Modelling Concession Across Speakers in Task-Oriented Dialogue

Kavita E. Thomas and Colin Matheson
School of Informatics
University of Edinburgh
{kavitat, colin}@cogsci.ed.ac.uk

Abstract

We determine criteria for modelling concession signalled via *plan-based* "but" in task-oriented dialogue (TOD) (following (Thomas, 2003)) based on some examples from the corpora, analysing where Speech Act (SA) and planning information fails to predict concession. In our approach we focus on cases involving cross-speaker concession, where the speaker accepts part of the previous speaker's turn but signals contrast via "but", and we argue that this contrast can (in the examples shown) be modelled as concessive. Then given a representation of task-plan history in the Information State (IS) model of the dialogue, we present an initial framework for an algorithm that predicts concessive interpretation across speakers in two situations encountered in the corpora. We motivate this work by showing how it updates beliefs in the PTT (Poesio and Traum, 1998) model of dialogue and can be used to facilitate recognition of planning mismatches and potentially avoid misunderstandings, and more generally improve discourse understanding and natural language generation (NLG).

1 Introduction

In this paper we determine criteria for modelling concession across speakers in task-oriented dialogue (TOD), extending the treatment of cross-speaker concession presented in (Thomas, 2003) to propose an algorithm that models two cases of cross-speaker concession signalled by the contrastive cue "but". We focus our analysis on what is communicated at the planning level in examples from the Maptask and TRAINS TOD corpora[1] and present an algorithm that interprets concession and generates an interpretation based on the presence of certain salient (predominantly plan-related) features. Due to the sparseness of cross-speaker concession data involving "but" (Thomas, 2003), we model two examples from the TRAINS and Maptask TODs rather than model a statistically frequent type of cross-speaker concession. We argue that semantic and speech act (SA) representations don't provide enough information to distinguish what's really being communicated at a task level, and motivate our work by claiming that understanding task-plan-related communication and inferences is essential for full interpretation and generation of utterances in TOD.

Our approach is based in the Information State (IS) framework, assuming the PTT (Poesio and Traum, 1998) model of dialogue. We motivate our work by showing how it updates beliefs in this model of dialogue and facilitates interpretation and generation by predicting the in-

[1]MAPTASK dialogues can be browsed interactively at www.ltg.ed.ac.uk/ amyi/maptask/demo.html and TRAINS dialogues can be found at www.cs.rochester/research/trains/.

ferences involved, which aids the interpretation of speakers' intentions and resolves potential planning mismatches. Our treatment of concession follows work by (Lagerwerf, 1998), extending the approach presented in (Thomas and Matheson, 2003) and (Thomas, 2003) which argued that Lagerwerf's treatment of concession can be extended to dialogue, and connects concessive cases in TOD to the planning information necessary to interpret them.

2 Determining Salient Characteristics of Cross-speaker Concession

Given the cue of contrast signalled by cross-speaker "but"[2], we start by determining what criteria are salient for distinguishing concession in a constructed example illustrating the sort of situation we're initially considering, and then analyse examples from the Maptask and TRAINS TOD corpora. We then determine how concession might be generated in the IS framework and focus on how the utterances connect with planning operations in the task, employing plan-recognition models like those described in (Litman and Allen, 1987) and the COLLAGEN project (Lesh et al., 1999). Our approach follows the concessive analysis presented in (Lagerwerf, 1998) and extended in (**?**) to TOD. Lagerwerf claims that concession requires a contextually available claim (or *tertium comparationis*, TC), for which both a positive and a negative argument are provided. Extending this approach to TOD, (Thomas, 2003) argues that the positive and negative arguments for and against the TC are presupposed as expectations launched by the relevant assertions in the adjacent speaker turns in these cross-speaker "but" cases with respect to the contextually relevant claim (i.e., the TC). Here we determine criteria for predicting concession in order to generate the TC and stance of the speakers with respect to the TC in cases where interpreting concession is reasonable. We motivate this work by demonstrating that recognising concession, the TC, and speakers' expectations regarding this TC, facilitates interpretation and generation of subsequent utterances in the dialogue.

