

Annotating Preferences in Negotiation Dialogues

Anais 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

Modeling user preferences is crucial in many real-life problems, ranging from individual and collective decision-making to strategic interactions between agents and game theory. Since agents do not come with their preferences transparently given in advance, we have only two means to determine what they are if we wish to exploit them in reasoning: we can infer them from what an agent says or from his nonlinguistic actions. In this paper, we analyze how to infer preferences from dialogue moves in actual conversations that involve bargaining or negotiation. To this end, we propose a new annotation scheme to study how preferences are linguistically expressed in two different corpus genres. This paper describes the annotation methodology and details the inter-annotator agreement study on each corpus genre. Our results show that preferences can be easily annotated by humans.

1 Introduction

Modeling user preferences is crucial in many real-life problems, ranging from individual and collective decision-making (Arora and Allenby, 1999) to strategic interactions between agents (Brainov, 2000) and game theory (Hausman, 2000). A web-based recommender system can, for example, help a user to identify (among an optimal ranking) the product item that best fits his preferences (Burke, 2000). Modeling preferences can also help to find some compromise or consensus between two or more agents having different goals during a negotiation (Meyer and Foo, 2004).

Working with preferences involves 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 the first task.

Handling preferences is not easy. First, specifying an ordering over acceptable outcomes is not trivial especially when multiple aspects of an outcome matter. For instance, choosing a new camera to buy may depend on several criteria (e.g. battery life, weight, etc.), hence, ordering even two outcomes (cameras) can be cognitively difficult because of the need to consider trade-offs and dependencies between the criteria. Second, users often lack complete information about preferences initially. They build a partial description of agents' preferences that typically changes over time. Indeed, users often learn about the domain, each others' preferences and even their own preferences during a decision-making process. Since agents don't come with their preferences transparently given in advance, we have only two means to determine what they are if we wish to exploit them in reasoning: we can infer them from what an agent says or from his nonlinguistic actions. In this paper, we analyze how to infer preferences from dialogue moves in actual conversations that involve bargaining or negotiation.

Within the Artificial Intelligence community, preference acquisition from nonlinguistic actions has been performed using a variety of specific tasks, including preference learning (Fürnkranz and

Hüllermeier, 2011) and preference elicitation methods (Chen and Pu, 2004) (such as query learning (Blum et al., 2004), collaborative filtering (Su and Khoshgoftaar, 2009) and qualitative graphical representation of preferences (Boutilier et al., 1997)). However, these tasks don't occur in actual conversations about negotiation. We are interested in how agents learn about preferences from actual conversational turns in real dialogue (Edwards and Barron, 1994), using NLP techniques.

To this end, we propose a new annotation scheme to study how preferences are linguistically expressed in dialogues. The annotation study is performed on two different corpus genres: the Verbmobil corpus (Wahlster, 2000) and a booking corpus, built by ourselves. This paper describes the annotation methodology and details the inter-annotator agreement study on each corpus genre. Our results show that preferences can be easily annotated by humans.

2 Background

2.1 What are preferences?

A preference is commonly understood as an ordering by an agent over outcomes, which are understood as actions that the agent can perform or goal states that are the direct result of an action of the agent. For instance, an agent's preferences may be defined over actions like *buy a new car* or by its end result like *have a new car*. The outcomes over which a preference is defined will depend on the domain or task.

Among these outcomes, some 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. Our aim is not to determine the most preferred outcome of an agent but follows rather the evolution of their commitments to certain preferences as the dialogue proceeds. To give an example, if an agent proposes to meet on a certain day X and at a certain time Y, we learn that among the agent's acceptable outcomes is a meeting on X at Y, even if this is not his most preferred outcome. We are interested in an ordinal definition of preferences, which consists in imposing a ranking over all (relevant) possible outcomes and not a cardinal definition which is based on numerical values that allow

comparisons.

