Developing a corpus of strategic conversation in *The Settlers of Catan*

Stergos Afantenos[1], Nicholas Asher[1], Farah Benamara[1], Anais Cadilhac[1],
Cedric Dégremont[1], Pascal Denis[3], Markus Guhe[2], Simon Keizer[4], Alex Lascarides[2],
Oliver Lemon[4], Philippe Muller[1], Soumya Paul[1], Verena Rieser[4] Laure Vieu[1]*

[1]IRIT, Univ. Toulouse, France; [2]School of Informatics, University of Edinburgh, UK
[3]Mostrare, INRIA NLE, France; [4]Herriot-Watt University, Edinburgh, UK
http://settlers.inf.ed.ac.uk/

Abstract. We describe a dialogue model and an implemented annotation scheme for a pilot corpus of annotated online chats concerning bargaining negotiations in the game *The Settlers of Catan*. We will use this model and data to analyze how conversations proceed in the absence of strong forms of cooperativity, where agents have diverging motives. Here we concentrate on the description of our annotation scheme for negotiation dialogues, illustrated with our pilot data, and some perspectives for future research on the issue.

Keywords: dialogue, game theory, corpus annotation

1 Introduction

The commonly held Gricean view of dialogue agents is that they are fully rational and adhere to the maxims of conversation that entail things like: normally one believes what one says and implicates, and one normally tries to help one's interlocutors achieve their goals [1–5]. The latter requires speakers to adopt shared intentions: their preferences are fully aligned. These assumptions are unwarranted in many conversational settings. Here we propose a different view and present some preliminary findings that buttress it.

We propose an approach to conversation that draws on the principle of rational behaviour from game theory: agents talk to maximize their expected utility (a measure combining belief and preference). So Gricean maxims of conversation hold only when they maximize utility.

While Grice's views still guide a majority of linguists, there are some alternatives. Traum [6] explored a route where cooperativity is determined only by the social conventions guiding conversation, obligations that do not presuppose speakers to adopt each other's goals. This leads to interesting analyses of persuasions and negotiations [7, 8]. For us, the social conventions that are foundational on Traum's account are based on the deeper rational basis of utility [9]. Utility is also the basis for training agents to behave in a certain way through reinforcement learning [10].

The following sections outline our model, describe the corpus we are gathering, and propose an annotation scheme for the corpus. The corpus consists of chat negotiations taking place during games of *The Settlers of Catan*, a popular boardgame.

2 The model

Conversations are dynamic and extensive, with in principle n unbounded number of possible moves: one can, in some sense, always say anything. They also don't have mandatory stopping points. The moves for each player consist in making a discourse contribution, which we characterize in terms of a discourse structure [2], consisting of discourse units linked to each other via discourse relations like *Elaboration*, *Question-Answer-Pair* (*QAP*) and *Explanation*. In particular, discourse relations link contributions from different participants, e.g., if an agent's repsonse to another agent's prior question may link the two contributions with the relation *QAP*.

Conversational participants are alternatively senders (*S*) or receivers (*R*). *S* sends a signal *s*, bearing in mind that receiver *R* has to figure out: (a) the message *m(s)* (i.e., *S*'s public commitments); (b) whether *m(s)* is credible; and (c) what signal *s'* should *R* send in return. *R* now becomes sender and *S* the receiver. With Farrell [11], we assume that at least part of the conventional meaning of the signal is determined prior to game play. In calculating (a), *R* must calculate the public commitments that S has made, including the discourse relations that connect his contribution to the prior discourse. Sometimes these involve strategic considerations: for instance, is S actually replying to the question asked in the prior turn or is she engaged in some other discourse move? If she is answering the question, is this something that *S* cannot plausibly later deny? Asher and Lascarides [12] investigate step (a) in detail and distinguish it from step (b); Asher and Quinley [9] argue that a trust game format is the right one for computing optimal moves in task (c).