[2] We consider that an occurrence of "but" involves contrast across speakers if the constituent it modifies (i.e., usually the main clause of the turn it appears in) is contrasted with something in a previous speaker turn. Cases where "but" is turn-initial or preceding the first clause of the turn are almost always cross-speaker cases. Here we look for contrast with something in the immediately preceding turns for simplicity.

2.1 A Motivating Example

Let's consider how concession might be interpreted in actual cases by tracing how we'd analyse a simple example first. The concessive interpretation algorithm proposed in section 4 describes two particular situations. In the first case, there is an opening question about which both speakers debate, as in the example below:

(1) B1: Should we add the mushrooms to the sauce now?

A2: That should be fine.

B3: But I thought we need to add the beans first.

Here the TC is the question asked in B1, and A argues for, while B argues against. Presumably B raises the question because of her own doubts about the validity of adding mushrooms at this point in the task, and wants to debate the issue with A in order to resolve what to do. Her doubt is emphasised by the objection she raises in B3. A would then need to respond to this objection by either denying B's concern that beans need to be added first, or by agreeing to B's implicit proposal to either add the beans first, or not add the mushrooms to the sauce. The other case that the algorithm addresses is the case in which A proposes a plan or goal (as in the example below), and B objects to this goal:

(2) A1: Let's add the mushrooms to the sauce.

B2: But don't we need to add the beans first?

B2': But the sauce is already too thick.

Notice that here the TC simply asks whether A's proposal is a valid one, with A implicitly supporting it and B objecting explicitly in B2 and indirectly in B2'. Sometimes (as B2' illustrates) the plan is proposed indirectly, as we shall later see in the TRAINS example. Similarly, interpreting B2' as an objection to the proposed plan requires recognising that B is pointing out a negative effect which will be exacerbated by A's proposal of adding something more to the sauce.

2.2 Analysing an Example from Maptask

Maptask dialogues involve two participants with slightly different versions of a map with various landmarks on it, and the goal is for the route giver (who has a route drawn on her map) to direct the follower along this route from start to finish, so a large part of the task involves establishing that the other participant has the same salient landmarks along the route in order to communicate clear directions. This is the sort of situation assessment we see going on in Maptask dialogue q3nc4 below (with disfluencies removed):

(3) F1: right so i'm going underneath this mountain and along to the abandoned truck ... not underneath the mountain ... above the mountain?

G2: above the mountain yeah.

F3: so it's to the right of the mountain that i'm going?

G4: the abandoned truck's to the left of the old temple

F5: yes, BUT the mountain's between the abandoned truck and the old temple

The DAMSL annotation scheme (Allen and Core, 1996), which annotates utterances by coding their multi-level direct SAs, would predict that G4 has a Forward-Looking Function (FLF) of a statement assertion, and F5 has a Backward-Looking Function (BLF) of partial-acceptance (under the agreement category) and also has a FLF of statement assertion. In terms of planning operator type, both G4 and F5 are facts. Neither semantic information nor direct SAs nor planning operator-type seem to shed much light on the issue of how concession might be interpreted here. Clearly we need to delve deeper into the actual problem the agents are trying to solve, namely their plans, in order to get at the underlying meaning of what's being conveyed by F5's "but". The follower (F) has the mountain on her map which the giver (G) doesn't have, and G only realises this discrepancy in maps after hearing F5.[3] This portion of the dialogue establishes the environment (and therefore path) that must be followed in this part of the map. I.e., the overall goal is to establish the route that must be traversed, which involves the subgoal of establishing salient landmarks around the given route that must be circumnavigated.

