Annotating Preferences in Chats for Strategic Games

Anaïs Cadilhac, Nicholas Asher and Farah Benamara
IRIT, CNRS and University of Toulouse
118, route de Narbonne
31062 Toulouse, France
{cadilhac, asher, benamara}@irit.fr

Abstract

This paper describes an annotation scheme for expressions of preferences in on-line chats concerning bargaining negotiations in the online version of the competitive game *Settlers of Catan*.

1 Introduction

Information about preferences is an important part of what is communicated in dialogue. A knowledge of ones own preferences and those of other agents are crucial to decision-making (Arora and Allenby, 1999), strategic interactions between agents (Brainov, 2000) (Hausman, 2000) (Meyer and Foo, 2004). Modeling preferences divides into three subtasks (Brafman and Domshlak, 2009): *preference acquisition*, which extracts preferences from users, *preference modeling* where a model of users' preferences is built using a preference representation language and *preference reasoning* which aims at computing the set of optimal outcomes.

We focus in this paper on a particular instantiation of the first task, extracting preferences from chat turns of actual conversation; and we propose an annotation scheme that is general enough to cover several domains. We extend the annotation scheme of (Cadilhac et al., 2012), which investigates preferences within negotiation dialogues with a common goal like fixing a meeting time (`Verbmobil` (C_V)) or making a hotel or plane reservation (`Booking` (C_B)) to a more complex domain provided by a corpus of on line chats concerning the game *Settlers of Catan*. In *Settlers*, players with opposing strategic interests bargain over scarce resources. Our results show that preferences can be easily annotated by humans and that our scheme adapts relatively easily to different domains.

2 Preferences in game theory

A preference is traditionally a complete ordering by an agent over outcomes. In traditional game theory (Osborne and Rubinstein, 1994), preferences or utilities over outcomes drive rational, strategic decision. They are the terminal states of the game, the end states of complete strategies, which are functions from the set of players $\mathcal{P}$ to the set of actions $\mathcal{A}$; by assigning end states a utility, strategies are thereby also assigned a preference. Game theory postulates that agents calculate their actions based on a common knowledge of all the players' preferences.

In real life, strategic interactions almost always occur under the handicap of various forms of imperfect information. People don't know what other relevant actors are going to do, because they typically don't know what they believe and what they want. In addition, the underlying game is so large that agents with limited computational power can't hope to compute in analytical fashion the optimal actions they should perform.

Because a knowledge of preferences is crucial to informed strategic action, people try to extract information about the preferences of other agents and often provide information about their own preferences when they talk. Almost always this information provides an ordinal definition of preferences, which consists in imposing a ranking over relevant possible outcomes and not a cardinal definition based on

Proceedings of the 6th Linguistic Annotation Workshop, pages 139–143,
Jeju, Republic of Korea, 12-13 July 2012.

numerical values. A *preference relation*, written $\succeq$, is a reflexive and transitive binary relation over elements of Ω. The preference orderings are not necessarily complete, since some candidates may not be comparable for a given agent. Let $o_1,\ o_2 \in \Omega$, $o_1 \succeq o_2$ means that outcome o_1 is equally or more preferred to the decision maker than o_2. *Strict preference* $o_1 \succ o_2$ holds iff $o_1 \succeq o_2$ and not $o_2 \succeq o_1$. The associated *indifference relation* is $o_1 \sim o_2$ if $o_1 \succeq o_2$ and $o_2 \succeq o_1$. Among elements of Ω, some outcomes are acceptable for the agent, i.e. the agent is ready to act in such a way as to realize them, and some outcomes are not. Among the acceptable outcomes, the agent will typically prefer some to others.

3 Data

Settlers of Catan is a competitive win-lose game that involves negotiations. The game is played online, and the state of the game is recorded and aligned with players' conversations. Each player acquires resources, hidden to the other players (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. They can get these resources from rolls of the dice or through trades with the other players. *Settlers* is a positional game with a combinatorial number of possible states. Agents often forget information, with the result that they are uncertain about the resources opponents have as well as about the scoring function other players are using. We have modified the online version of the game so that agents have to converse to carry out trades, using a chat interface. So far we have twenty pilot games involving mostly casual players; each game transcript contains 30 or more self-contained bargaining conversations, for a total of around 2000 dialogue turns.