The relevant type of game-theoretic analysis at each stage should be carefully identified. In *The Settlers of Catan* players have neither perfect information nor can they have knowledge of complete preferences over all possible sequences of moves. Conversational agents must act based on partial knowledge of other agents' preferences and even their own preferences. Preferences evolve during the conversation; agents revise their plans as they talk and learn new information or forget what has been said. Thus conversations are games of *incomplete information*, *imperfect information* and *imperfect recall*.

Within applications of signaling games, cooperativity and truth-telling is usually assumed or even claimed as a consequence of the framework [13]. In our model, cooperativity comes in several different flavors. Neither sincerity nor full cooperativity involving shared intentions apply when people trade resources to win the game. Consider (1). (1b) is true but misleading because it implicates that *B* doesn't have rock (note that players do not know the full hands of their opponents).

(1)   a.   A: Do you have rock?
        b.   B: I've got lots of wheat   [in fact, B has a rock]

Though outright lies are rare in the chats, speech acts that plausibly involve misleading implicatures are frequent. Cooperativity at the level of sincerity has broken down to some degree: since *A* and *B* both want to win, their preferences cannot be completely aligned. Players also often do not respond to questions like (1a). But cooperativity has not broken down entirely in (1): in in contrast to an assertion like *I won't answer* the move (1b) is one that appears to comply with *A*'s intention—to now an answer. Further, without a minimum of linguistic cooperativity, dialogue wouldn't take place at all, and bargaining transactions could not take place if the bargainers did not cooperate at least to

the extent of agreeing on a trade. In each bargain a coalition cooperates within an overall game of strategic competition.

To support and refine our general model we are collecting a corpus of negotiation dialogues and are particularly interested in the interactions between game strategy and conversation. The intentional model of Grosz and Sidner [3] postulated a close parallelism between domain level intentions and conversational intentions, which we will show doesn't hold. Below we describe our annotation model. As discourse coherence is a crucial basic level of cooperativity, our annotations supply the discourse structure of a dialogue. From this rhetorical structure, we calculate the commitments to preferences that speakers reveal through their utterances, cf. [14]. If these commitments to preferences are judged credible, then dialogue participants can use them to guide their linguistic and non-linguistic actions.

3   Developing a corpus of strategic conversation

To study strategic conversations, we are developing a corpus of online chats between agents playing *The Settlers of Catan*, a competitive win–lose game involving negotiations. The state of the game is recorded and aligned with players' conversations, which allows theorists and annotators to access the players' hands and to identify insincere moves.

In *The Settlers of Catan* each player has some resources of 5 types (ore, wood, wheat, clay, sheep), which they use in different combinations to build roads, settlements and cities, which in turn give them points towards winning the game. Players' hands are hidden to the other players. Although in each turn players can receive some resources, in practice they will not win the game without trading with other players. *The Settlers of Catan* is a positional game with a combinatorial number of possible states, which poses a problem for standard game-theoretic solution concepts. Computationally bounded agents cannot always analytically compute a winning strategy, so they will use scoring functions: computable ways of evaluating board positions. Even the smartest human agents are only able to read ahead a bounded number of positions. As a preliminary model of such agents, we assume that agents are using top down pruning, e.g. reading off all positions within a finite horizon from the current position, and applying backward induction given a common scoring function. Furthermore, agents have limited memory, resulting in uncertainty about the resources opponents are holding and thus the actual position of the game. They are also uncertain about other players' scoring functions. Negotiations in this game thus mirror complex real life negotiations.

We have modified an existing implementation of the game so that players use a chat interface to converse and have annotated 21 games, collected via this online game interface. Each game contains a few dozen self-contained bargaining conversations; a total of ca. 2000 dialogue turns. We have now started an online league (`http://settlers.inf.ed.ac.uk/`) and have collected ca. 40 games in total. We plan to collect hundreds of games with players of different skill levels. The corpus will be made freely available to the research community.