Concessive analysis F5 can be seen as expressing concession, since F agrees with G4's assertion but asserts previously unmentioned information which is contrastive because it indicates that despite G4's assertion being acceptable, G4's attempt to respond to her question (F3) hasn't answered it. Notice also that F3 indirectly proposes a route around the mountain, and that G4's indirect answer (i.e., assertion of a fact) can only be interpreted with respect to F3, since G must believe that the orientation of the truck with respect to the temple helps F to recognise the route she is trying to navigate. F5 accepts this information, but implies that it fails to answer F3 via the introduction of a contrastively cued (i.e., "but") assertion. This assertion then needs to be checked against the speaker's plan in order to determine where the contrast lies. In this case, the plan (which includes relative locations of objects on the map as constraints and preconditions of navigational actions) indicates that the path needs to be defined with respect to the mountain (i.e., she needs to know whether to go above or below the mountain).

So the concession expresses acceptance of G4's assertion but indicates that this assertion fails to answer the question, which is a Speech Act (SA) level relation rather than something at the planning level, but planning information is necessary for determining this. Determining where the contrast lies with this planning information also enables detection of discrepancies and deficiencies in their maps and facilitates alignment of their perceived environments. After hearing F5, G understands F's difficulty (they've aligned their maps), and communicates the route to F taking into account the mountain that's on F's map. The TC in this case can be framed as the proposal raised in F3 to go to the right of the mountain, and G4 can be assumed to be against this TC since it doesn't answer F3's question satisfactorily (according to F5), leaving F5 to support the TC.

Is this really argumentative? However, thinking of the two turns as supporting and denying the TC doesn't add much useful information here, and

[3] G2 thinks F1 refers to the white mountain (a different one), which is further along on the route, and responds accordingly.

may not even be accurate, since the speakers are not disputing a claim but resolving a confusion; it might be more useful instead to think of the two turns as expressing a different perspective on the environment. That is, in G's map, because there's no mountain nearby, he expresses the salient features he thinks F must circumnavigate, while in F's map, the route must be specified with respect to the mountain that's between the two features (i.e., the truck and the temple) that G mentions. So while argumentative stance *per se* is not an issue here, determining where the discrepancy lies is crucial, and viewing G4 and F5 as conveying different perceived environments with respect to the salient feature specified in the TC (i.e., the mountain) resolves this confusion. In other words, determining the TC and the speakers' perspective with respect to this TC facilitates this realignment by isolating the discrepancy in their plans and establishing how the preconditions raised in G4 and F5 both need to be accounted for in their joint plan.

2.3 Analysing the TRAINS example

In the TRAINS TODs[4], the agents' plans are considerably more complicated than in many other TODs (e.g., Maptask) and involve planning future actions rather than interleaving task and speech actions. Here the agents plan how to achieve a transportation task together where one agent knows additional constraints and guides the user. In the dialogue subsection below (from d.93-15.2, with disfluencies removed), one needs to understand where in the plan the speakers are in order to recognise that the user (U) is actually proposing a new subplan with her assertion of an effect, and that the system (S) agrees that this effect holds but rejects the validity of the indirectly proposed subplan (on the basis that the desired task goal, in this case getting the maximum amount of oranges to Bath by 7 a.m., wouldn't be met).

(4) U: well actually I'll already have an engine in Bath after I unload the boxcars right

S: right but you wouldn't have it in time

Where semantic and SA analyses fail Semantic representations of the utterances clearly can't capture the inferences that must be made in order to determine that U is actually indirectly proposing a plan. Considering the speech act (SA) level, the DAMSL annotation scheme would predict for U the FLF of information-request and statement assertion and the BLF of reject-part (cued by the "actually"). (Indirect SAs are not marked, though if they were, then the utterance would be annotated as an action-directive.) DAMSL would classify S' utterance as involving the BLFs of both partial-acceptance and partial-rejection (for the "right" and the "but you wouldn't have it in time" parts of the utterance respectively), along with the FLF of statement-assertion. Simply mapping the assertional content of the turns to planning operators maps both turns to effects in this case, which doesn't give enough information to infer that U is proposing a new plan. Clearly we need to account for where in the joint plan these operators occur and how they are related, i.e., we need to account for prior planning actions when annotating cases like this, which we argue can be facilitated by accessing a task-plan history (TPH) for the dialogue so far that can be maintained in a dynamic semantic representation like the IS framework.

