

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Trading in a multiplayer board game: Towards an analysis of non-cooperative dialogue

Permalink

<https://escholarship.org/uc/item/9zt506xx>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 34(34)

ISSN

1069-7977

Authors

Guhe, Markus
Lascares, Alex

Publication Date

2012

Peer reviewed

Trading in a multiplayer board game: Towards an analysis of non-cooperative dialogue

Markus Guhe (m.guhe@ed.ac.uk), Alex Lascarides (alex@inf.ed.ac.uk)

School of Informatics, Informatics Forum, 10 Crichton St, Edinburgh EH1 2HX, Scotland

Abstract

We collected and analysed a small dialogue corpus of people playing *The Settlers of Catan*. Dialogues are trading negotiations where Gricean maxims of cooperativity often break down as players adopt conflicting intentions in their attempt to win the game. This has consequences for what information players are sharing and for the sincerity of their contributions.

In this paper, we motivate and describe a two-level scheme for analysing non-cooperative dialogues, where both levels are interdependent. Each dialogue move is a move in the game (e.g., an offer to trade), and a coherent contribution to the dialogue so far, connected to a prior segment with a coherence relation, such as indirect answerhood or rejection. Parsing and generating coherence relations is computationally feasible (e.g., Baldridge & Lascarides, 2005), and here we'll argue that their semantics help to identify the game move, even when it is implicated rather than linguistically explicit.

Keywords: non-cooperative dialogue; dialogue move; negotiation; *The Settlers of Catan*; JSettlers; multiplayer game.

Non-cooperative dialogue

Standard accounts of dialogue usually assume that the agents cooperate with one another on at least two levels (Grice, 1975): they coordinate the conventions that govern linguistic meaning (basic cooperativity); and they share attitudes towards what is said, including shared intentions (content cooperativity). Tutoring is a widely investigated example where content cooperativity is a reasonable assumption.

But there are many dialogue scenarios where conflicting preferences compel agents to perform dialogue moves that aren't content cooperative. Such situations range from differences in preferences with relatively low stakes (e.g., whether to go to a French or an Italian restaurant) to very high, life changing, stakes (e.g., a cross examination of a defendant by a prosecutor in court). Lying in court has potential high rewards (if the lie is not uncovered) as well as high penalties (if it is). But telling the truth or lying are just the two extremes in a spectrum of non-cooperative behaviour. To illustrate this, consider the following cross-examination of a defendant by a prosecutor (from Solan & Tiersma, 2005):

- (a) Prosecutor: Do you have any bank accounts in Swiss banks, Mr. Bronston?
- (b) Bronston: No, sir.
- (c) Prosecutor: Have you ever?
- (d) Bronston: The company had an account there for about six months, in Zurich.

The locutionary content of both (b) and (d) are true, but Bronston succeeds in *deflecting* the prosecutor's enquiry with a misleading implicature with (d): (d) implicates that Bronston never had any Swiss bank account, and this is false.

In this paper, we describe dialogue moves in non-cooperative situations, drawing on data from a dialogue cor-

pus of people playing *The Settlers of Catan*. In the corpus, each dialogue move is aligned with a machine readable version of the game state. We present a scheme of the types of dialogue moves that occur in our corpus. Our eventual goal is to build a dialogue system that plays *The Settlers of Catan*, engaging in non-cooperative dialogues like humans do. For such a system, the agent playing the game will use the dialogue moves from that scheme to interact with the other players so as to forward its chosen game moves. This will require us to parse and generate the semantic content of these dialogue moves; we will exploit coherence relations from SDRT for this purpose (Asher & Lascarides, 2003).

The Settlers of Catan

The Settlers of Catan is a multiplayer board game set on the fictional island *Catan*. *Catan* is represented by a map consisting of hexes (see Fig. 1; see catan.com for the full set of rules). Two to four players settle on the island by building settlements and cities connected by roads. It is a zero-sum game: the first player to win 10 victory points wins and all others lose. One obtains victory points by, for example, building a settlement (1 point) or a city (2 points). To build, one needs resources: clay, wood, rock, ore and sheep.

