial models at this stage. In particular, models which do not contain any information beyond some instant of time. In other words, these models allow us to model that the future is under-determined. Such models seem more realistic, but are not essential for the point which I want to make in this paper.

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

is just as determined as the aforementioned models except for the fact that the actions of a particular agent, say John, are not specified in this model at any transition. Now, a policy provides us with a rule which takes such a partial model and turns it into a fully determined, i.e., total, one. Thus the policy determines which actions an agent does and does not carry out at any given instant of time. Whether a policy is successful can now be decided by checking whether the total model which it delivers always satisfies the constraint (13). In other words, a successful policy is one which allows an agent to make sure that all his or her commitments are true. Let us now examine a particular policy:

(14) POLICY FOR ACTION
If at time t agent A's commitment state is c and $Int(c)$ entails that A carries out some action a at t, then A should carry out a. Furthermore, A should only carry out actions which are required by this policy.

Before examining how successful this policy is, let us look at some of its concrete predictions. Consider an agent whose commitment state consists only of the following commitment:

(15) $(\forall v\, t\, t'\, e.(vampire(v) \wedge encounter(john, v, e) \wedge occur(e, t) \wedge successor(t', t) \wedge after(t, t_{utt})) \rightarrow (\exists e'.make\_cross(john, e') \wedge occur(e', t')))$

According to the policy (14), this agent will undertake no actions at all. This explains why we consider it strange for an agent to avoid vampires solely on the basis of the imperative (8). Alternatively, suppose that John's commitment state corresponds with the following formula:

(16) $vampire(v_1) \wedge encounter(john, v_1, e_1) \wedge occur(e_1, t_3) \wedge$
$(\forall v\, t\, t'\, e.(vampire(v) \wedge encounter(john, v, e) \wedge occur(e, t) \wedge successor(t', t) \wedge after(t, t_{utt})) \rightarrow (\exists e'.make\_cross(john, e') \wedge occur(e', t')))$

In words, John is committed to v_1 being a vampire and encountering him at t_3. Now, John's commitment state entails that at t_4 John makes a cross. And hence, according to policy (14), John should make a cross at t_4.[13] Finally, consider a situation in which John is committed to not making a cross:

(17) $\neg\exists e\, t.(make\_cross(john, e) \wedge occur(e, t)) \wedge$
$(\forall v\, t\, t'\, e.(vampire(v) \wedge encounter(john, v, e) \wedge occur(e, t) \wedge successor(t', t) \wedge after(t, t_{utt})) \rightarrow (\exists e'.make\_cross(john, e') \wedge occur(e', t')))$

This commitment state entails, as is to be expected, that John should avoid vampires:

(18) $\forall v\, t\, e.(vampire(v) \rightarrow \neg(encounter(john, v, e) \wedge occur(e, t) \wedge after(t, t_{utt})))$

[13] In fact, the commitment state also needs to include the information that t_4 is the successor of t_3. We assume that the commitment state alway contains necessary information to derive that one instant succeeds another.

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

Let us now look more closely at how successful the policy (14) is in making sure that an agent satisfies all his or her commitments. It turns out that we need to make a number of further assumptions to guarantee its success.

First, we need to impose a constraint on how the commitment state of an agent changes over time. It should not be possible that an agent has a commitment state c_{t_x} at t_x and c_{t_y} at t_y such that c_{t_x} and c_{t_y} are not consistent with each other. It is impossible to satisfy both states in one and the same world/model and therefore (13) would be violated.[14]

Second, if a commitment state c_{t_y} entails an action by an agent at t_x, then c_{t_x} should also entail this action. Thus, in line with the policy (14), the agent will actually carry out the action at t_x. Without this assumption, s/he would not do so, since according to the policy an agent only carries out the actions at time t to which s/he is committed at time t. We could, of course, remove this clause from the policy, but then there would still be no guarantee that the agent would actually carry out the required action.

Third, we need to assume that we are only considering the actions which the agent itself is supposed to carry out. Of course, s/he might also have commitments pertaining to states, events and actions by others. However, for the moment, we need to assume that these commitments are by definition true in the model.

Fourth, we need to assume that the agent's commitment state at each point in time is *complete*. For instance, suppose an agent is only committed to (15). Furthermore, suppose that in that world the agent encounters a vampire at t_4, but does not realize this and therefore is not committed to encountering the vampire. In that case, according to policy (14) the agent will not make a cross. But, of course, he should have done so in order to satisfy constraint (13). In other words, for the policy to be successful the agent's commitment state should be complete in the following sense: there should be no true information which is not part of an agent's commitment state, but which if it were, would cause the agent to perform an action which s/he would not undertake if this information was not present.