Where plans can help This particular example involves an implicit plan proposed by U to take the engine already in Bath back up to Corning with the boxcars to collect more oranges and return to Bath. This implicit proposal is recognised by S, who then criticises the effect of this plan with respect to achieving the agents' main goal of getting to Bath by 7 a.m. with the maximum amount of oranges by indicating that they can't get the oranges back to Bath by 7. In our approach, we assume that a planner that can recognise the implicit proposal in U's assertion and the criticism of this proposal's effect in S's turn is available[5] and communicates with the dialogue manager in order to keep the IS updated with what's happening in terms of the agents' salient planning actions. We also assume that an interpreter is available to map utterances into planning operators and SAs, which is feasible in restricted TOD domains and has been implemented in other systems (e.g., the

[4] See www.cs.rochester.edu/research/trains/

[5] Via scripts we write that communicate with the planner and request forward-chaining or plan-recognition with various inputs.

BEE tutorial dialogue project[6] and the COLLAGEN project (Lesh et al., 1999).

Concessive analysis In this example the TC asks whether U's proposed plan is valid, and U argues for it (inferred since U proposes it) while S argues against. Notice that in order to determine this TC we need to infer U's proposed plan and record it in the IS, requiring a dynamic representation of plan history for the dialogue so far. In fact, given an evolving record of the plan history in the IS, we can interpret this example as involving Denial of Expectation (DofE, following (Thomas and Matheson, 2003)) with the expectation $having\_engine\_at\_Bath\_at\_current\_stage > plan\_of\_getting\_more\_oranges\_to\_Bath-in\_time\_succeeds$, which is more informative than the DofE interpretation proposed using their algorithm, which would predict $having\_engine\_at\_Bath > having\_it\_there\_in\_time$, since their approach doesn't have the benefit of prior planning history. The bonus of being able to predict useful expectations for DofE in these indirect TOD cases provides an additional argument in favour of maintaining planning history in the IS. In the next two sections, we will describe how to derive concessive interpretations in plan-based cases like this.

3 Incorporating Task-Plan History in the IS

In (Thomas, 2003) we saw that annotation based on shallow analysis didn't give conclusive results or indicate trends that help predict rhetorical relations, so we bypass the problems of sparse and diverse TOD data by resorting to a "deeper" analysis that requires a dynamic representation of task-plan history in the IS. Consider the IS shown in fig. 5 which contains salient task-plan information that helps us interpret the TRAINS example above.

We omit irrelevant fields and acts here for brevity, and just show what the relevant part of the IS[7] (following the IS structure given in (Matheson et al., 2000))would look like for S (after hearing U utter CA1[8]) in Ex.4. Notice that TPH contains Task Actions (TA), mirroring Conversational Acts (CAs) in the Dialogue History (DH). The recognition of CAs as signalling specific TAs (e.g., TA4 and TA5) would occur within a CA interpretation process which is called by the dialogue manager, and would involve plan-recognition (e.g., (Litman and Allen, 1987) and (Lesh et al., 1999)). This process would require communication with the planner, since many CAs communicate information about planning operators (e.g., the assertions of effects in the TRAINS example) rather than directly making plan proposals. In these cases the interpreter would need to call the planner with the specific planning operator communicated and have the planner forward-chain from this operator to see if the TC is achievable given this input. These processes would be called by update rules in the dialogue manager that call the interpreter and planner in order to determine what intentions are being communicated in the utterances.