Players take turns and attempt to obtain resources and to build. A turn starts with the roll of dice. Each player potentially obtains or loses resources through dice rolls. This depends on a combination of: the number rolled, the location of game pieces on the board, and the resources currently held. (So, future game states are non-deterministic, necessitating players to calculate the risks of moves). The player whose turn it is can trade resources with the bank or other players and can use resources to build roads, settlements or cities.

Dialogues where the players trade resources surface during game play as the players try to exchange resources with one another. Their decisions about what resources they want and what they are prepared to give up are influenced by what they need to build, e.g. a road requires 1 clay and 1 wood. A player can only build on a location if it touches one of his already existing pieces (road, settlement or city); all settlements and cities must be separated by at least 2 roads (of any player). Thus, players' decisions about trades are not only determined by the decisions about what they want to build and where they want to build it, but also by their estimates of what will most advance, or undermine, the building strategies of the other players (see Thomas, 2003). Players can agree any trade; they can even lie or bluff, as found in classic conceptions of bargaining games (Osborne & Rubinstein, 1990).

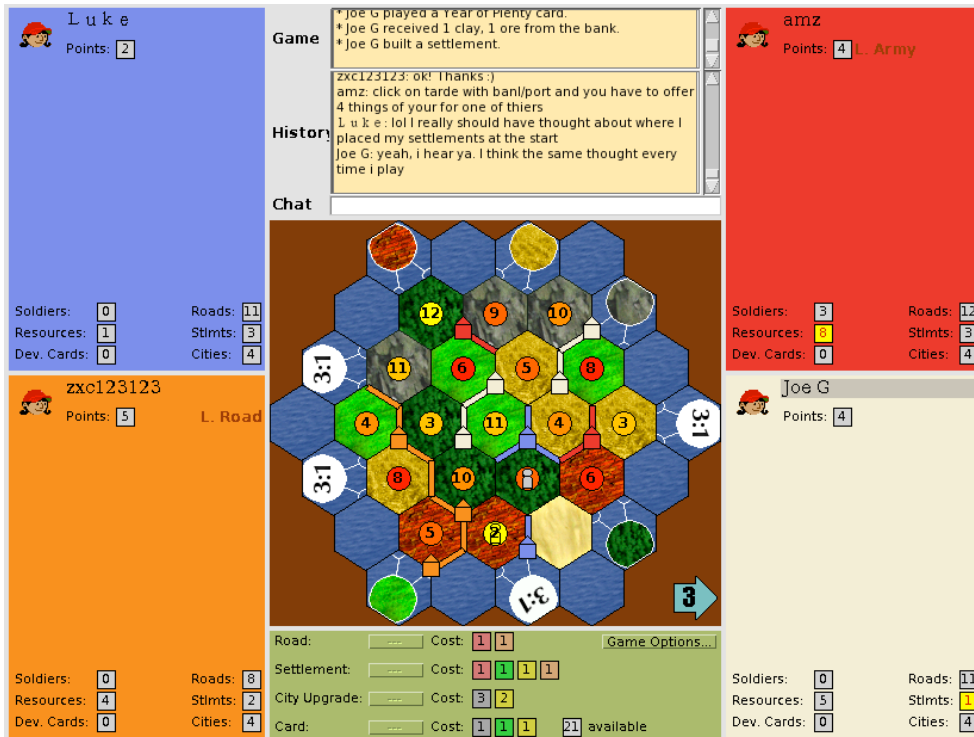


Figure 1: A game of *The Settlers of Catan* in the JSettlers interface.

Modifying JSettlers for Corpus Collection

To analyse trading dialogues and to perform inferences about an agent’s underlying game strategy, one needs a corpus where each utterance is aligned with a machine readable version of its synchronous game state. To this end, we adapted an existing open source implementation of Settlers, cf. Fig. 1 (JSettlers2, version 1.1.12). JSettlers mimics the physical version of the game in that each player can see the whole board but only his own hand (in a colour-coded panel). In JSettlers, trades are done graphically via the interface, see Fig. 2a. The player whose turn it is has two panels containing ‘I Give’, ‘I Get’ and colour coded buttons for the other players – he can offer a trade by adding appropriate values for these three items. This prompts the players in question, who can then accept or reject the trade through button clicking.

