

# ACTION RELATIONS IN RATIONALE CLAUSES AND MEANS CLAUSES\*

CECILE T. BALKANSKI  
Aiken Computation Lab, Harvard University  
Cambridge, MA 02138, U.S.A

## Abstract

Communication among agents collaborating on a task often involves complex utterances about multiple actions. In this paper, we analyze two particular multi-action constructions, utterances with means clauses and utterances with rationale clauses. We present the distinctive features of these constructions, describe their logical form, and define interpretation rules for deriving their meaning in terms of the generation, enablement and Contributes relations between actions which have been motivated independently by research in plan recognition. We also demonstrate that these rules yield the appropriate properties of rationale clauses and means clauses and show how these constructions can be distinguished algorithmically.

## 1 Introduction

Communication among agents collaborating on a task often involves complex utterances about multiple actions, as illustrated in the following example:

By going and getting the key and then coming and unlocking the grate, for that you get a certain number of points, and then as you collect treasures, you get more points. [1]

Multi-action utterances like these present two particular challenges for natural language interpretation systems: determining an appropriate representation of their meaning, and defining the process by which this representation can be derived from natural language utterances. Substantial work has been done on the development of logical forms as a necessary step toward the goal of representing the meaning of an utterance. However, very little is yet established on the representation of the information communicated in complex action utterances like the one given above, and still less on the interpretation of the relations that connect the actions described in these utterances.

In this paper, we analyze two particular multi-action constructions, utterances with rationale clauses and utterances with means clauses, illustrated respectively in (1a) and (1b) below, with the goal of meeting, at

\*I would like to thank Barbara Grosz, Andrew Kehler, Karen Lochbaum, Stuart Shieber and Candy Sidner for helpful comments regarding this paper. This research has been supported by U S WEST Advanced Technologies, by the Air Force Office of Scientific Research under Contract No.AFOSR-89-0273, and by an IBM Graduate Fellowship.

least in part, these two challenges. As their names suggest, means clauses express the means by which an action is performed, while rationale clauses express the reason, or rationale, of the main clause action. In the last section of this paper, we show how to distinguish rationale clauses from two superficially similar constructions: purpose clauses, as in "Mary bought a suit to wear at the meeting", and infinitival relative clauses, as in "John found the book to give to his sister".

- (1) a. Mary pressed the button to reset the printer.  
b. Mary reset the printer by pressing the button.

Rationale clauses and means clauses are of interest not only as exemplars of multi-action utterances, but also because of the subtle differences in the information that can be felicitously inferred from their use. For example, despite the similarities between (1a) and (1b), the continuation in (2), below, is felicitous after (1a) but not after (1b). The interpretation task therefore requires determining the relations between the actions described in multi-action utterances and drawing appropriate inferences about them. This particular need is one of the issues addressed in this paper.

- (2) But that didn't work because she forgot to plug it in.

We start by presenting distinctive properties of utterances with rationale clauses and means clauses. After describing their logical form, we then present the generation, enablement and Contributes relations between actions [3, 13], arguing for their use in the interpretation of these constructions. We finally define interpretation rules for deriving the meaning of the logical form predicates representing means clause and rationale clauses. We also show how these rules predict the properties of these constructions mentioned above. We end the paper by discussing the scope of the analysis and showing how the constructions to which these rules apply can be distinguished algorithmically. A companion paper [5] examines the mental states of the speaker and performing agent of such utterances, and describes the interpretation model.

## 2 Properties of rationale and means clauses

A detailed analysis of multi-action utterances in task-oriented dialogues [3] and a subsequent examination of a set of selections from the Associated Press news wire

revealed the following distinctive features displayed in utterances with rationale clauses and means clauses.

## 2.1 Semantic properties

**Overlap:** In utterances with means clauses or rationale clauses, the performance of one of the actions referred to in the utterance may result in the simultaneous performance of the other (but need not, see below). The overlap in meaning between utterances (1a) and (1b) exemplifies this property: in both utterances, Mary's pressing the button may result in her resetting the printer at the same time.