More formally, let Ω be a set of possible outcomes. 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 by a given agent. Given the two outcomes o_1 and o_2 , $o_1 \succeq o_2$ 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 $o_2 \succeq o_1$.

2.2 Preferences vs. opinions

It is important to distinguish preferences from opinions. While opinions are defined as a point of view, a belief, a sentiment or a judgment that *an agent may have about an object or a person*, preferences, as we have defined them, involve an ordering on behalf of an agent and thus are *relational and comparative*. Hence, opinions concern absolute judgments towards objects or persons (positive, negative or neutral), while preferences concern relative judgments towards actions (preferring them or not over others). The following examples illustrate this:

- (a) The movie is not bad.
- (b) The scenario of the first season is better than the second one.
- (c) I would like to go to the cinema. Let's go and see *Madagascar 2*.

(a) expresses a direct positive opinion towards the movie but we do not know if this movie is the most preferred. (b) expresses a comparative opinion between two movies with respect to their shared features (scenarios) (Ganapathibhotla and Liu, 2008). If actions involving these movies (e.g. seeing them) are clear in the context, such a comparative opinion will imply a preference, ordering the first season scenario over the second. Finally, (c) expresses two preferences, one depending on the other. The first is that the speaker prefers to go to the cinema over other alternative actions; the second is, given that preference, that he wants to see *Madagascar 2* over other possible movies.

Reasoning about preferences is also distinct from reasoning about opinions. An agent's preferences

determine an order over outcomes that predicts how the agent, if he is rational, will act. This is not true for opinions. Opinions have at best an indirect link to action: I may hate what I'm doing, but do it anyway because I prefer that outcome to any of the alternatives.

3 Data

Our data come from two corpora: one already-existing, *Verbmobil* (C_V), and one that we created, *Booking* (C_B).

The first corpus is composed of 35 dialogues randomly chosen from the existing corpus *Verbmobil* (Wahlster, 2000), where two agents discuss on when and where to set up a meeting. Here is a typical fragment:

π_1 *A*: Shall we meet sometime in the next week?

π_2 *A*: What days are good for you?

π_3 *B*: I have some free time on almost every day except Fridays.

π_4 *B*: Fridays are bad.

π_5 *B*: In fact, I'm busy on Thursday too.

π_6 *A*: Next week I am out of town Tuesday, Wednesday and Thursday.

π_7 *A*: So perhaps Monday?

The second corpus was built from various English language learning resources, available on the Web (e.g., www.bbc.co.uk/worldservice/learningenglish). It contains 21 randomly selected dialogues, in which one agent (the customer) calls a service to book a room, a flight, a taxi, etc. Here is a typical fragment:

π_1 *A*: Northwind Airways, good morning. May I help you?

π_2 *B*: Yes, do you have any flights to Sydney next Tuesday?

π_3 *A*: Yes, there's a flight at 16:45 and one at 18:00.

π_4 *A*: Economy, business class or first class ticket?

π_5 *B*: Economy, please.

Our approach to preference acquisition exploits discourse structure and aims to study the impact of discourse for extracting and reasoning on preferences. Cadilhac et al. (2011) show how to compute automatically preference representations for a whole stretch of dialogue from the preference representations for elementary discourse units. Our annotation here concentrates on the commitments to pref-

erences expressed in elementary discourse units or EDUs. We analyze how the outcomes and the dependencies between them are linguistically expressed by performing, on each corpus, a two-level annotation. First, we perform a segmentation of the dialogue into EDUs. Second, we annotate preferences expressed by the EDUs.

The examples above show the effects of segmentation. Each EDU is associated with a label π_j . For *Verbmobil*, we rely on the already available discourse annotation of Baldridge and Lascarides (2005). For *Booking*, the segmentation was made by consensus.

We detail, in the next section, our preference annotation scheme.