We adapted this interface to enforce trades to be negotiated linguistically prior to an actual exchange of resources. Specifically, the buttons required to exchange resources are hidden by a ‘Register a Trade’ button (Fig. 2b), and only one other player can be selected via the colour coded buttons. Players were instructed to first agree a trade via the chat interface and only then ‘register’ the trade. Clicking ‘Register a Trade’ revealed the trading interface (Fig. 2a), and enacting the trade is achieved in the same way as the original version of JSettlers.

Negotiation dialogues

In contrast to other task-oriented dialogues, e.g., the map task (Anderson et al., 1991; Guhe & Bard, 2008), *The Settlers*

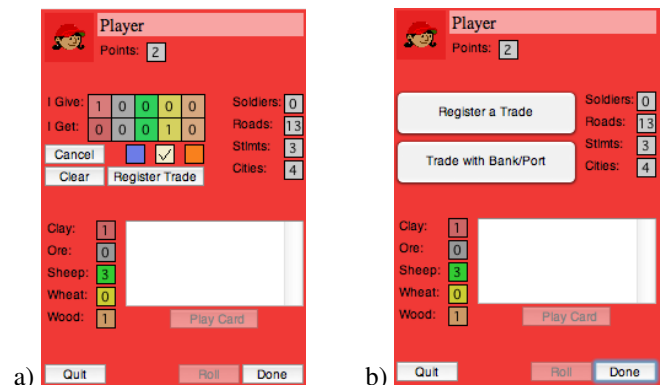


Figure 2: (a) A player’s hand panel in JSettlers. (b) The modified hand panel during normal game play. Trade is agreed via the chat interface prior to using the panel to register the trade.

of *Catan* is not a cooperative game but a zero sum game: only one player can win, all others lose. Even so, it is sometimes necessary for players to form alliances to block another player from winning, e.g., *don’t give tomm wheat, he can win the game if he gets wheat* (from dialogue 1). Because of this, players avoid dialogue moves that might undermine their trustworthiness; they like to *appear* to be cooperative even if they aren’t actually cooperating at the content level (i.e., adopting a shared intention), e.g., they may reject a trade even when they are able to perform it by saying *no sorry* (when asked *can I buy wood for 1 ore?* and the player having wood; from dialogue 2). Joint action takes place on multiple levels

simultaneously (Clark, 1996), and we find cooperativity on many levels in our corpus, e.g. all players share a common language, they follow the rules of the game, and they do actually trade with each other.

To analyse trade interactions, it is not necessary to treat the dialogue exchanges within a single game as a single dialogue. Instead, they appear as chunks, especially since the game state changes between trades (thanks to dice rolls, for instance), leading the players to adopt different preferences about what resources they need at the start of a new *dialogue game*, as compared to their preferences at the end of the last one. In addition, the game is too complex for players to reason about the complete game tree anyway; so, they structure their game space, aiming to achieve shorter term goals by using heuristics about how to advance their goal of winning (Thomas, 2003). Because of this structuring of the game space, it is natural to analyse the dialogues in discrete chunks, each addressing a short term goal, e.g. to obtain wood. We adopted this subdivision into dialogue games from the analysis of map task dialogues (Kowtko, Isard, & Doherty, 1993).

We collected an exploratory pilot corpus of 21 games, each containing multiple dialogue games, which we will make available publicly in due course. All participants were native speakers of English and naive to the purpose of the study. Most were students at the University of Edinburgh.