The data in *Settlers* is more complex than that in C_V or C_B because the dialogues typically involve three or more agents, each with incompatible overall goals. The need to trade requires players to form coalitions in which the participants negotiate the bargain over resources. Thus, there are preferences over which coalition to form, as well as over actions like giving or receiving resources.

Most of the turns in the chats involve negotiation and represent offers, counteroffers, and acceptances or rejections of offers. The example from our corpus in Table 1 involves some creative vocabulary (*alt tab* as a lexical verb) or V ellipsis without a surface antecedent (*I can wheat for clay*) with imperfect knowledge/recall amply evident (Euan's *what's up?*). There are also strategic comments, a persuasion move (49), and underspecified bargaining moves that get specified as more information becomes common knowledge.

While in this paper we concentrate on the annotation of preferences of chat turns, our annotated example shows that our corpus incorporates four layers of annotations: (1) the pre-annotation involves a segmentation of the dialogue into chat lines and the author of each chat line is automatically given, (2) the addressee of the turn, (3) the discourse structure and (4) the players' preferences. The discourse structure of most of the dialogues in *Settlers*, established by consensus, is relatively straightforward. The discourse structure is needed to specify the underspecified elements in our preference annotation.

4 Preference annotation layer

As for C_V and C_B (Cadilhac et al., 2012), our annotation of expressed preferences in each turn involves two steps: identify the set Ω of outcomes, on which the agent's preferences are expressed, and then identify the dependencies between the elements of Ω by using a set of specific non-boolean operators. Preferences in *Settlers* can be atomic preferences or complex preferences.

Atomic preference statements are of the form "I prefer X" where X paradigmatically is identified with a verb phrase ("to trade" or "to give wheat for sheep") or an entire clause describing an action. Sometimes X is identified by a definite noun phrase ("some of your sheep"). The action in question is determined by taking into account of the verb to which X is an argument to specify the action and the full outcome. Agents may also express preferences using questions. That is, in "Do you want to trade?", the agent implicates a preference for trading with the addressee. For negative and wh-interrogatives, the implication is even stronger. A negative preference expresses an unacceptable outcome, i.e. what the agent does not prefer. It can be explicitly expressed ("I have no wood") or inferred from the con-

Speaker	Id	Turn	addressee	Rhet. function
Euan	47	And I alt tab back from the tutorial. What's up?	ALL	
Joel	48	do you want <to trade>_1 ** 1	EUAN	Q-elab(47, 48)
Card.	49	<joel>_1 fancies <a bit of your clay>_2 ** receive(1, Euan, <2,?>)	EUAN	Expl*(48, 49)
Joel	50	yes <>_1 ** 1	CARD	Ackn(49, 50)
Joel	51	!	EUAN	Comment(50, 51)
Euan	52	Whatcha got? <>_1 ** 1	JOEL	Q-elab([48-50], 52)
Joel	53	<wheat>_1 or <wood>_2 ** offer(Joel, Euan, <1,?> $\triangledown$ <2,?>)	EUAN	QAP(52, 53)
Euan	54	I can <wheat>_1 for <1 clay>_2. ** receive(Euan, Joel, <1,?>) $\mapsto$ offer(Euan, Joel, <2,1>	JOEL	Elab([52,53], 54)
Joel	55	awesome <>_1 ** 1	EUAN	Ackn(54, 55)

Table 1: Example negotiation with discourse annotation

text ("no"), which means that the player rejects an offer and thus does not want to trade.