**Distinction:** It follows from past tense utterances with means clauses, but not from those with rationale clauses, that the agent's "goal action" was successfully performed. This action is described in the main clause of an utterance with a means clause, but in the adjunct clause of an utterance with a rationale clause. This property is reflected in the difference between utterances (1a) and (1b) with respect to the continuation in (2): only from (1b) can it be concluded that Mary successfully reset the printer<sup>1</sup>. Thus (2) is not felicitous in the context of (1b), but it is in that of (1a) because in that context, Mary might have failed to reset the printer.

**Ambiguity:** Utterances with rationale clauses are ambiguous with respect to the relative timing of the two actions referred to: the performance of the main clause action may result in the simultaneous performance of the adjunct clause action, or it may only make possible its subsequent performance. This property explains why, in the examples below, response (3b) is felicitous whereas the same utterance in the context of (4a) is not.

- (3) a. John: I inserted a diskette to back up my file.  
 b. Mary: Okay, now back up the file.  
 (4) a. John: I typed 'cp new.txt' to back up my file.  
 b. Mary: ✗ Okay, now back up the file.

## 2.2 "Algebraic" properties

Drawing an analogy with the algebraic properties of relations, utterances with means clauses and rationale clauses are irreflexive, asymmetric and transitive, as reflected in the following examples (the first utterance or utterance pair of each set includes a means clause, the second, a rationale clause).

- (5) a. \*Mary reset the printer by resetting the printer.  
 \*Mary reset the printer to reset the printer.  
 b. Mary reset the printer by pressing the button.  
 \*Mary pressed the button by resetting the printer.  
 Mary pressed the button to reset the printer.  
 \*Mary reset the printer to press the button.  
 c. Mary fixed her problem by resetting the printer by pressing the button.  
 Mary fixed her problem by pressing the button.

<sup>1</sup>The companion paper [5] addresses the issue of embedding inferences like this one in belief contexts.

Mary pressed the button to reset the printer to fix her problem.  
 Mary pressed the button to fix her problem.

Similar examples may be found to show that utterances with rationale clauses about sequential actions have the same properties, e.g., for transitivity:

- (6) John borrowed the book to bring it back home.  
 John brought the book back home to look for papers on adjunct clauses.  
 John borrowed the book to look for papers on adjunct clauses.

Transitivity affects utterances with means clauses and rationale clauses in different ways. The transitivity of utterances with means clauses, illustrated in the first utterance pair of (5c), makes it possible to describe actions at lower levels of detail, i.e., with greater granularity. The transitivity of utterances with rationale clauses, and in particular those that describe sequential actions as illustrated in (6), increases not the granularity of the description, but the "distance" between the actions described, i.e., the number of steps in the sequence of actions linking them.

## 3 Actions and action relations in logical forms

Logical forms represent the literal meaning of an utterance and are derived compositionally by semantic interpretation rules based on the syntactic structure of the utterance. Our logical forms are based on Davidson's formalism [6], because action reification facilitates the representation of utterances about multiple actions and relations among them. These logical forms are existentially quantified sentences of first-order logic, with predicates that include an additional argument position for the action being described. Davidson's framework is extended for multi-action utterances by introducing multiple action variables and by representing connectives like "by" (introducing means clauses) and "to" (introducing rationale clauses) using two-place predicates ranging over these action variables, as illustrated in (7)<sup>2</sup>.