The corpus is not currently annotated because the purpose of collecting it was to obtain a sufficient sample so as to inform decisions about an annotation scheme for the corpus collection proper, which we are engaged in at the moment. For this reason, we have not analysed this pilot corpus in quantitative detail; e.g. with respect to differences between individual players, linguistic features, strategies, or level of cooperativity in the dialogue moves. Instead, we report here on the range of moves that surfaced in this corpus. There are several types of dialogue games (sub-dialogues), e.g., clarifying the rules of the game, discussions of the game state, small-talk, and, most importantly, trade negotiation dialogue games, which also are the most frequent in the corpus. Each trade negotiation constitutes a dialogue game.

We analyse dialogues on two levels, a game-strategic level and a dialogue level. Put differently, we distinguish the level of *what* game move is being made and *how* that move is realised through lexical and compositional semantics plus an assumption that the contribution is coherent in context. This distinction is particularly useful for building an agent that will play *The Settlers of Catan* and be able to communicate with other players using natural language. For building this agent, the game-strategic computations should be separate from the computations needed for understanding and generating dialogue moves for achieving those game strategies.

Please be aware of the potential confusion between the terms *dialogue move* (or: *move in a dialogue game*), which we adopt from the map task, and *game move*. The latter refers to a move in the game *The Settlers of Catan*, e.g., building a road, or trading 1 ore for 3 wheat, while the former refers to

a move in the dialogue or equivalently a type of speech act; this might be a *property* of the utterance (e.g., to question) or a *relation* (e.g., to give an indirect answer), where the successful performance of the speech act depends on the content of both the current utterance and a prior segment of the dialogue. In trading dialogues, game moves (such as to get 1 ore for 3 wheat) are achieved by performing dialogue moves (e.g., by *accepting* a prior offer through uttering *It's a deal*). In trading dialogues, game moves are performed *by* performing a dialogue move: so accepting a trade may be performed by performing the speech act *Accept*. If not further specified, we are talking of dialogue games and dialogue moves.

Dialogue Games

Dialogue moves in the map task

We will take the twelve types of dialogue moves from the map task (Kowtko et al., 1993) as our starting point. Six dialogue moves initiate a new game: **Instruct**: Give an instruction to the fellow player of what to do. **Check**: Request confirmation of whether a previous utterance has been understood correctly. **Query-YN, Query-W**: A yes/no or wh-question about some aspect of the task, except clarification of an instruction (which is a Check move). **Explain**: Provide new information about the goal of the task with respect to the current state.

Then, there are six further moves: **Clarify**: Clarify or repeat a prior utterance or known information. **Reply-Y, Reply-N**: Elicited response to Query-YN, Check or Align. **Reply-W**: Elicited response to Query-W or Check. **Acknowledge**: Acknowledge understanding (can close a dialogue game.) **Ready**: Indicate intention to start a new game.

Not all dialogue moves can be uniquely assigned to just one of these classes; some fulfil multiple functions.

Dialogue games can contain sub-games, creating a nested structure. Again, we follow Kowtko et al. (1993): 'Nesting is considered to occur when a sub-game is initiated whose purpose can clearly be seen as contributing to the goal of the current game . . . A break, such as the announcement of a misunderstanding, puts the status of the current game in doubt; the game might continue if the misunderstanding is cleared up, or it might be abandoned.' (p. 3)

We also follow Kowtko et al. in defining dialogue moves as 'an utterance, a partial utterance or a group of utterances, which convey the same specific intent.' (p. 3) In other words, a dialogue move can consist of one constituent or several, and if it is several constituents they may likewise be individual constituents in a further game, thereby creating a hierarchical structure of moves. Dialogue moves are thus defined by their function rather than linguistic form. We provide the analysis of this more detailed level, i.e. how utterances are related, in terms of SDRT (Asher & Lascarides, 2003) below.

Dialogue moves in *The Settlers of Catan*

The dialogues in our corpus bear many similarities with map task dialogues, but there are also significant differences.

The fundamental structure of a trade negotiation consists

of one player making an offer and another player accepting or rejecting it. Thus, we replace the map task's Instruct move with the Offer and Counteroffer moves, and the moves Acknowledge and Ready are supplemented by Accept, Reject and Retract.¹ The full scheme, including the dialogue moves retained from the map task scheme, is given in Fig. 3.