4 Preference annotation scheme

To analyze how preferences are linguistically expressed in each EDU, we must: (1) identify the set Ω of outcomes, on which the agent's preferences are expressed, and (2) identify the dependencies between the elements of Ω by using a set of specific operators, i.e. identifying the agent's preferences on the stated outcomes. Consider the segment "Let's meet Thursday or Friday". We have $\Omega = \{meet\ Thursday, meet\ Friday\}$ where outcomes are linked by a disjunction that means the agent is ready to act for one of these outcomes, preferring them equally.

Within an EDU, preferences can be expressed in different ways. They can be atomic preference statements or complex preference statements.

4.1 Atomic preferences

Atomic preference statements are of the form "I prefer X", "Let's X", or "We need X", where X describes an outcome. X may be a definite noun phrase ("Monday", "next week", "almost every day"), a prepositional phrase ("at my office") or a verb phrase ("to meet"). They can be expressed within comparatives and/or superlatives ("a cheaper room" or "the cheapest flight").

Preferences can also be expressed in an indirect way using questions. Although not all questions entail that their author commits to a preference, in many cases they do. That is, if *A* asks "can we meet next week?" he implicates a preference for meeting. For negative and wh-interrogatives, the implication

is even stronger. Expressions of sentiment or politeness can also be used to indirectly introduce preferences. In *Booking*, the segment “economy please” indicates the agent’s preference to be in an economy class.

EDUs can also express preferences via free-choice modalities; “I am free on Thursday” or “I can meet on Thursday” tells us that Thursday is a possible day to meet, it is an acceptable outcome.

A negative preference expresses an unacceptable outcome, i.e. what the agent does not prefer. Negative preference can be expressed explicitly with negation words (“I don’t want to meet on Friday”) or inferred from the context (“I am busy on Monday”).

While the logical form of an atomic preference statement is something of the form $Pref(X)$, we abbreviate this in the annotation language, using just the outcome expression X to denote that the agent prefers X to the alternatives, i.e. $X \succ \bar{X}$. If X is an unacceptable outcome, we use the non-boolean operator *not* to denote that the agent prefers *not* X to other alternatives, i.e. $\bar{X} \succ X$. In our *Verbmobil* annotation, X is typically an NP denoting a time or place; X as an outcome is thus shorthand for *meet on X* or *meet at X*. For *Booking*, X is short for *reserve* or *book X*.

4.2 Complex preferences

Preference statements can also be complex, expressing dependencies between outcomes. Borrowing from the language of *conditional preference networks* or CP-nets (Boutilier et al., 2004), we recognize that some preferences may depend on another action. For instance, given that I have chosen to eat fish, I will prefer to have white wine over red wine—something which we express as *eat fish : drink white wine* \succ *drink red wine*.

Among the possible combinations, we find conjunctions, disjunctions and conditionals. We examine these conjunctive, disjunctive and conditional operations over outcomes and suppose a language with non-boolean operators $\&$, ∇ and \mapsto taking outcome expressions as arguments.

With conjunctions of preferences, as in “Could I have a breakfast and a vegetarian meal?” or in “Mondays and Fridays are not good?”, the agent expresses two preferences (respectively over the ac-

ceptable outcomes breakfast and vegetarian meal and the non acceptable outcomes not Mondays and not Fridays) that he wants to satisfy and he prefers to have one of them if he can not have both. Hence $o_1 \& o_2$ means $o_1 \succ \bar{o}_1$ and $o_2 \succ \bar{o}_2$.

The semantics of a disjunctive preference is a free choice one. For example in “either Monday or Tuesday is fine for me” or in “I am free Monday and Tuesday”, the agent states that either Monday or Tuesday is an acceptable outcome and he is indifferent between the choice of the outcomes. Hence $o_1 \nabla o_2$ means $o_2 : o_1 \sim \bar{o}_1, \bar{o}_2 : o_1 \succ \bar{o}_1$ and $o_1 : o_2 \sim \bar{o}_2, \bar{o}_1 : o_2 \succ \bar{o}_2$.