- (7) a. Mary reset the printer by pressing the button.  
 LF:  $\exists x_1, x_2, \text{reset}(x_1) \wedge \text{past}(x_1) \wedge \text{agt}(x_1, \text{Mary})$   
 $\wedge \text{obj}(x_1, \text{Printer}) \wedge \text{by}(x_1, x_2)$   
 $\wedge \text{press}(x_2) \wedge \text{obj}(x_2, \text{Button})$   
 b. Mary pressed the button to reset the printer.  
 LF:  $\exists x_1, x_2, \text{press}(x_1) \wedge \text{past}(x_1) \wedge \text{agt}(x_1, \text{Mary})$   
 $\wedge \text{obj}(x_1, \text{Button}) \wedge \text{in.order.to}(x_1, x_2)$   
 $\wedge \text{reset}(x_2) \wedge \text{obj}(x_2, \text{Printer})$

The following subsections discuss the four main components of these representations: action variables

<sup>2</sup>Extensional representations of propositions, like this one, have difficulties with respect to certain intensional phenomena such as substitution of identicals. Hobbs discusses this problem and a number of possible solutions [10], but these questions lie beyond the scope of the present paper.

(e.g.,  $x_1$  and  $x_2$ ), action descriptions (e.g., the predicates *reset* and *press*), tense predicates (e.g., the predicate *past*), and action relations (e.g., the predicates *by* and *in\_order.to*).

### 3.1 Action variables

While Davidson's formalism assumes that action variables denote "real" action tokens, action variables in our formalism correspond to action entities in the discourse model. The ability to refer to an entity anaphorically has been suggested as a test for "discourse entity-hood" [20]. Actions referred to in means clauses and rationale clauses pass this test, as illustrated in (8). The pronoun 'that' in this utterance is ambiguous, but the possibility of its coreferring with the adjunct clause action is sufficient justification for attributing discourse entity-hood to that action.

- (8) Mary will clone the rhinovirus to make a vaccine  
for the common cold and that will make her famous.

Action entities may represent actual, or "real", actions, i.e., actions that have occurred (e.g., the reset action in (7a)), or unrealized actions, i.e., actions that have not yet occurred or might never occur (e.g., the reset action in (7b) that may or may not result from the press action) [4]. Quantification in the logical form is therefore over a universe that contains everything that can be talked about, with no commitment to existence in the real world (as in, e.g., [10, 17]). Discourses in collaborative settings often involve discussions about unrealized actions. Agents, for example, might make suggestions for future actions, or discuss the advantages and disadvantages of a particular course of action. It is thus crucial to adopt a representation that will account for these actions as well.

### 3.2 Action descriptions

Actions are represented by predicates taking a single action variable, e.g., *reset*( $x_1$ ), and every action argument is represented by a separate predicate, e.g., *agt*( $x_1$ , *Mary*). This type of representation has been proposed by a number of researchers (e.g., [21]) as an alternative to Davidson's representation, in which only the optional arguments are represented as separate predicates. One reason is the arbitrariness that is sometimes involved in deciding on the argument structure of various action predicates. For the present analysis, this type of representation is useful because it avoids having either to introduce existential variables for missing arguments or to determine how these arguments are recovered during the process of building the logical form. Missing subject and time arguments of adjunct clause actions in the logical forms of means clauses and rationale clauses (e.g., those in (7)) can be recovered through the interpretation of the action relations (see Section 5.2). Not all rationale clauses, however, have empty subject positions, e.g., "Mary prepared lobster (in order) for her guests to have a taste of New England cuisine."

### 3.3 Tense predicates

The tense predicates in our logical forms capture two pieces of information: temporal information and information about action occurrence. This distinction is necessary because actions, whether their time is past, present, or future, need not occur. For example, agents may discuss a future action without ever performing it. Similarly, agents may talk about a past action that never occurred, as in "John said that Mary left yesterday, but he was wrong".

If the logical form predicate *past* holds of some action, then that action occurred (in some world or belief context) and its occurrence was in the past. In (7), *past* is predicated of the action variable  $x_1$ , denoting the actions described in the main clauses of both (7a) and (7b). If there is no information to the contrary (e.g., the utterance is not embedded in an intensional context, established either directly within the utterance, as in "Mary thought she reset the printer", or through the discourse context, as when describing a dream), then we can conclude that the action occurred in the real world.