Moves initiating a dialogue game

- Check
- Query-YN, Query-W
- Explain
- Align
- Offer
 - Make Complete Offer
 - Make Incomplete Offer
 - Give Specified
 - Get Specified
 - Undo Trade
- Counteroffer
 - Cannot Accept Proposal
 - Could Accept Proposal
- Game-Strategic Comment

Other moves

- Clarify
 - Reply-Y, Reply-N
 - Reply-W
 - Acknowledge
 - Ready
 - Accept
 - Complete Offer
 - Reject
 - Cannot Trade
 - Won't Trade
 - Retract
 - Null-move
-

Figure 3: Dialogue moves in *The Settlers of Catan*.

The Offer, Counteroffer and Reject moves have subtypes. For **Offers**, apart from the rare **Undo Trade** move, where a player requests to reverse a trade, the main distinction is whether an offer is complete. With **Make Complete Offer** a player specifies precisely both what he is prepared to give and what he wants in return, and from whom. In **Make Incomplete Offer** a player only specifies some values of the trade (e.g., the resource she wants, but not how much, nor from whom, nor for what). **Give Specified** dialogue moves only specify what the player wants to trade (*Does anybody want clay?*) **Get Specified** (*I need clay.*) is similar.

As we are investigating the non-cooperative aspect of dialogue, we subdivide **Reject** moves into those where the player **Cannot Trade** because he lacks the required resources, and **Won't Trade**, where he could make the trade but chooses not to. Because we record the game states, we can distinguish Cannot Trade from Won't Trade by looking at the resources of the speaker. Analogously, **Counteroffers** can be made be-

¹*Instruct* is a very map task specific dialogue move – in the map task one person instructs another person to draw a route on a map. While in *The Settlers of Catan* a player may give an instruction to another player, no player is obliged to follow it. Particularly in trade negotiations, no player can instruct another player to do something, which is why we are dropping it here.

cause the player **Cannot Accept the Proposal** or **Could Accept the Proposal** but does not want to do so.

As in the map task scheme, dialogue moves aren't mutually exclusive, e.g., a Counteroffer (of either type) can also be a Make Incomplete Offer (from dialogue 19): *Hardie: i'll trade 2 wheat for 1 wood – Gytis: I need clay or ore*. Taken by itself, Gytis's dialogue move is simply a Make Incomplete Offer, but in this context, it is also a Counteroffer (he wants clay or ore instead of wheat). Note that the Counteroffer is drawn from both what Gytis says and what Hardie says, and to infer this particular Counteroffer one needs to draw on the analysis at the second (the discourse) level, see the next section. It is through recognising the coherence relation between Gytis's contribution and Hardie's, and reasoning about that relation's *semantics* that one infers that Gytis has implicated acceptance of a part of Hardie's offer, i.e. give Hardie 1 wood.

We add some more dialogue moves to this scheme. Players can start a new sub-game by making **Game-Strategic Comments** in the context of trade dialogues, e.g. *don't give tomm wheat, he can win the game if he gets wheat* (dialogue 1).

A **Complete Offer** move completes the trade offer of a Make Incomplete Offer move by specifying the missing information: i.e., it differs from Counteroffer in that no part of the original offer is rejected. Further, as long as a trade has not been concluded, a player who has offered a trade may **Retract** the offer. Finally, because *The Settlers of Catan* is a multiplayer game, a player can decide to stay silent and not to respond to a trade offer. This is a **Null-move**.

Because dialogues are entirely unrestricted, trade dialogue games can be interrupted by other types of dialogue games, e.g. a dialogue game on how to use the hand panel, cf. Fig. 2. We do not consider these cases here.