Most of the turns in the chats represent offers, counteroffers, and acceptances or rejections of offers. Many of the negotiation moves in the dialogue, however, do not express completely specified bargaining moves:

(2)   Anybody have any sheep for wheat?

If such moves are to be analyzed strategically, we need to handle them as part of the game strategy. In general the structure of such games is not directly reflected in a typical negotiation dialogue of our corpus. We interpret (2) as a pre-negotiation question or a signal in a sender–receiver game. Such moves are essentially information seeking moves, giving evidence that for real agents *The Settlers of Catan* is a game of imperfect information; in (2), it's incomplete information about other agents' preferences. Our agents' cognitive limitations force them to rely on local scoring functions, thus allowing for Pareto-improving deals in the overall context of a conflicting strategic situation. The fact that our agents play with imperfect information of several kinds and the fact that communication need not be credible might threaten the existence of strategically stable deals. In our corpus many trade offers result in no bargain, which again is not predicted in bargaining games with perfect information. Observed negotiation failure, which would be puzzling under perfect information, is, however, no surprise in our framework.

4   Annotations

Our annotation scheme provides a multi-layered annotation of the chat data, reflecting the multi-dimensional nature of the game. Each turn is segmented into *elementary discourse units* (*EDUs*), and annotators provide a characterization of each EDU in terms of a basic theory of speech acts as well as its effects on the game. One basic feature specifies the turn's addressee, which is sometimes unclear. Other features deal with the domain aspect of the dialogue: what trade offers and counteroffers are made and what offers are (un)successful. We are also interested in comments on strategic play, as they may give information on strategic reasoning. At the next level, we annotate the discourse structure and thus detail the moves in the conversational game as well as the overall discourse structure of an agent's contribution. At a final level, we encode the commitments to preferences, many of which can be computed from the domain level acts.

Discourse annotation, like any semantic annotation, is complex and difficult for naive annotators and usually has low inter-annotator agreement. We therefore divided the annotation task into several layers, ranging from relatively simple (determining the addressee and surface speech act of an EDU) to the complex: identifying discourse relations that connect EDUs, and identifying *complex discourse units* (CDUs)—that is, extended discourse units consisting discourse-related sub-units—that are themselves arguments to discourse relations. The different layers allow us to perform consistency checks on more complex annotation levels using the simpler levels.

4.1   Domain level acts

A pre-annotation automatically segments turns. Each turn contains features specifying what a player enters in the chat window, his resources, the state of the game board and a time stamp. The pre-annotation divides each turn into EDUs but associates no features with them. Prior work has led to a stable view of EDUs and a variety of features exploiting syntactic structure and semantic and discursive function for detecting them.

The annotators specify the type of the *Segment* units, customizing them according to their type of task-level dialogue act: *Offer*, *Counteroffer*, *Accept*, *Refusal*, *Strategic comment* and *Other* units. Annotators also specify the addressee of the act and the surface type of the utterance (*Question*, *Request*, *Assertion*).

Strategic comments label units that comment on strategic moves that have been or should be made. EDUs without a direct role in the game are labelled *Other* on this layer of the annotation model, which also provides some basic discourse structure information and information about preferences. EDUs should be labelled as *Contains preference* if the dialogue move contains an explicit preference, as *Correction* or as *Clarification request*. These serve to constrain and check more detailed rhetorical structure annotations.

Each domain level act has associated features. That is, one encodes the type and quantity of resources that are: declared to be possessed, not possessed, giveable, receivable, or not receivable. For example, an offer may take the form of a question, as in *Markus, will you trade 1 wheat for 1 ore?* For this *Offer* unit the features should specify that *Markus* is the addressee, that there is a receivable *Wheat* resource of quantity 1 and a giveable *Ore* resource of quantity 1. Feature values may be empty when values are unknown, e.g. *Does anyone have any wheat?* is an offer where a receivable resource is specified (*Wheat* without quantity) but addressee and giveable resource are unspecified. Agents typically refine the offer through counteroffers. Boolean combinations of resource names handle cases like: *Anyone have wheat or clay?* Since we record the state of the game, we know the resources of all players, but the players might not. Thus, responses like *Yes* to the above question is an important clue for modeling players' strategizing.