Utterances with rationale clauses and means clauses differ in what information they communicate about the occurrence of the action described in the adjunct clause. In neither construction does this clause involve a tensed verb, which explains why *past* is not predicated of  $x_2$  in the LFs of (7a) and (7b). It is possible, however, to infer from (7a) that the action denoted by that variable did occur, and from (7b) that it might not have occurred, as illustrated in (1) and (2) in the Introduction. This information is captured by the interpretation rules for means clauses and rationale clauses, as explained in Section 4.

### 3.4 Action relations

Finally, the LFs in (7) include predicates representing the connectives that link the two clauses in each of the utterances, namely *by* and *in\_order.to*. The meaning of these predicates is specified by interpretation rules, presented in the next section, that state necessary conditions on the use of these constructions. It is important to note that these predicates are used only to represent the connectives "by" and "to" introducing means clauses and rationale clauses. Other uses of these lexical items are recognizable in a number of ways (e.g., syntactically: "by" followed by a noun phrase, or "to" introducing an infinitival complement (argument) to a verb, as in "John wanted to leave"), as discussed in Section 5.3; these are assumed to be represented differently in the LF.

## 4 Interpreting action relations

Three relations between actions — generation, enablement and Contributes — that have been motivated independently by research in plan recognition [15, 3, 13] can be employed in the interpretation of utterances

with means clauses and rationale clauses. In this section we demonstrate this use of these relations, and also provide a definition of enablement.

#### 4.1 Generation

The generation relation, initially defined by Goldman [8], holds of two actions when the performance of the first results in the simultaneous performance of the second. Generation, however, encompasses more than simply two simultaneous actions. For example, Mary can reset the printer by pressing a button only in certain circumstances, namely when the electricity is on, the printer is plugged in, etc. Performing the generating action when these conditions do not hold does not result in the performance of the generated action. Generation, therefore, is a relation that depends upon contextual constraints, called generation-enabling conditions [15], that must hold during performance time. This relation was formalized by Pollack [15], and subsequently revised in previous work of the author's [3]. The following English gloss of the definition suffices for the present discussion and avoids an unnecessary digression into an explanation of the formalization; A and B are actions, comprising an act-type, agent and time.

- (9) A generates B if and only if
1. A and B are actions performed by the same agent at the same time
  2. there is a set of conditions, C, such that
    - a. these conditions hold during the performance time of A and
    - b. the act-type of A conditionally generates the act-type of B under C.

Clause (2b) states the requirement that there be a *conditional generation* relation between the two act-types and the generation-enabling conditions, C. This relation abstracts away from specific actions to capture the regularities of a generation relation across occurrences. It holds between act-types  $\alpha$  and  $\beta$  and conditions C if (a) whenever an action of type  $\alpha$  occurs while these conditions hold, an action of type  $\beta$  occurs at the same time and (b) neither an occurrence of an action of type  $\alpha$ , nor C holding, is by itself sufficient to guarantee the occurrence of an action of type  $\beta$  [15]. An important property of the generation relation is that if it holds of two actions A and B, and A occurs, then so does B. This follows from the generation relation requiring the relevant conditions to hold during performance time (as stated in clause (2a) in the definition).

In light of the generation relation, we can reconsider the examples in (1). Both these utterances refer to an underlying conditional generation relation: in the right circumstances, pressing the button results in simultaneously resetting the printer. They differ, however, in that they do so from different perspectives: utterance (1b), with the means clause, asserts the occurrence of the generated action (the resetting action), while utterance (1a), with the rationale clause,

does not (cf. example (2)). Therefore, the relevant generation-enabling conditions must have held during performance time in the context of (1b) but not necessarily in that of (1a). Because the generation relation requires these conditions to hold during performance time, it correctly describes the relation between the two actions in (1b) but not necessarily that in (1a).

Utterances with rationale clauses, like (1a), only have the *potential* for realizing a generation relation, and this potentiality rests on two different factors: (a) the possibility that there is no conditional generation relation between the two actions, as in "Mary ate an apple to reset the printer", or (b) the possibility that the generation-enabling conditions do not hold during performance time, as in (1a) where the two actions are related by a conditional generation relation, but the associated generation-enabling conditions might not hold during performance time.