Complex preference statements express dependencies between outcomes (Boutilier et al., 2004)). Among the possible combinations, we find conjunctions, disjunctions and conditionals. We examine operations over outcomes and suppose a language with non-boolean operators *&*, $\triangledown$ and $\mapsto$ respectively, taking outcome expressions as arguments. With conjunctions of preferences, as in "Can I have one sheep and one ore?", the agent expresses two preferences (respectively over the acceptable outcomes of his getting one sheep and his getting one ore) that he wants to satisfy and he prefers to have one of them if he cannot have both. The semantics of a disjunctive preference is a free choice one. For example in "I can give wheat or sheep", the agent states that giving sheep or wheat is an acceptable outcome and he is indifferent between the choice of the outcomes. Finally, some turns express conditional among preferences. In our corpus, all offers and counteroffers express conditional preferences; "I can wheat for sheep", there are two preferences: one for receiving sheep, and, given the preference for receiving sheep, one for the giving of wheat.

In *Settlers*, an outcome X can play a role in several actions: a preference for the speaker's receiving or offering the resource X, a preference for a trade, a preference for performing the action X, etc. To specify these different actions, we use, in addition to the vocabulary of our previous annotation language, two functions: *receive(o, a, <r,q>)* and *offer(o, a, <r,q>)* such that: o is the preference owner, a is the addressee, r is the resource and q is the quantity of the resource needed (or offered). If some of these arguments are underspecified, we put *?*. Outcomes, which are closed under our non-boolean operators, can specify one or more arguments of our new predicates, or range over an action description. In addition, we have decided to annotate anaphoric and unspecified bargaining moves using an empty outcome (50). Table1 shows how the example is annotated (*<outcome>*_i indicates outcome number i in the turn; preference annotation is given after **).

5 Inter-annotator agreements

Two judges manually annotated two games from our corpus of 20 *Settlers* dialogues using the previously described annotation scheme. The two games contain 74 bargaining conversations for a total of 980 turns with 632 outcomes, 147 of which are unacceptable (*not* operator). There are 20 instances of &, 27 of $\triangledown$ and 80 of $\mapsto$. We computed four inter-annotator agreements on: (a) outcome identification, (b) outcome acceptance, (c) outcome attachment and (d) operator identification.

For (a), we compute a *lenient* match between annotations using Cohen's Kappa (i.e. there is an overlap between their text spans as in "sheep" and "some sheep"). We obtain a Kappa of 0.92 for *Settlers* while for both C_V and C_B we obtained a Kappa of 0.85. As in C_V and C_B, the main case of disagreement concerns redundant preferences which we decided not to keep in the gold standard because the player just wants to insist by repeating already stated preferences. In *Settlers*, we observed four additional cases of disagreement: (1) sometimes judges do not annotate underspecified preferences which are often used to introduce new, to make current preferences more precise or to accept preferences. Hence, we decided to annotate them in the gold standard. (2)

annotators sometimes forget to annotate a resource when it is lexicalized by a synonym (as "dolly" and "sheep"), (3) annotators often fail to decide if the action is about receiving or offering a resource (as in "ore for clay") mainly because the same lexicalizations do not always lead to the same actions, (4) judges do not always annotate preferences that are not directly related to the action of trading, offering or receiving a resource.

For (b), the aim is to compute the agreement on the *not* operator, that is if an outcome is acceptable, as in Dave: "I will give <you>_1 <wheat>_2", or unacceptable, as in Tomm: "No <ore>_1, sorry". We get a Kappa of 0.97 for *Settlers* while we obtained a Kappa of 0.90 for C_V and 0.95 for C_B. As in C_V and C_B, the main case of disagreement concerns negations that are inferred from the context.

For (c), since the structure of the bargaining packages outcomes in a very predictable way, it is quite intuitive, and simpler than for C_V and C_B, to decide how options are integrated in the preference annotation in *Settlers* which includes functions (offer and receive). We computed annotator agreement using the F-score measure because this task involves structure building as in "Joel wants to trade wheat for clay, or wheat for ore", which gives us: *(receive(Joel,?,<clay,?>)* $\mapsto$ *offer(Joel,?,<wheat,?>))* $\triangledown$ *(receive(Joel,?,<ore,?>)* $\mapsto$ *offer(Joel,?,<wheat,?>))*. The agreement concerns turns that contain at least three outcomes and was computed on the previously built gold standard once annotators discussed cases of outcome identification disagreements. We obtain an agreement of 93% for C_V, 82% for C_B and perfect agreement for *Settlers*.