Sincerity and Informativeness

Players can be non-cooperative in at least two ways. First, a player P_2 's response to P_1 's prior move may reject the goal that underlies P_1 's prior utterance, even though he was in a position to fulfil that goal: a Won't Trade move, e.g., P_1 asks for wheat and P_2 responds *no, sorry* despite having wheat. Second, a speaker's utterance may semantically entail and/or implicate content they know to be false. While a player can in principle perform an insincere dialogue move other than Reject, e.g. Offer, in practice such moves are rare (they have little foreseeable benefit but foreseeable penalties should the lie be exposed). So insincere moves tend to be rejections; indeed, they tend to be Won't Trade (and not Can't Trade) rejections.

Our corpus exhibits the whole range from sincere statements to outright lies. We mainly distinguish sincere rejections from deflections and lies. In the following examples from our corpus, we specify the resources a player has when making an utterance by using the abbreviations c (clay), o (ore), s (sheep), wh (wheat) and wo (wood).

Sincere rejection. Sincere rejections of both subtypes are possible. A move is sincere only if declarations of the player's

resources are truthful. For example, amz’s utterance below is a sincere Can’t Trade move; Joe’s and Luke’s utterances are sincere Won’t Trade moves (all taken from dialogue 8):

amz[c=0|o=0|s=1|wh=1|wo=1] i dont have any clay
 Luke[c=1|o=0|s=0|wh=1|wo=0] I have one clay but I want it
 for myself?
 Joe G[c=1|o=0|s=1|wh=0|wo=0] sorry, i’m holding on to
 mine

Deflection. Deflections are misleading implicatures as described earlier; the following extract from dialogue 2 yields a misleading scalar implicature (since Tomm has wheat):

Cat[c=0|o=2|s=0|wh=2|wo=1] anyone for ore or wheat?
 Thomas[c=2|o=2|s=0|wh=2|wo=3] i have ore

Lie. Lies are moves where a player declares false information, e.g., this Reject from dialogue 7:

Lorelei1292[c=0|o=1|s=3|wh=9|wo=1] can anyone give me
 some clay for some wheat?
 AM123[c=2|o=0|s=0|wh=2|wo=1] sorry have none of that!

Even assuming AM123 misunderstood the direction of the trade offer, the only available antecedents for resolving the anaphoric expression *none of that* are *wheat* and *clay*; thus, AM123’s utterance declares content she knows to be false.

Null-move. Trade offers can also be rejected by not answering at all (from dialogue 11):

Britt[c=1|o=1|s=5|wh=0|wo=0] anyone in need of sheep?
 Simon[c=0|o=0|s=2|wh=4|wo=0] no thanks
 David A[c=0|o=1|s=1|wh=2|wo=0] ⟨null⟩
 Din[clay=0|ore=1|s=1|wh=1|wo=1] ⟨null⟩

Limited informativeness. It is strategically important to decide how much information about one’s own hand a player is willing to reveal. Deflections are a prime example for this.

Volunteering information. In contrast to the limited cooperativity exhibited in, for instance, Null Moves, there are also many instances of players behaving much more cooperatively than strictly speaking they need to in order to advance the domain-level game. For instance (from dialogue 3):

Euan[c=0|o=0|s=2|wh=2|wo=1] Anyone got any ore?
 Cardlinger[c=0|o=1|s=0|wh=7|wo=3] only the 1
 Joel[c=0|o=0|s=2|wh=2|wo=1] havent had it in awhile

One possible explanation for why players might choose to provide more specific content than simply a direct answer to the question is that by doing so they are more likely to be perceived as cooperative and providing positive face to their interlocutors (Brown & Levinson, 1978), which may be useful later in the game.

The Dialogue Level: SDRT

The previous section detailed dialogue moves on the agent level. Now we describe how those moves interact with con-

tent at the dialogue level. This link between dialogue moves on the one hand and coherence relations and semantics on the other, is needed for identifying which dialogue moves (or communicative intentions) an agent has revealed during language interpretation, and it is needed for deciding which utterances (or dialogue moves) to make during language production.

Consider the utterance *I need clay*. Without any context, this is a Make Incomplete Offer move. But in the exchange *A: I want wood. B: I need clay*, the same utterance has to be understood in relation to the content of the previous utterance. Reasoning about its coherent use in context leads one to infer that it is a Complete Offer dialogue move.