The following sections, however, will show that an analysis of rationale clauses based solely in terms of the generation relation is too restrictive.

#### 4.2 Enablement

The interpretation of rationale clauses also requires appealing to the *enablement* relation; for example, in (10) below (cf. (1a)), the performance of the main clause action does not (potentially) result in the *simultaneous* performance of the adjunct clause action. Rather, it only makes possible the *subsequent* performance of the adjunct clause action. The generation relation, therefore, does not correctly describe the action relation in this utterance. Instead, the relationship is one of enablement, a relation we define as holding of two actions when the performance of the first brings about a condition that is necessary for the subsequent performance of the second [3].

- (10) Mary inserted a diskette to back up her file.

Much of the literature that makes use of the enablement relation has been vague about its intended interpretation, using it at different times to mean different things and not explaining why one action must be performed before the other (e.g., [18]). In the following definition, the dependency between the enabling and enabled actions is explained by the role of an intervening conditional generation relation<sup>3</sup>. A closer examination of utterance (10), for example, shows that the effect of the enabling action, namely that a diskette be in place, provides the generation-enabling condition of a conditional generation relation between typing some command and backing up the associated file.

- (11) A enables B if and only if
1. the time of A is prior to the time of B,
  2. there are a set of conditions, C, and a third action A' such that
    - a. one of the conditions in C, C<sub>1</sub>, holds as a result of the performance of A, and
    - b. the act-type of A' conditionally generates the act-type of B under C.

<sup>3</sup>This is not the only possible type of dependency [3].

If  $C_i$ , the condition brought about by the performance of A, is the only element of the set of generation-enabling conditions C, then the performance of A is both necessary and sufficient for the subsequent performance of B.

### 4.3 Contributes

So far, utterances with rationale clauses have been interpreted as instances of the generation or enablement relation. This construction, however, can also relate two actions when the chain of relations between them involves both generation and enablement relations, as illustrated in (12a), or a complex action constructor [3], as illustrated in (12b):

- (12) a. Mary built the table to sell it at the market.  
         Mary sold the table at the market to earn money.  
         Mary earned money to pay her rent.  
         Mary built the table to pay her rent.  
         (enablement + generation + enablement)
- b. I picked up the foot end for us to lift the piano.  
         (simult + generation)

Utterances with rationale clauses thus allow for a wide range of relations between the main clause and adjunct clause actions. To capture this phenomenon, we interpret them as instances of the *Contributes* relation. This relation, also originally introduced for plan recognition purposes, is defined as the transitive closure of the generation and enablement relations, and the complex action constructors [13]. It therefore covers the entire range of relations that can exist between the actions in this type of utterance.

## 5 Analysis

### 5.1 Interpretation rules

The inference rules defining the meaning of means clauses and rationale clauses are given below. The *by* and *in\_order\_to* predicates in the left hand side of the rules are the logical form predicates representing the connectives introducing these clauses<sup>4</sup>. The verbs “generates” and “contributes” in the right hand side are used here, as elsewhere in the paper, in their technical sense, corresponding to the action relations defined in the previous sections. Note that the order of the arguments in the *by* predicate and generation relation is reversed.

- (13) a. LF1:  $by(A,B) \rightarrow$   
         B generates A and B occurred.
- b. LF2:  $in\_order\_to(A,B) \rightarrow$   
         A contributes to B,  
         within restricted belief contexts.

The LF1 rule maps the LF representation of means clause to the generation relation and specifies that the

<sup>4</sup>This paper does not address the reference problem and assumes that the mapping of action variables in the logical form to action entities in the discourse model has already been applied.

generating action occurred. This occurrence is a feature of the means clause construction and not of the use of the gerund form to describe the action. Given an utterance with a means clause, it can therefore be inferred that the generated action occurred as a result of the generating one and that the relevant contextual conditions held during performance time.