How the TPH helps resolve anaphoric actions For example, before update CA3 would contain $assert[S, not\_enough\_time(do(X))]$; in the IS above, the argument X of the anaphoric action communicated by S in CA3 has already been resolved to TA4 in the update process by recognising that U is making a proposal. U's indirect plan proposal can be inferred via the interpreter, planner and the TPH; i.e., TA4 is inferred by the planner upon receiving CA3, thereby enabling recognition (in the planner) that an alternative to TA3 is being proposed, and resulting in TA5. This resolution should occur before the update module determines whether or not to generate concession, and the algorithm proposed in the next section takes as input an IS that's been updated with this information, which simply requires the dialogue manager to call the algorithm after applying the update rules which interpret the input and updates the IS with the TAs recognised by the planner. So to resolve anaphoric actions like $do(X)$ in CA3, we need to incorporate into the update module a

[6] See www.cogsci.ed.ac.uk/ jmoore/tutoring/index.html.

[7] The abbreviated fields are: Previous and Current Dialogue Unit (PDU and CDU), Ground (GND) and Conditions (COND). TPH is omitted in CDU for brevity, since it's empty.

[8] The field values for the $take$ action in TPH are: *with (where E1-3 are engines), from, to and via.*

(5)
$$\left[\begin{array}{ll} \text{PDU} & \left[\begin{array}{ll} \text{GND} & [\text{TA3: } take(E3, Elmira, Bath, Corning)] \\ \text{DH} & [\text{CA1: } assert(U, have[engine, in(Bath), after(unload(boxcars))])] \\ \text{TPH} & \left[\begin{array}{ll} \text{TA4:} & propose\_plan(take([E2, 2boxcars], Bath, Bath, Corning)) \\ \text{TA5:} & propose\_alternative(TA4, TA3) \end{array}\right] \end{array}\right] \\ \text{CDU} & \left[\begin{array}{ll} \text{DH} & \left[\begin{array}{l} \text{CA2: } acknowledge(S, CA1) \\ \text{CA3: } assert(S, not\_enough\_time(TA4)) \\ \text{CA4: } concession(S, TC[willTA4succeed?], for(TC, U), \\ against(TC, S)) \end{array}\right] \\ \text{COND} & \left[\begin{array}{l} accept(U, CA4) \rightarrow scp(U, concession(S, TC[willTA4succeed], \\ for(TC, U), against(TC, S))) \end{array}\right] \end{array}\right] \end{array}\right]$$

rule of the form:[9]

1. If PDU.TOGND.TPH contains a TA that matches $propose\_alternative(TA_i, TA_j)$ then
   (a) If PDU.TOGND.DH contains a CA_n of the form $assert(Y)$ where Y contains $do(X)$, then replace $do(X)$ with TA_i in Y.

3.1 The Benefits of Modelling Cross-speaker Concession

If U accepts the concession communicated by S, e.g., by saying "Yeah, I guess you're right, going back up to Corning for more oranges wouldn't work", then she is socially committed (Matheson et al., 2000) to the concession relation being communicated with the given TC and S's expectation that her proposed plan wouldn't work. In fact, she needs to recognise that S is responding to her implicit plan in order to know that S recognised it. On the other hand, if she doesn't accept S's concession, she can indicate this by saying "But it only take an hour to travel between Bath and Corning, so I should get it there in time", which requires that she understands that S is signalling a rejection of the implicit plan she proposed, and counters S's argument that her proposed plan wouldn't work. So both recognising these implicit plans and also the concession relation S communicated is essential for her to generate the appropriate response in cases like these.

4 An Initial Framework for an Algorithm for Concession

We propose an update algorithm to model the cases of cross-speaker concession discussed above. (We will take into account other cases after more examples from the corpora are analysed.) This algorithm would be called by an update rule in the dialogue manager triggered by "but". If the required circumstances (outlined in the algorithm) apply in the given case, concession is predicted. How do we distinguish the simplicity of the Maptask example above, where the TC for concessive interpretation arose from the immediately preceding question in F(3) from the plan-dependendent interpretations in the TRAINS example? As a first step we could check for questions to which alternative answers are being given in the turns containing and preceding the "but", and if so, predict that the question or proposal is the TC and generate concessive interpretations if the planner verifies that the relevant turns argue for and against (or favour and disfavour–given the planner) the TC (case 1 in the algorithm below).