Finally, in our *Settlers* corpus, the most frequent operators are not and $\mapsto$ because the main purpose of the players in this corpus is to propose, accept or reject a trade. The other two operators & and $\triangledown$ are equally split. The most frequently used binary operators were $\mapsto$ in C_V and & and $\mapsto$ in C_B. The Cohen's Kappa for (d), averaged over all the operators, is 0.93 for C_V, 0.75 for C_V and 0.95 for *Settlers*. In C_V and C_B, we observed two main cases of disagreement: between $\triangledown$ and &, and between & and $\mapsto$. These cases were more frequent for C_B, accounting for the lower Kappa there than for C_V. In *Settlers*, the main case of disagreement concerns the confusion between $\triangledown$ and &. The high agreement on $\mapsto$ reflects the fact that $\mapsto$ occurs in the description of an offer which is easy to annotators to spot.

The same linguistic realizations do not always lead to the same annotations. The coordinating conjunction "or" is a strong predictor for recognizing a disjunction of preferences, at least when "or" is clearly outside of the scope of a negation. In C_V and C_B, the coordinating conjunction "and" can also give a disjunction, especially when it is used to link two acceptable outcomes that are both of a single type (e.g., day, type of room) between which an agent wants to choose a single realization. In *Settlers*, the connector "and" generally links two outcomes that the agent wants to satisfy simultaneously and involves a conjunction of preferences, as in Dave: "I can give <you>_1 <one wheat>_2 and <ore>_3 for <wood>_4" where we have: *receive(Dave, 1, <4, ?>)* $\mapsto$ *offer(Dave, 1, <2, 1> & <3, ?>)*. When "and" links two outcomes and one at least is unacceptable, it gives a conjunction of preferences, as in Dave: "I dont have <any ore>_1, but i do have <plenty clay>_2" where we have: *not offer(Dave, ?, <1, ?>) & offer(Dave, ?, <2, ?>)*.

6 Conclusion and Future Work

We have proposed a linguistic approach to preference acquisition that aims to infer preferences from chats concerning bargaining negotiations in an online version of the game *Settlers of Catan*. The described annotation scheme extends the scheme of (Cadilhac et al., 2012), which investigated preferences within negotiation dialogues with a common goal like fixing a meeting time or making a hotel or plane reservation to the more complex domain of *Settlers*, where the types of actions were more diverse. The next step is to automate the process of preference extraction from turns or elementary discourse units using NLP methods, while at the same time pursuing the annotation and automation of the discourse parsing process. We also plan to study the evolution of these preferences *vis à vis* strategies of the underlying game, giving us an insight into how humans strategize within complex games like *Settlers* or real life situations, for which standard game theoretic solution concepts are not feasible for limited agents like us.

References

Neeraj Arora and Greg M. Allenby. 1999. Measuring the influence of individual preference structures in group decision making. *Journal of Marketing Research*, 36:476–487.

Craig Boutilier, Craig Brafman, Carmel Domshlak, Holger H. Hoos, and David Poole. 2004. Cp-nets: A tool for representing and reasoning with conditional *ceteris paribus* preference statements. *Journal of Artificial Intelligence Research*, 21:135–191.

Ronen I. Brafman and Carmel Domshlak. 2009. Preference handling - an introductory tutorial. *AI Magazine*, 30(1):58–86.

Sviatoslav Brainov. 2000. The role and the impact of preferences on multiagent interaction. In *Proceedings of ATAL*, pages 349–363. Springer-Verlag.

Anaïs Cadilhac, Nicholas Asher, and Farah Benamara. 2012. Annotating preferences in negotiation dialogues. In *Proceedings of *SEM*.

Daniel M. Hausman. 2000. Revealed preference, belief, and game theory. *Economics and Philosophy*, 16(01):99–115.

Thomas Meyer and Norman Foo. 2004. Logical foundations of negotiation: Strategies and preferences. In *In Proceedings of the Ninth International Conference on Principles of Knowledge Representation and Reasoning (KR04*, pages 311–318.

Martin Osborne and Ariel Rubinstein. 1994. *A Course in Game Theory*. MIT Press.