We are using the framework of *Segmented Discourse Representation Theory* (SDRT, Asher & Lascarides, 2003; Lascarides & Asher, 2009) for specifying such coherence relations. In SDRT’s model of dialogue, each dialogue agent *publicly commits* to coherence relations between his own utterance and prior ones, even if they were said by another agent. Shared public commitments then reveal agreed content; inconsistent public commitments reveal disputes. For instance, above, *B* commits to the coherence relation *Plan-Elab(a,b)* (where *a* is *A*’s utterance and *b* is *B*’s). This means that *b* provides information that elaborates a plan for achieving the communicative intention underlying *a*. *Plan-Elab* is a veridical relation: a public commitment to *Plan-Elab(a,b)* entails a commitment to the contents of *a* and *b*. Thus *B* is committed to *a*: he implicates endorsement of it by the illocutionary effects of the speech act *Plan-Elab* that he performed.

Such moves reveal commitments to preferences as well; in our domain, these preferences are (partial) information about offers that an agent would prefer to perform. Cadilhac, Asher, Benamara, and Lascarides (2011) describe a detailed algorithm for computing such preferences from the dialogue’s content. It exploits the recursion over the discourse structure that is engendered by coherence relations to build a partial description of each player’s preferences, as it evolves through the dialogue exchange, e.g.:

A: I am prepared to give you rock for wood.
 B: I need clay.

The relation between these utterances is a *Plan-correction*, because *B* corrects *A*’s offer. The Cadilhac algorithm for constructing *B*’s declared preference in context yields that the implied trade is clay for wood. Thus, it computes *which part* of the prior trade was rejected (i.e., getting rock) and captures implicatures about what is accepted; here, the implicature that *B* accepts giving wood. While the *Plan-correction* relation is between whole utterances, the bits that are accepted and rejected are computed on the basis of the semantics of the relation and the semantics (and information structure) of the arguments (Lascarides & Asher, 2009). In contrast, in

A: [I want wood.]_{a₁} [What do you want?]_{a₂}
 B: [I need clay.]_b

two discourse relations are at work *Q-Elab(a₁,a₂)* and *IQAP(a₂,b)*. *Q-Elab(a₁,a₂)* means that *a₂* is a question

whose possible answers all elaborate a plan for achieving the underlying communicative intention of a_1 . *IQAP* (Indirect Question Answer Pair) relates a question (a_2) to a segment (b) whose content, together with shared knowledge, defeasibly implies a direct answer to the question.

Default principles for identifying the scope of implicit endorsements and denials from Lascarides and Asher (2009) also mean that B accepts $Q\text{-Elab}(a_1, a_2)$. In other words, by committing to the question (via the veridical *IQAP*), B implicates a commitment also to the question's *illocutionary effects*, here $Q\text{-Elab}(a_1, a_2)$. This relation, not the content of a_2 itself, yields an interpretation of the question equivalent to *What do you want in exchange for wood?* Thus B's trade offer b is a combination of A's and B's contributions.

A typical case from our pilot corpus (dialogue 19) demonstrates that such interactions can be even more complex:

Hardie: [anyone got ore?] $_{a_1}$ [1-1 for either clay or wheat?] $_{a_2}$
Gytis: [I could give 1 ore for 2 clay] $_b$

Gytis's utterance is related in multiple ways to Hardie's $Plan\text{-Elab}(a_1, a_2)$. First, through b , Gytis commits to a *Plan-correction* on Hardie's $Plan\text{-Elab}(a_1, a_2)$: he chooses clay over wheat and wants 2 ore instead of 1. Second, he commits to $IQAP(a_1, b)$: he indirectly answers he has ore. Finally, he commits to $Plan\text{-Elab}(a_1, b)$: a_1 reveals Hardie's incomplete offer for getting ore, which Gytis's response completes.

Thus, without a principled account of discourse relations like SDRT and an algorithm like the one by Cadilhac et al. (2011) for extracting declared preferences from dialogue content, it will be far from straightforward to map actual dialogue into a representation of the underlying trade, or generate dialogues that express one's intended trade.