If there's no preceding question then one should check whether the PDU proposes a plan, goal or action[10], and if so, whether the "but" turn counters this proposal also via communication with the planner (case 2). Here we only address cases where the planning operator communicated in CDU (the "but" turn) is an effect or precondition (referred to as *Operator(TA$_i$)* below, where TA_i is the task action proposed in PDU) that can't be met given the proposal made in PDU, and we'll focus on other possible planning operator combinations in future work. So the initial distinctions an algorithm modelling these cases should make are as follows, given a dialogue with B uttering "but" in CDU (i.e., turn i):

1. If turn B(i − 2) contains a question in CA_k in

[9]This rule is an example of many similar rules for other situations involving anaphoric reference to actions that need to be resolved in order to interpret the dialogue.

[10]Here we only deal with assertions and not indirect proposals, though this will be addressed in later work. We will also address preconditions and effects (both future ones and failed past ones) in PDU in future work, since speakers can object to actions they think are forthcoming or point out problems they think have occurred.

UDU.TOGND.DH,[11]

(a) Call interpreter with (CA_l, CA_j, CA_k) and get back (TA_l, TA_j, TA_k), where CA_l comes from the assertion in turn A$(i-1)$ found in PDU.TOGND.DH and CA_j comes from the assertion in B(i) found in CDU.TOGND.DH where both are uttered by different speakers. TA_k contains the task action queried by the question contained in CA_k; if CA_k queries an effect or precondition, then replace TA_k by CA_k.[12]

(b) Pass $alternatives(TA_l, TA_j, TA_k)$ to the planner, which returns "yes" if CA_l and CA_j provide alternate answers, arguments or comments w.r.t. the question in CA_k and also returns $for[TC[CA_k], speaker(CA_l)]$ and $against[TC[CA_k], speaker(CA_j)]$ if the arguments in TA_l and TA_j support and counter[13] the question in CA_k and vice versa otherwise.

   i. If "yes" then

      A. Update the IS by adding to CDU.TOGND.DH a new CA of the form $concession(TC[CA_k],$ $for[TC[CA_k], speaker(CA_l)],$ $against[TC[CA_k], speaker(CA_j)])$ or vice versa the arguments of for and against depending on which turn supports and counters the TC and resolving the speakers by determining where CA_k and CA_l occur in the IS.

   ii. Else do nothing.

2. Else if [[PDU.TOGND.DH contains a CA_n of the form $assert(Y)$[14] (where Y corresponds to TA_j, which was returned by the planner after the interpreter passed it Y)] and [PDU.TOGND.TPH contains $TA_j = propose\_plan(Y, Plan\_Y)$ (where the planner instantiates $Plan\_Y$ with the plan Y proposes directly or implicitly)] and [CDU.TOGND.TPH contains $CA_k = Operator(TA_i)$, where $Effect$ can be either an effect or precondition and is instantiated in CA_k (e.g., here it's not_enough_time).]] then

   (a) Update the IS by adding to CDU.TOGND.DH a new CA of the form $concession(TC[``IsPlan\_Yvalid?''],$ $for[TC, speaker(CA_n)],$ $against[TC, speaker(TA_k)])$ where speaker gets resolved appropriately by determining whose turn it is; the speaker of PDU is assumed to favour the proposal they put forth in CA_n while the speaker of the "but" utterance in CDU presumably argues against this proposal.