The LF2 rule maps the LF representation of rationale clauses to a *Contributes* relation within a belief context, capturing the fact that (a) the relation can be any combination of the generation and enablement relations, as well as the complex action constructors and (b) that the relation between the two actions is only potential. A new set of action relations is not needed for representing the potentiality of the action relation in these utterances. Instead, as described in the companion paper [5], it is best expressed in terms of the different beliefs and intentions of the speaker (S) and performing agent (G). For example, G could believe a *Contributes* relation between actions A and B, while S does not; or G and S could both believe the relation, but have different beliefs regarding associated conditions and actions, e.g., G may believe A is necessary and sufficient for B whereas S believes A is only necessary.

When utterances with rationale clauses involve simultaneous actions, then the *Contributes* relation derived by the LF2 rule specializes to the generation relation. Exactly when such specialization is possible is a plan recognition problem that goes beyond the scope of this paper. Utterances sometimes provide temporal or agency information that can be used to optimize the plan recognition process. In particular, if the performance times, or performing agents, of the two actions are known to be different, then a specialization to generation is ruled out. When temporal and agency information is not available, then *recipes* provide another source of information for disambiguating rationale clauses [3, 14].

### 5.2 Consequences of the rules

The two interpretation rules LF1 and LF2 predict all of the properties of utterances with rationale clauses and means clauses that were presented in Section 2.

**Semantic properties:** The LF2 rule accounts for the Ambiguity property of rationale clauses by making use of the *Contributes* relation which can relate either simultaneous or sequential actions. Because the *Contributes* relation may specialize to a generation relation, the LF1 and LF2 two rules account for the Overlap property. Finally, because the LF2 rule is based on potentiality, while the LF1 rule is not, these rules also account for the Distinction property.

**Algebraic properties:** It can be shown that the generation and *Contributes* relations are irreflexive, asymmetric and transitive, reflecting the algebraic properties of the NL utterances that realize them. By describing simultaneous actions performed by the same agent, the generation relation also captures the increased granularity of the action description: by de-

scribing sequential (as well as simultaneous) actions, the Contributes relation captures the increased distance between the actions described.

In addition, these interpretation rules account for the missing agent and time arguments of the actions described in means clauses and rationale clauses. Because the generation relation requires the agents and times of the two actions to be identical, the missing arguments in a means clause can be recovered on the basis of the agent and time information associated with the main clause action. The Contributes relation, however, imposes no temporal or agency constraint. Therefore, without additional information, it is not possible to determine whether the agent (when missing) and time of a rationale clause are the same as those of the main clause, cf. "I brought this miserable Morgon to enjoy with our dinner" [2].

### 5.3 Scope of the analysis

The LF1 and LF2 interpretation rules predict specific facts about English, namely a large set of properties characterizing the differences and similarities between rationale clauses and means clauses. They do so by defining the meanings of the *by* and *in.order.to* LF predicates in terms of independently motivated action relations. These predicates are useful for representing not only the meaning of rationale clauses and means clauses, but also that of related constructions. For example, the *in.order.to* predicate can be used to represent the connective "so that" in utterances such as "Mary got up early so that she could catch the first train to Montreal". The *by* predicate can also be used for non-clausal constructions, as in "John got the information he needed by means of an interview". Furthermore, these predicates are meant to be language independent. In French, for example, the *in.order.to* predicate would be used to represent the construction "pour + infinitive", as in "Marie a regardé des films pour préparer son voyage au Japon" (Mary watched movies to prepare for her trip to Japan), and the *by* predicate, to represent the construction "en + gerund", as in "Marie a préparé son voyage au Japon en regardant des films" (Mary prepared for her trip to Japan by watching movies). Unlike the corresponding English construction, however, the French "en-clause" is ambiguous between a means clause reading (as given in the English translation) and a temporal reading that can be translated as "while watching movies". Determining how to disambiguate these two interpretations is an interesting research question.