4.2 Rhetorical function

Annotating rhetorical function and preferences is more difficult. At the level of rhetorical function, an EDU as well as an entire turn has one or more discourse relations to other EDUs or CDUs that characterize its discourse role in the text. Prior annotation levels place constraints on this rhetorical level, e.g., if something is labeled as *Accept* on the domain level, an *Acknowledge* relation should hold with particular arguments at this level. For example, an acknowledgment at the level of content in the sense of *I accept your offer* as opposed to a more superficial acknowledgment (*I heard and understood what you said*) is encoded using both *Accept* and *Acknowledge*. Prior annotation levels can thus disambiguate rhetorical moves. Frequent relations in our corpus are *Elaboration*, *Explanation* and *Acknowledgement*.

Elaboration. In our corpus, elaborations typically occur, when an agent makes an offer and then further specifies it. For example:

(3)   a.   Does anyone have any wheat?
          b.   I have ore

In (3b) the speaker is further specifying the offer given in (3a); we represent this as *Elaboration(3a, 3b)*. Sometimes speakers will follow up one question with another one. These questions often stand in an *Elaboration* relation, as in:

(4)   a.   Anyone want sheep for ore?
          b.   2 sheep for 1 ore?

Explanation. Explanation(α, β) holds when β explains why, or gives the cause of, what happened in α. In our corpus Explanations often occur without markers as in (5):

(5)　　[I don't want any sheep.]$_1$ [I'm already drowning in sheep.]$_2$

Acknowledgement. Acknowledgements like *OK, Right, Good, Fine* are common in our corpus. It is often difficult to determine what is acknowledged: an understanding of what was said, an acceptance or agreement, or a signal to change the topic of conversation and move on. Therefore, this feature can be partially specified or unspecified. For example, in (6a–c), it is hard to tell what *OK* acknowledges, but (6d) makes it clear.

(6)　　a.　A: Does anyone have any ore?
　　　　b.　A: I'm offering sheep.
　　　　c.　B: OK.
　　　　d.　B: How about 1 ore for 2 sheep?

(6d) is a *counteroffer* of the offer made in (6a–b). So the acknowledgment in (6c) is an acceptance of what was said and hence of the offer made in (6a–b). So we would annotate this example with *Elaboration(6a, 6b)*, *Acknowledge([6a, 6b], 6c)*, *Q-elab(6c, 6d)* and *Q-elab([6a, 6b], 6d)*. Usually, an acknowledgement is linked to the previous segment.

4.3　Preferences

The final level codes commitments to speakers' preferences, e.g., offers convey complex preferences. Occasionally, explicit preference statements are not related to an offer, such as: *joel fancies a bit of your clay*. Annotators mark the EDU as expressing a preference. Following [14] we will automatically detect such preferences in future work.

5　Preliminary results, example annotation and conclusions

Annotators received training on 22 negotiation dialogues and ca. 560 turns. Inter-annotation agreement at EDU and rhetorical structure levels for this training will be used to refine the guidelines. Ignoring one dialogue without any negotiation and considering 787 instances of doubly-annotated EDUs, kappa is 0.62, a moderate level due to the many *Other* acts. For rhetorical structure, using an exact match criterion of success (easy to compute but harsher than necessary), kappa is 0.45. All figures are very preliminary.

The table on the next page shows an annotated negotiation dialogue with imperfect knowledge and recall amply evident. There are also strategic comments, a persuasion move (49), and underspecified bargaining moves that get specified as more information becomes common knowledge. The discourse structure of the text is relatively straightforward. The preference commitments can be computed from the discourse structure. There are interesting interactions with the game. In this case, both players gain a resource that enables them to build a road segment, and this is not possible for the other players given their resources. This coalition and bargain is optimal, at least without deep look ahead, consistent with the use of local solutions in this complex game.