Finally, some EDUs express conditional among preferences. For example, in the sentence “What about Monday, in the afternoon?”, there are two preferences: one for the day Monday, and, given the Monday preference, one for the time afternoon (of Monday), at least for one syntactic reading of the utterance. Hence $o_1 \mapsto o_2$ means $o_1 \succ \bar{o}_1$ and $o_1 : o_2 \succ \bar{o}_2$.

For each EDU, annotators identify how outcomes are expressed and then indicate if the outcomes are acceptable, or not, using the operator *not* and how the preferences on these outcomes are linked using the operators $\&$, ∇ and \mapsto .

4.3 Example

We give below an example of how some EDUs are annotated. $\langle o \rangle_i$ indicates that o is the outcome number i in the EDU, the symbol $//$ is used to separate the two annotation levels and brackets indicate how outcomes are attached.

π_1 : \langle Tuesday the sixteenth \rangle_{-1} I got class \langle from nine to twelve \rangle_{-2} // $1 \mapsto$ not 2

π_2 : What about \langle Friday afternoon \rangle_{-1} , \langle at two thirty \rangle_{-2} or \langle three \rangle_{-3} , // $1 \mapsto (2 \nabla 3)$

π_3 : \langle The room with balcony \rangle_{-1} should be equipped \langle with a queen size bed \rangle_{-2} , \langle the other one \rangle_{-3} \langle with twin beds \rangle_{-4} , please. // $(1 \mapsto 2) \& (3 \mapsto 4)$

In π_1 , the annotation tells us that we have two outcomes and that the agent prefers outcome 1 over any other alternatives and given that, he does not prefer outcome 2. In π_2 , the annotation tells us that the agent prefers to have one of outcome 2 and outcome 3 satisfied given that he prefers outcome 1. In this example, the free choice between outcome 2 and

outcome 3 is lexicalized by the coordinating conjunction “or”. On the contrary, π_3 is a more complex example where there is no discursive marker to find that the preference operator between the couples of outcomes 1 and 2 on one hand, and 3 and 4 on the other hand, is the conjunctive operator &.

5 Inter-annotator agreements

Our two corpora (*Verbmobil* and *Booking*) were annotated by two annotators using the previously described annotation scheme. We performed an intermediate analysis of agreement and disagreement between the two annotators on two *Verbmobil* dialogues. Annotators were thus trained only for *Verbmobil*. The aim is to study to what extent our annotation scheme is genre dependent. The training allowed each annotator to understand the reason of some annotation choices. After this step, the dialogues of our corpora have been annotated separately, discarding those two dialogues. Table 1 presents some statistics about the annotated data in the gold standard.

	C_V	C_B
No. of dialogues	35	21
No. of outcomes	1081	275
No. of EDUs with outcomes	776	182
% with 1 outcome	71%	70%
% with 2 outcomes	22%	19%
% with 3 or more outcomes	8%	11%
No. of unacceptable outcomes (not)	266	9
No. of conjunctions (&)	56	31
No. of disjunctions (∇)	75	29
No. of conditionals (\mapsto)	184	37

Table 1: Statistics for the two corpora.

We compute four inter-annotator agreements: on outcome identification, on outcome acceptance, on outcome attachment and finally on operator identification. Table 2 summarizes our results.

5.1 Agreements on outcome identification

Two inter-annotator agreements were computed using Cohen’s Kappa. One based on an *exact* matching between two outcome annotations (i.e. their corresponding text spans), and the other based on a *le-*

	C_V	C_B
Outcome identification (Kappa)	exact : 0.66 lenient : 0.85	
Outcome acceptance (Kappa)	0.90	0.95
Outcome attachment (F-measure)	93%	82%
Operator identification (Kappa)	0.93	0.75

Table 2: Inter-annotator agreements for the two corpora.

nient match between annotations (i.e. there is an overlap between their text spans as in “2p.m” and “around 2p.m”). This approach is similar to the one used by Wiebe, Wilson and Cardie (2005) to compute agreement when annotating opinions in news corpora. We obtained an exact agreement of 0.66 and a lenient agreement of 0.85 for both corpus genres.