Finally, we need to assume that the actions of an agent are independent of each other or that if there are any constraints governing the interactions between different actions, then these should be included in the commitment state of the agent. For instance, suppose that there is a constraint which says that if an agent does not do action a_1 at t_1, then he cannot do a_2 at t_2, such as: one cannot submit a paper without first writing it. Now suppose an agent is committed to submitting a paper at t_2. If this were the only commitment that the agent maintained, then policy (14) would tell him/her to do nothing at t_1. But the aforementioned constraint would then prevent him/her from submitting a paper at t_2. Thus, the agent would not be able to live up to its commitments. To avoid such situations, an agent would need to be committed to the information expressed by the constraint. Suppose the agent were committed to the

[14] Note that the proposed assumption automatically rules out that an individual commitment state is inconsistent.

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

information that (1) if you do not write a paper at t_x, then you do not submit a paper at t_{x+1} and (2) the agent submits a paper at t_2. These two propositions together entail that the agent writes a paper at t_1. And, hence, at t_1 the agent has to write a paper according to (14).

According to (14) an agent will not carry out any actions which are not licensed by (14). This puts a serious constraint on the behaviour of an agent. The condition is intended to prevent an agent from performing some action which prevents him from carrying out some other action which s/he is committed to at a later time. For instance, one might, just for the fun of it, throw away a key, which one needs later on to open a lock. However, if we assume that an agent's commitment state captures such regularities (i.e., that the lock cannot be opened with the key), than we can formulate a more liberal policy:

(19) Policy for Action: Second Version
If at time t agent A's commitment state is c and $Int(c)$ entails that A carries out some action a at t, then A should carry out a. Furthermore, A should only carry out actions which if they were to be added to his or her commitment state would not lead to an inconsistent commitment state.

I have presented a policy and a modified version of it. Both are only successful with respect to (13) when a number of further conditions/constraints hold. Many of these constraints may be seen as too restrictive. However, I hope that by making these constraints explicit we have gained a starting point for the systematic study of agent behaviour. The next step would be to investigate how an agent can make sure that all the constraints are actually satisfied, for instance, by means of an active policy for gathering information, such that all true information which is relevant to the commitments is available to the agent (in other words, a policy for maintaining a complete commitment state).

I conclude this section by drawing attention to a further problem in the analysis of imperatives which goes back to Ross (1941) and which we can now be solved. Ross noted that whereas p entails $p \vee q$ in classical logical, the imperative $p!$ does not entail the imperative $p \vee q!$ For instance, when I tell you to *Post a letter*, then I do not tell you to *Burn it or Post it.* In my approach, we maintain the Classical Logical inference from *Somebody posts a letter* to *Somebody posts a letter or burns it.* However, how information affects the behaviour of an agent is regulated by (19). An agent who is committed to *Posting a letter*, as a result of accepting the imperative 'Post the letter', will in line with (19) simply post the letter, since this is an action which is entailed by his or her commitment. The entailed proposition that s/he post or burn it has no effects on his or her actions. Note that if it did, this could lead to a violation of (13). Suppose the agent had a policy of also making complex sentence true directly. Then, s/he might choose to make the disjunction true by burning the letter. But this would prevent him or her from satisfying the commitment to posting the letter.

Alternatively, a person who accepts to *Burn or post a letter* will have a commitment state which contains this proposition. From this proposition the agent can infer no action whatsoever. Here we need to invoke the completeness

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

assumption to get the agent moving. This assumption says that there should be no true information which is not part of an agent's commitment state, but which if it were, would cause the agent to an action which s/he would not undertake if this information was not present. Suppose our agent considered doing nothing at first. That would mean that s/he would not burn the letter. But this information together with the commitment to burn or post it entails that the agent should post it. Alternatively, if the agent were to first consider not posting the letter, this would cause him or her to burn it. If we invoke the completeness assumption, then it turns out that that the agent either has to post the letter or burn it.

4 Conclusions and Further Issues

An analysis of imperatives has been provided which takes into account the influence of background information on the use of (complex) imperatives. In particular, I have described a model of communicating agents in which imperatives modify the commitments of these agents. The commitments in turn are expected to influence the behaviour of the agent. I formulated a constraint which required that the commitments of the agent should be true. Subsequently, I examined a policy which an agent might adopt in order to satisfy this constraint. It turned out that various further constraints are needed to make the policy successful, i.e., to guarantee that all of the agents commitments are made true.

The proposed analysis of imperatives uses Classical Predicate Logic to represent the propositional content of imperatives. I showed that such an analysis does not need to run into certain problems which have been associated with Classical Logical analyses of imperatives. Therefore, we did not depart from Classical Logic.[15] In this respect, my analysis of imperatives is Gricean in nature (Grice, 1975): we retain the Classical Logical core of the analysis and account for certain features of the use of imperatives in terms of rules which govern the relation between the content/information which is conveyed by imperatives and the role which this information plays in the behaviour of the addressee. In other words, we make a distinction between the questions 'What does the propositional content of an imperative entail?' and 'What can we conclude when somebody utters an imperative with a particular content?'.