Speaker	Id	Turn	Dom. function	Rhet. function	Prefs
Euan	47	[And I alt tab back from the tutorial.]_1	OTHER		
		[What's up?]_2	OTHER	Result*(47_1,47_2)	
Joel	48	[do you want to trade?]	OFFER	Q-elab(47_2, 48)	
			⟨Joel,?,?,Euan⟩		
Card.	49	[joel fancies a bit of your clay]	STRAT.-COMMENT	Expl*(48, 49)	Pref(joel)
Joel	50	[yes]	OTHER	Ackn(49, 50)	
Joel	51	[!]	OTHER	Comment(50, 51)	
Euan	52	[Whatcha got?]	COUNTEROFFER	Q-elab([48-50], 52)	
			⟨Euan,?,?,Joel⟩		
Joel	53	[wheat]	HAS-RESOURCES	QAP(52, 53)	
			⟨Joel,wheat⟩		
Euan	54	[I can wheat for clay.]	COUNTEROFFER	Elab([52,53], 54)	
			⟨Euan,wheat,clay,Joel⟩		
Joel	55	[awesome]	ACCEPT(54)	Ackn(54, 55)	

Concluding, we have offered a model of strategic conversation and introduced a corpus of dialogues with negotiations for resources in a win–lose game, *The Settlers of Catan*. We have also described an annotation model for our corpus of negotiation dialogues. Our project aims eventually to automate the annotation process providing a large quantitative analysis for some kinds of strategic conversation.

References

1. Allen, J., Litman, D.: A plan recognition model for subdialogues in conversations. Cognitive Science **11**(2) (1987) 163–200
2. Asher, N., Lascarides, A.: Logics of Conversation. Cambridge University Press (2003)
3. Grosz, B., Sidner, C.: Attention, intentions and the structure of discourse. Computational Linguistics **12** (1986) 175–204
4. Grosz, B., Sidner, C.: Plans for discourse. In P. R. Cohen, J.M., Pollack, M., eds.: Intentions in Communication. MIT Press (1990) 365–388
5. Lochbaum, K.E.: A collaborative planning model of intentional structure. Computational Linguistics **24**(4) (1998) 525–572
6. Traum, D., Allen, J.: Discourse obligations in dialogue processing. In: ACL94, Las Cruces, New Mexico (1994) 1–8
7. Traum, D., Marsella, S., Gratch, J., Lee, J., Hartholt, A.: Multi-party, multi-issue, multi-strategy negotiation for multi-modal virtual agents. In Prendinger, H., Lester, J., Ishizuka, M., eds.: Intelligent Virtual Agents. LNCS 5208. Springer, Berlin (2008) 117–130
8. Traum, D., Swartout, W., Gratch, J., Marsella, S.: A virtual human dialogue model for non-team interaction. In Dybkjr, L., Minker, W., Ide, N., eds.: Recent Trends in Discourse and Dialogue. Springer Netherlands (2008) 45–67
9. Asher, N., Quinley, J.: Begging questions, their answers and basic cooperativity. In: LENLS 8, Japan (2011)
10. Frampton, M., Lemon, O.: Recent research advances in reinforcement learning in spoken dialogue systems. Knowledge Eng. Review **24**(4) (2009) 375–408
11. Farrell, J.: Meaning and credibility in cheap-talk games. Games and Economic Behaviour **5** (1993) 514–531
12. Asher, N., Lascarides, A.: Strategic conversation. (to appear in Semantics and Pragmatics)
13. Lewis, D.: Convention: A Philosophical Study. Harvard University Press (1969)
14. Cadilhac, A., Asher, N., Benamara, F., Lascarides, A.: Commitments to preferences in dialogue. In: Proc. of the 12th Annual SIGDIAL Meeting. (2011) 204–215