We made the gold standard after discussing cases of disagreement. We observed four cases. The first one concerns redundant preferences which we decided not to keep in the gold standard. In such cases, the second EDU π_2 does not introduce a new preference, neither does it correct the preferences stated in π_1 ; rather, the agent just wants to insist by repeating already stated preferences, as in the following example:

π_1 A: Thursday, Friday, and Saturday I am out.

π_2 A: So those days are all out for me,

The second case of disagreement comes from anaphora which are often used to introduce new, to make more precise or to accept preferences. Hence, we decided to annotate them in the gold standard. Here is an example:

π_1 A: One p.m. on the seventeenth?

π_2 B: That sounds fantastic.

The third case of disagreement concerns preference explanation. We chose not to annotate these expressions in the gold standard because they are used to explain already stated preferences. In the following example, one judge annotated “from nine to twelve” to be expressions of preferences while the other did not :

π_1 A: Monday is really not good,

π_2 A: I have got class from nine to twelve.

Finally, the last case of disagreement comes from preferences that are not directly related to the action of fixing a date to meet but to other actions, such as having lunch, choosing a place to meet, etc. Even though those preferences were often missed by annotators, we decided to keep them, when relevant.

5.2 Agreements on outcome acceptance

The aim here is to compute the agreement on the *not* operator, that is if an outcome is acceptable, as in “<Mondays>_1 are good // 1”, or unacceptable, as in “<Mondays>_1 are not good // not 1”. We get a Cohen’s Kappa of 0.9 for *Verbmobil* and 0.95 for *Booking*. The main case of disagreement concerns anaphoric negations that are inferred from the context, as in π_2 below where annotators sometimes fail to consider “in the morning” as unacceptable outcomes:

π_1 A: Tuesday is kind of out,

π_2 A: Same reason in the morning

Same case of disagreement in this example where “Monday” is an unacceptable outcome:

π_1 well, I am, busy <in the afternoon of the twenty sixth>_1, // not 1

π_2 that is <Monday>_1 // not 1

5.3 Agreements on outcome attachment

Since this task involves structure building, we compute the agreement using the F-score measure. The agreement was computed on the previously built gold standard once annotators discussed cases of outcome identification disagreements. We compare how each outcome is attached to the others within the same EDU. This agreement concerns EDUs that contain at least three outcomes, that is 8% of EDUs from *Verbmobil* and 11% of EDUs from *Booking*. When comparing annotations for the example π_1 below, there is three errors, one for outcome 2, one for 3 and one for 4.

π_1 <for the next week>_1 the only days I have open are <Monday>_2 or <Tuesday>_3 <in the morning>_4.

- Annotation 1 : $1 \mapsto (2 \nabla (3 \mapsto 4))$
- Annotation 2 : $1 \mapsto ((2 \nabla 3) \mapsto 4)$

We obtain an agreement of 93% for *Verbmobil* and 82% for *Booking*.

5.4 Agreements on outcome dependencies

Finally, we compute the agreements for each couple of outcomes on which annotators agreed about how they are attached.

In *Verbmobil*, the most frequently used binary operator is \mapsto . Because the main purpose of the agents in this corpus is to schedule an appointment, the preferences expressed by the agents are mainly focused on concepts of time and there are many conditional preferences since it is common that preferences on specific concepts depend on more broad temporal concepts. For example, preferences on hours are generally conditional on preferences on days. In *Booking*, there are almost as many $\&$ as \mapsto because independent and dependent preferences are more balanced in this corpus. The agents discuss preferences about various criteria that are independent. For example, to book a hotel, the agent express his preferences towards the size of the bed (single or double), the quality of the room (smoker or nonsmoker), the presence of certain conveniences (TV, bathtub), the possibility to have breakfast in his room, etc. Within an EDU, such preferences are often expressed in different sentences (compared to *Verbmobil* where segments’ lengths are smaller) which lead annotators to link those preferences with the operator $\&$. Conditionals between preferences hold when decision criteria are dependent. For example, the preference for having a vegetarian meal is conditional on the preference for having lunch. There also are conditionals between temporal concepts, for example, to choose the time of a flight.