Conclusion

We have collected a pilot corpus of trading dialogues within the zero-sum game *The Settlers of Catan*. These dialogues reveal non-cooperative dialogue to be a nuanced affair: while Gricean maxims of cooperativity break down, basic cooperativity is adhered to and speakers often share intentions.

We argued for two interleaved levels of analysis: a game level on which dialogue moves are particular types of game moves; and a dialogue level where the agents' public commitments to coherent interpretations of their contributions are recorded. Both levels are needed for a computationally feasible mapping between utterances and the agents' preferences during language interpretation and production. Thus, for an agent playing *The Settlers of Catan*, dialogue moves are one type of game move the agent can make and must extract from other players' utterances. The semantic representation that is computed by means of coherence relations is the interface to the linguistic subsystems.

Due to the non-cooperativeness of such dialogues, the dialogue moves that players make vary in their sincerity and informativeness. We highlighted moves from our corpus that demonstrate the spectrum, ranging from fully sincere moves through deflections to lies.

In the STAC project (*STRAtegic Conversation*, see acknowledgments), we are currently in the process of collecting the corpus proper, which our research partners in Toulouse will fully annotate with SDRT relations. An implementation of the Cadilhac algorithm will make it possible to extract preferences from discourse structure to semi-automatically annotate the corpus with the agents' preferences (expressed as offers). The relations among preferences, plus the game state (in particular, a player's resources) then serve to determine the game moves. This will be a rich resource for building a dialogue agent that plays *The Settlers of Catan*, empirically grounding the agent's decisions to match what people do.

Acknowledgments

This work is supported by ERC grant 269427 (STAC). We are particularly grateful for discussions with our STAC research partners in Toulouse: Stergos Afantenos, Nicholas Asher, Farah Benamara, Anais Cadilhac, Cedric Dégremont, Philippe Muller, Soumya Paul and Laure Vieu.

References

- Anderson, A. H., Bader, M., Bard, E. G., Boyle, E., Doherty, G., Garrod, S., et al. (1991). The HCRC Map Task Corpus. *Language and Speech*, 34(4), 351–366.
- Asher, N., & Lascarides, A. (2003). *Logics of conversation*. Cambridge: Cambridge University Press.
- Baldridge, J., & Lascarides, A. (2005). Probabilistic head-driven parsing for discourse structure. In *Proc. of the 9th Conf. on Comp. Natural Language Learning* (pp. 96–103).
- Brown, P., & Levinson, S. (1978). *Politeness: Some universals and language usage*. Cambridge University Press.
- Cadilhac, A., Asher, N., Benamara, F., & Lascarides, A. (2011). Commitments to preferences in dialogue. In *Proc. of SIGDIAL 2011* (pp. 204–215).
- Clark, H. H. (1996). *Using language*. Cambridge, MA: Cambridge University Press.
- Grice, P. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.), *Speech acts* (Vol. 3, pp. 41–58). New York: Academic Press.
- Guhe, M., & Bard, E. G. (2008). Adapting referring expressions to the task environment. In *Proc. of the 30th CogSci* (pp. 2404–2409).
- Kowtko, J. C., Isard, S. D., & Doherty, G. M. (1993). *Conversational games within dialogue*. http://groups.inf.ed.ac.uk/hcrc_publications/rp-31.ps.gz. Retrieved on 15 December 2011.
- Lascarides, A., & Asher, N. (2009). Agreement, disputes and commitments in dialogue. *J. of Semantics*, 26(2), 109–158.
- Osborne, M., & Rubinstein, A. (1990). *Bargaining and markets*. Academic Press.
- Solan, L., & Tiersma, P. (2005). *Speaking of crime: The language of criminal justice*. Chicago: U. of Chicago Press.
- Thomas, R. S. (2003). *Real-time decision making for adversarial environments using a plan-based heuristic*. Unpublished doctoral dissertation, Northwestern University.