3. Else do nothing. (This case will be elaborated after analysing more examples.)

Our approach addresses cases like the Maptask and TRAINS examples via update rules in the dialogue manager which call the interpreter (which then sends the planner the planning operators communicated to get back TAs). The interpreter then passes the TAs it receives back to the update rules which update the IS accordingly with the given TAs. An update rule in the dialogue manager triggered by "but" then calls the concessive interpretation algorithm, which tests whether the given situation can be interpreted as concession. The algorithm orders the concessive cases so that we search for the simpler situation of a leading question first (the Maptask case), before exploring whether an alternative plan is being proposed (as in the TRAINS example). The algorithm then spells out how to derive concessive interpretations in either of these cases. We leave for future work (via the "else do nothing" statements) examples that convey concession in different ways. Future work will also address possible social obligations (following (Matheson et al., 2000)) launched by proposals raised (e.g., the TC).

5 Conclusions and Future Directions

In this paper we present a novel treatment of cross-speaker concession when signalled by "but" in two specific cases. We motivate our work by analysing how the examples modelled communicate concession, and show that recognising concession across speakers gives useful information about speakers' expectations and goals and helps avoid potential planning mismatches by enabling the speakers to recognise each others' intentions and respond appropriately and meaningfully. While the sorts of cases we argued were concessive here are far-removed from examples displaying the sort of argumentative force one usually associates with concession (*epistemic* relations, following (Sweetser, 1990)), we call these TOD cross-speaker "but"

[11] UDU was left out of the IS shown in (5) for brevity, and contains ungrounded dialogue units prior to PDU. See (Matheson et al., 2000; Poesio and Traum, 1998) for more information on the IS model and fields involved.

[12] For now we will only allow the question to lie in the turn before the turn preceding the "but", i.e., B$(i-2)$, though a more developed treatment should allow for (possibly via search) salient unanswered questions which have been possibly interrupted by (e.g.) clarificatory or subgoal-resolving subdialogue.

[13] We determine whether TA_l and TA_j support or counter the TC (CA_k) by performing forward-chaining from both TAs in turn to see whether the TC is achievable given that input.

[14] We don't require Y to be an effect, as occurred in the TRAINS example to allow for some generalisability.

examples concessive because they display different stances with respect to some proposal, and identifying these stances gives useful information about what the speakers believe, how they perceive their environment, and what their plans are, and as shown in the Maptask example, they facilitate alignment and grounding and help expedite clarification of misunderstandings.

In future work we will extend this analysis to more cases of concession across speakers found in task-oriented corpora to expand our coverage, and we'll use these examples to also enlarge upon our current conception of concession across speakers in dialogue. Exploring how the combinations of planning operator-type pairs in concessive cases convey potential or previous problems and how these can be resolved will be a practical benefit of this work.

References

J. Allen and M. Core. 1996. DAMSL: Dialog Act Markup in Several Layers. University of Rochester.

L. Lagerwerf. 1998. Causal Connectives Have Presuppositions. *Catholic Univ. of Brabant, Holland Academic Graphics, The Hague, The Netherlands*.

N. Lesh, C. Rich, and C. Sidner. 1999. Using Plan Recognition in Human-Computer Collaboration. *International Conference on User Modelling*.

D. Litman and J. Allen. 1987. A Plan Recognition Model for Subdialogues in Conversation. *Cognitive Science*, volume 11.

C. Matheson, M. Poesio, and D. Traum. 2000. Modelling Grounding and Discourse Obligations Using Update Rules. *Proceedings of the North American Association for Computational Linguistics*.

M. Poesio. and D. Traum. 1998. Towards an Axiomatization of Dialogue Acts. *Proceedings of Twente Workshop*.

E. Sweetser. 1990. From Etymology to Pragmatics. Metaphorical and Cultural Aspects of Semantic Structure. *Cambridge University Press, Cambridge*.

K. Thomas. 2003. Modelling 'but' in Task-Oriented Dialogue. *4th International and Interdisciplinary Conference on Modeling and Using Context*.

K. Thomas and C. Matheson. 2003. Modelling Denial of Expectation in Dialogue. *5th International Workshop on Computational Semantics*.

K. Thomas. 2003. Criteria for Modelling "But" in Task-Oriented Dialogue. *To appear in Proceedings of ESSLLI Workshop on Discourse Particles*.