Table 3 shows the Kappa for each operator on each corpus genre. The Cohen’s Kappa, averaged over all the operators, is 0.93 for *Verbmobil* and 0.75 for *Booking*. We observe two main cases of disagreement: between ∇ and $\&$, and between $\&$ and \mapsto . These cases are more frequent for *Booking* mainly because annotators were not trained on this corpus. This is why the Kappa was lower than for *Verbmobil*. We discuss below the main two cases of disagreement.

Confusion between ∇ and $\&$. The same linguistic realizations do not always lead to the same operator. For instance, in “<Monday>_1 and <Wednesday>_2 are good” we have $1 \nabla 2$ whereas in “<Monday>_1 and <Wednesday>_2 are not

	C_V	C_B
$\&$	0.90	0.66
∇	0.97	0.89
\mapsto	0.92	0.71

Table 3: Agreements on binary operators.

good” or in “I would like a <single room>.₁ and a <taxi>.₂” we have respectively *not 1 & not 2* and *1 & 2*.

The coordinating conjunction “**or**” is a strong predictor for recognizing a disjunction of preferences, at least when the “or” is clearly outside of the scope of a negation¹, as in the examples below (in π_1 , the negation is part of the wh-question, and not boolean over the preference):

- π_1 Why don’t we <meet, either Thursday the first>.₁, or <Thursday the eighth>.₂ // *1 ∇ 2*
- π_2 Would you like <a single>.₁ or <a double>.₂? // *1 ∇ 2*

The coordinating conjunction “**and**” is also a strong indication, especially when it is used to link two acceptable outcomes that are both of a single type (e.g., day of the week, time of day, place, type of room, etc.) between which an agent wants to choose a single realization. For example, in *Verbmobil*, agents want to fix a single appointment so if there is a conjunction “and” between two temporal concepts of the same level, it is a disjunction of preference (see π_3 below). It is also the case in *Booking* when an agent wants to book a single plane flight (see π_4).

- π_3 <Monday>.₁ and <Tuesday>.₂ are good for me // *1 ∇ 2*
- π_4 You could <travel at 10am.>.₁, <noon>.₂ and <2pm>.₃ // *1 ∇ (2 ∇ 3)*

The acceptability modality distributes across the conjoined NPs to deliver something like $\diamond(\textit{meet Monday}) \wedge \diamond(\textit{meet Tuesday})$ in modal logic (clearly acceptability is an existential rather than universal modality), and as is known from studies of free choice modality

¹When there is a propositional negation over the disjunction as in “I don’t want sheep or wheat”, which occurs frequently in a corpus in preparation, we no longer have a disjunction of preferences.

(Schulz, 2007), such a conjunction translates to $\diamond(\textit{meet Monday} \vee \textit{meet Tuesday})$, which expresses our free choice disjunction of preferences, $o_1 \nabla o_2$.

On the other hand, when the conjunction “and” links two outcomes referring to a single concept that are not acceptable, it gives a conjunction of preferences, as in π_5 . Once again thinking in terms of modality is helpful. The “not acceptable” modality distributes across the conjunction, this gives something like $\Box\neg o_1 \wedge \Box\neg o_2$ (where \neg is truth conditional negation) which is equivalent to $\Box(\neg o_1 \wedge \neg o_2)$, i.e. *not o₁ & not o₂* and not equivalent to $\Box(\neg o_1 \vee \neg o_2)$, i.e. *not o₁ ∇ not o₂*.