The work reported in this paper can be seen as complementary to much work in AI on planning (see, e.g., the collection of papers in Cohen et al., 1990; Grosz & Kraus, 1996 and Lochbaum, 1994) which focuses on commitments to actions which are complex in the sense that these actions can be decomposed into a collection of sub-actions. The actions in question are, however, typically not part of a commitment which is logically complex. In this paper, *logically* complex commitments have been at the centre of our investigation as opposed to complex actions. Due to the focus on logically simple commitments which involve complex actions, work in AI has tended to pay little attention to the relation between an agent's commitment state and the actions which s/he carries out on the basis of

[15] As is, for instance, advocated in Perez Ramirez (2000).

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

that commitment state. Again, this issue has been addressed specifically in this paper.

A possible avenue for further research is an investigation into how the account of logically complex commitments which is argued for in this paper can be combined with AI accounts of complex actions. Further issues of investigation are the inherent temporal vagueness of simple imperatives. For instance, how do we analyse the propositional content of 'open the door'? which might be paraphrased as follows: *Open the door as soon as possible after now but not too long after now.* And finally, what should an agent do if s/he discovers that s/he violated a particular commitment? Such an agent would end up in an inconsistent commitment state if s/he simply added that information to his or her commitment state. Is it possible to define a sensible commitment revision strategy for such agents?

Acknowledgments

The problems concerning the relation between imperatives and action first attracted my attention when working, from 1994 until 1998, on the DenK project (see Bunt et al., this volume). I would like to thank my colleagues on the DenK project for the stimulating environment which put me on a track which eventually led to the work which is reported in this paper. Furthermore, thanks are due to Kees van Deemter and Miguel Perez Ramirez for helpful discussions on the issues which are addressed in this paper.

References

Austin, J. (1962) *How To Do Things With Words.* Oxford: Clarendon Press.

Beun, R.J. (1992) A framework for Cooperative Dialogue. In: Taylor, M., F. Néel & D. Bouwhuis (eds.), *The Structure of Multimodal Dialogue II*, Maratea, Italy: Proceedings of the Venaco Workshop on Multimodal Dialogue.

Beun, R.J. (1994) Mental state recognition and communicative effects. *Journal of Pragmatics* **21**, 191–214.

Biber, D., S. Johansson, G. Leech, S. Conrad & E. Finegan (1999) *Longman Grammar of Spoken and Written English.* Harlow: Longman.

Bunt, H. & Van Katwijk (1979) Dialogue acts as elements of a language game. In: W. Zonneveld & F. Weerman (eds), *Linguistics in the Netherlands 1977–1979*, Dordrecht: Foris Publications, 264–282.

Bunt, H. & W. Black (eds.) (2000), *Computational Pragmatics: Abduction, Belief and Context.* Amsterdam/Philadelphia: John Benjamins.

Cohen, P., J. Morgan & M. Pollack *Intentions in Communication.* Cambridge Massachusetts: The MIT Press.

Grice, H. (1975) Logic and Conversation. In: P. Cole & J. Morgan (eds.), *Speech Acts. Syntax and semantics*, Vol. 3. 41–58, New York: Academic Press.

Grosz, B. & S. Kraus (1996) Collaborative Planning for Complex Group Action. *Artificial Intelligence* **86**(2), 269-357.

Hamblin, C. (1971) Mathematical Models of Dialogue. *Theoria* **37**, 130–155.

In: Bunt, H. and R.J. Beun (2001), Cooperative Multimodal Communication, Lecture Notes in Artificial Intelligence Series 2155, Springer, Berlin/Heidelberg. pp.197-214.

Hamblin, C. (1987) *Imperatives.* Oxford: Blackwell.

Kamp, H. & U. Reyle (1993) *From Discourse to Logic.* Dordrecht: Kluwer Academic Publishers.

Lewis, D. (1979) Scorekeeping in a Language Game. In: Bäuerle, R., U. Egli & A. von Stechow (eds.), *Semantics from Different Points of View*, Berlin: Springer 172–187

Lochbaum, K. (1994) *Using Collaborative Plans to Model the Intentional Structure of Discourse.* PhD. thesis, Cambridge, Massachusetts: Harvard University

Parsons, T. (1990) *Events in the Semantics of English.* Cambridge, Massachusettes: The MIT Press.

Perez Ramirez, M. (2000) Imperatives, State of the Art. Proceedings of CLUK 3, Universities of Brighton and Sussex, 18–25.

Piwek, P. (2000) Imperatives, Commitment and Action: Towards a Constraint-based Model. In: *LDV Forum*, Special Issue on Communicating Agents (To appear). Available at: `http://www.itri.bton.ac.uk/~Paul.Piwek`

Power, R. (1979) The Organisation of Purposeful Dialogues. *Linguistics* **17**, 107–152.

Ross, A. (1941) Imperatives and Logic. *Theoria* **7**, 53–71.

Searle, J. (1969) *Speech Acts: An Essay in the Philosophy of Language.* Cambridge: Cambridge University Press.

Stenius, E. (1967) Mood and language-game. *Synthese* **17**, 254-274.