In the remainder of this section, we show how to distinguish algorithmically the constructions to which the LF1 and LF2 rules apply, thereby making it possible for an interpretation system to build the correct LF representations.

**Rationale clauses:** These clauses are optional infinitival adjuncts, introduced either by the infinitival marker "to", or, when the subject is not missing, by the lexical item "for" followed by a noun phrase (realizing the subject) followed by "to". Both these con-

structions are optionally preceded by "in order". Note that the LF2 rule applies to rationale clauses with or without explicit subjects. For example, given "Mary prepared lobster (in order) for her guests to have a taste of New England cuisine", Mary's preparing lobster contributes to her guests having a taste of New England cuisine.

A large number of to-infinitival clauses (with or without subjects) to which the LF2 rule does not apply can be distinguished on syntactic grounds, namely by not being optional infinitival adjuncts. Utterances that are syntactically ambiguous show a corresponding semantic ambiguity, and the LF2 rule applies to the rationale clause reading; e.g.: "Bush warned his press secretary to prevent a leak", where the infinitival clause may function as an optional adjunct (i.e., realizing a rationale clause) or a complement.

Two types of adverbials to which the LF2 rule does not apply, but which are syntactically indistinguishable from rationale clauses, are comment clauses, as in "John is not sure what to do, to put it briefly", and outcome clauses, as in "Mary awoke one morning to find the house in an uproar" [16]. Given their status as parentheticals, comment clauses are recognizable in text by a comma, and in speech by intonation (note that without the comma, the infinitival clause becomes a rationale clause). Outcome clauses are restricted to verbs of discovery or to the verbs "be" or "become" (e.g., "She lived to be 100"), and are therefore also distinguishable.

Finally, there are purpose clauses, e.g., "Mary bought a suit to wear at the meeting". Purpose clauses are also optional infinitival adverbials, but differ from rationale clauses in the number and type of gaps they allow and in their control pattern<sup>5</sup> [2, 12]. They also differ from the superficially similar infinitival relative clauses by not being attached to a noun phrase. The LF2 rule, in fact, makes correct predictions for this construction as well, e.g., Mary's buying a suit contributes to her wearing the suit at the meeting. Purpose clauses with inanimate subjects, however, as in "John bought a shelf to hold his books", require some adjustment with respect to the action's role in the Contributes relation. Actions relations are between actions, and actions are assumed to have human agents. Therefore, rather than deriving a Contributes relation between John's buying a shelf and the shelf's holding his books, the system should derive a relation between John's buying a shelf and his having something (namely the shelf) to hold his books. This issue is currently being investigated.

**Means clauses:** These clauses are very easily distinguishable by their form: the lexical item "by" followed by an -ing clause (in this context, "by" may be treated as a preposition or a subordinator). Apparent counterexamples to the LF1 rule are utterances of the form "Mary pays her rent by building furniture",

<sup>5</sup>From a descriptive point of view, rationale clauses differ from purpose clauses in that only the former allow the paraphrase with the words "in order to" [12].

in which the two actions are sequential. However, the fact that LF1 does not apparently apply to this example, is a sign that additional pragmatic reasoning (e.g., conversational implicatures [9]) is needed. In this particular example, LF1 does apply, though indirectly through the utterance of which it is a metaphor, namely, "Mary makes a living by building furniture". The need for further Gricean reasoning of this sort is reflected in the markedness of such examples.

## 6 Conclusion

In this paper, we analyzed utterances with rationale clauses and means clauses, presented their logical form, and defined interpretation rules for these two constructions that make use of independently motivated action relations. These rules make it possible to draw appropriate inferences about the actions described in these utterances, thereby accounting for the similarities and differences between rationale clauses and means clauses. The algebraic properties of the action relations were shown to reflect the properties of the utterances that realize them. We also showed how to distinguish the constructions to which these rules apply.