The connector “and” also involves a conjunction of preferences when it links two independent outcomes that the agent wants to satisfy simultaneously. For example, in π_6 , the agent wants to book two hotel rooms, and so the outcomes are independent. In π_7 , the agent expresses his preferences on two different features he wants for the hotel room he is booking.

- π_5 <Thursday the thirtieth>.₁, and <Wednesday the twenty ninth>.₂ are, booked up // *not 1 & not 2*
- π_6 Can I have one room< with balcony>.₁ and <one without balcony>.₂? // *1 & 2*
- π_7 <Queen>.₁ and <nonsmoking>.₂ // *1 & 2*

Confusion between & and \mapsto . In this case, disagreements are mainly due to the difficulty for annotators to decide if preferences are dependent, or not. For example, in “I have a meeting <starting at three>.₁, but I could meet <at one o’clock>.₂”, one annotator put *not 1 \mapsto 2* meaning that the agent is ready to meet at one o’clock because he can not meet at three, while the other annotated *not 1 & 2* meaning that the agent is ready to meet at one o’clock independently of what it will do at three.

Some connectors introduce contrast between the preferences expressed in a segment as “**but**”, “**although**” and “**unless**”. In the annotation, we can model it thanks to the operator \mapsto . When it is used between two conflicting values, it represents a correction. Thus, the annotation $o_1 \mapsto \textit{not } o_1$ means we need to replace in our model of preferences $o_1 \succ \overline{o_1}$ by $\overline{o_1} \succ o_1$. And vice versa for *not o₁ \mapsto o₁*.

- π_8 I have class <on Monday>.₁, but, <any time, after one or two>.₂ I am free. // *not 1 \mapsto (1 \mapsto 2)*

- π_9 <Friday>_1 is a little full, although there is some possibility, <before lunch>_2 // not $1 \mapsto (1 \mapsto 2)$
- π_{10} we're full <on the 22nd>_1, unless you want <a smoking room>_2 // not $1 \mapsto (1 \mapsto 2)$

However, it is important to note that the coordinating conjunction “but” does not always introduce contrast, as in the example below, where it introduces a conjunction of preferences.

- π_{11} I am busy <on Monday>_1, but <Tuesday afternoon>_2, sounds good // not $1 \& 2$

The subordinating conjunctions “if”, “because” and “so” are indications for detecting conditional preferences. The preferences in the main clause depend on the preferences in the subordinate clause (if-clause, because-clause, so-clause), as in the examples below.

- π_{12} so if we are going to be able to meet <that, last week in January>_1, it is going have to be <the, twenty fifth>_2 // $1 \mapsto 2$
- π_{13} <the twenty eighth>_1 I am free, <all day>_2, if you want to go for <a Sunday meeting>_3 // $3 \mapsto (2 \mapsto 1)$
- π_{14} it is going to have to be <Wednesday the third>_1 because, I am busy <Tuesday>_2 // not $2 \mapsto 1$
- π_{15} I have a meeting <from eleven to one>_1, so we could, meet <in the morning from nine to eleven>_2, or, <in the afternoon after one>_3 // not $1 \mapsto (2 \vee 3)$

Whether or not there are some discursive markers between two outcomes, to find the appropriate operator, we need to answer some questions : does the agent want to satisfy the two outcomes at the same time ? Are the preferences on the outcomes dependent or independent ?

We have shown in this section that it is difficult to answer the second question and there is quite some ambiguity between the operators $\&$ et \mapsto . This ambiguity can be explained by the fact that both operators model the same optimal preference. Indeed, we saw in section 4.2 that for two outcomes o_1 and o_2 linked by a conjunction of preferences ($o_1 \& o_2$), we have $o_1 \succ \overline{o_1}$ and $o_2 \succ \overline{o_2}$. For two outcomes o_1 and o_2 where o_2 is linked to o_1 by a conditional preference ($o_1 \mapsto o_2$), we have $o_1 \succ \overline{o_1}$ and $o_1 : o_2 \succ \overline{o_2}$. In both cases, the best possible world for the agent is the one where o_1 and o_2 are both satisfied at the same time.

6 Conclusion and Future Work

In this paper, we proposed a linguistic approach to preference acquisition that aims to infer preferences from dialogue moves in actual conversations that involve bargaining or negotiation. We studied how preferences are linguistically expressed in elementary discourse units on two different corpus genres: one already available, the `Verbmobil` corpus and the `Booking` corpus purposely built for this project. Annotators were trained only for `Verbmobil`. The aim is to study to what extent our annotation scheme is genre dependent.

Our preference annotation scheme requires two steps: identify the set of acceptable and non acceptable outcomes on which the agents preferences are expressed, and then identify the dependencies between these outcomes by using a set of specific non-boolean operators expressing conjunctions, disjunctions and conditionals. The inter-annotator agreement study shows good results on each corpus genre for outcome identification, outcome acceptance and outcome attachment. The results for outcome dependencies are also good but they are better for `Verbmobil`. The difficulties concern the confusion between disjunctions and conjunctions mainly because the same linguistic realizations do not always lead to the same operator. In addition, annotators often fail to decide if the preferences on the outcomes are dependent or independent.

This work shows that preference acquisition from linguistic actions is feasible for humans. The next step is to automate the process of preference extraction using NLP methods. We plan to do it using an hybrid approach combining both machine learning techniques (for outcome extraction and outcome acceptance) and rule-based approaches (for outcome attachment and outcome dependencies).

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.
- Jason Baldrige and Alex Lascarides. 2005. Annotating discourse structures for robust semantic interpretation. In *Proceedings of the 6th IWCS*.
- Avrim Blum, Jeffrey Jackson, Tuomas Sandholm, and

- Martin Zinkevich. 2004. Preference elicitation and query learning. *Journal of Machine Learning Research*, 5:649–667.
- Craig Boutilier, Ronen Brafman, Chris Geib, and David Poole. 1997. A constraint-based approach to preference elicitation and decision making. In *AAAI Spring Symposium on Qualitative Decision Theory*, pages 19–28.
- 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.
- Robin Burke. 2000. Knowledge-based recommender systems. In *Encyclopedia of Library and Information Science*, volume 69, pages 180–200. Marcel Dekker.
- Anaïs Cadilhac, Nicholas Asher, Farah Benamara, and Alex Lascarides. 2011. Commitments to preferences in dialogue. In *Proceedings of SIGDIAL*, pages 204–215. ACL.
- Li Chen and Pearl Pu. 2004. Survey of preference elicitation methods. Technical report.
- Ward Edwards and F. Hutton Barron. 1994. Smarts and smarter: Improved simple methods for multiattribute utility measurement. *Organizational Behavior and Human Decision Processes*, 60(3):306–325.
- Johannes Fürnkranz and Eyke Hüllermeier, editors. 2011. *Preference Learning*. Springer.
- Murthy Ganapathibhotla and Bing Liu. 2008. Mining opinions in comparative sentences. In *Proceedings of the 22nd International Conference on Computational Linguistics - Volume 1, COLING '08*, pages 241–248, Stroudsburg, PA, USA. Association for Computational Linguistics.
- 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.
- Katrin Schulz. 2007. *Minimal Models in Semantics and Pragmatics: Free Choice, Exhaustivity, and Conditionals*. PhD thesis, ILLC.
- Xiaoyuan Su and Taghi M. Khoshgoftaar. 2009. A survey of collaborative filtering techniques. *Advances in Artificial Intelligence*, 2009:1–20.
- Wolfgang Wahlster, editor. 2000. *Verbmobil: Foundations of Speech-to-Speech Translation*. Springer.
- Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. *Language Resources and Evaluation*, 39(2-3):165–210.