This analysis fills an important gap in the natural language interpretation literature. Surprisingly little research has examined multi-action utterances such as those with means clauses and rationale clauses. There is a large body of linguistics research on purpose clauses, but it focuses on syntactic aspects of the construction and issues of control (e.g., [2, 12]). In the computational linguistics literature, Huettner et al's work on the generation of adjunct clauses [11] and Di Eugenio's analysis of instructions [7, 19] both examine purpose constructions, but from different perspectives than that of this paper [5].

This research is being extended in two main directions. One is to examine future tense utterances and other types of utterances about unrealized actions. Another is to cast the interpretation rules in a general processing model that takes into account the mental states of the agents involved. Beliefs and intentions are necessary for the treatment of unrealized actions, as well as for capturing other aspects of rationale clauses and means clauses. In particular, utterances with rationale clauses, but not necessarily those with means clauses, communicate the speaker's beliefs about the performing agent's intention to perform the actions described in the utterance. Finally, mental states are indispensable in any interpretation model that purports to account for inferences that are drawn on the basis of utterances. Initial progress in these two areas is reported in the companion paper [5].

## References

- [1] Albert, D. 1990. Cooperative dialogs while play-

ing adventure. TR-02-90, Harvard Univ.

- [2] Bach, E. 1982. Purpose clauses and control. In Jacobson, P. & Pullum, G., eds, *The Nature of Syntactic Representation*. D. Reidel Publ. Co.
- [3] Balkanski, C. 1990. Modelling act-type relations in collaborative activity. Technical Report TR-23-90, Harvard Univ.
- [4] Balkanski, C. 1991. Logical form of complex action sentences in task-oriented dialogs. In *Proc. ACL'91, Student Session*.
- [5] Balkanski, C. 1992. Actions, beliefs and intentions in rationale clauses and means clauses. In *Proc. AAAI'92*.
- [6] Davidson, D. 1967. The logical form of action sentences. In Rescher, N., ed, *The Logic of Decision and Action*. University of Pittsburgh Press.
- [7] Di Eugenio, B. 1992. Goals and action in natural language instructions. Technical Report MS-CIS-92-07, Univ. of Pennsylvania.
- [8] Goldman, A. 1970. *A Theory of Human Action*. Princeton University Press, Princeton, NJ.
- [9] Grice, H.P. 1975. Logic and conversation. In *Syntax & Semantics, vol 3, Speech Acts*, Academic Press, NY.
- [10] Hobbs, J. 1985. Ontological promiscuity. In *Proc. ACL'85*.
- [11] Huettner, A., Vaughan, M., and McDonald, D. 1987. Constraints on the generation of adjunct clauses. In *Proc. ACL'87*.
- [12] Jones, C. 1991. *Purpose clauses: syntax, thematics and semantics of English purpose constructions*. Kluwer Academic Publishers.
- [13] Lochbaum, K., Grosz, B., and Sidner, C. 1990. Models of plans to support communication: An initial report. In *Proc. AAAI'90*.
- [14] Lochbaum, K. 1991. An algorithm for plan recognition in collaborative discourse. In *Proc. ACL'91*.
- [15] Pollack, M. 1986. Inferring domain plans in question-answering. Technical Report 403, SRI International, Menlo Park, CA.
- [16] Quirk, R., Greenbaum, S., Leech, G., and Svartvik, J. 1985. *A Comprehensive Grammar of the English Language*. Longman.
- [17] Schubert, L. K. and Hwang, C. H. 1990. An episodic knowledge representation for narrative texts. Technical Report 345, Univ. of Rochester.
- [18] Schuster, F. 1988. Anaphoric reference to events and actions: a representation and its advantages. In *Proc. COLING'88*.
- [19] Webber, B. and Di Eugenio, B. 1990. Free adjuncts in natural language instructions. In *Proc. COLING'90*.
- [20] Webber, B. 1983. So what can we talk about now? In Brady, M., and Berwick, R. C., eds, *Computational Models of Discourse*. MIT Press.
- [21] Wilensky, R. 1991. The ontology and representation of situations. In *Proceedings of KR'91*.