Extracting and Modelling Preferences from Dialogue

Nicholas Asher[1], Elise Bonzon[2], and Alex Lascarides[3]

[1] IRIT, CNRS, Université Paul Sabatier, 118 route de Narbonne, F-31062 Toulouse Cedex 4, France
asher@irit.fr
[2] LIPADE, Université Paris Descartes, 45 rue des Saints Pères, 75006 Paris, France
elise.bonzon@parisdescartes.fr
[3] School of Informatics, University of Edinburgh, 10 Crichton Street, Edinburgh, EH8 9AB, Scotland, UK
alex@inf.ed.ac.uk

Abstract. Dialogue moves influence and are influenced by the agents' preferences. We propose a method for modelling this interaction. We motivate and describe a recursive method for calculating the preferences that are expressed, sometimes indirectly, through the speech acts performed. These yield partial *CP-nets*, which provide a compact and efficient method for computing how preferences influence each other. Our study of 100 dialogues in the Verbmobil corpus can be seen as a partial vindication of using CP-nets to represent preferences.

1 Introduction

It is well accepted that dialogues are structured by various moves that the participants make—e.g., answering questions, asking follow-up questions, elaborating and defending prior claims, and so on. Such moves often affect the way interlocutors view a speaker's preferences and consequently influence how they respond. Dialogue (1) from the Verbmobil corpus [13] illustrates this.

(1) π_1 *A*: Shall we meet sometime in the next week?
π_2 *A*: What days are good for you?
π_3 *B*: Well, 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*: Well next week I am out of town Tuesday, Wednesday and Thursday.
π_7 *A*: So perhaps Monday?

Intuitively, *A*'s question π_1 reveals his preference for meeting next week but it does so indirectly: the preference is not asserted and accordingly responding with *I do too* (meaning "I want to meet next week too") would be highly anomalous. Nevertheless, *B*'s response π_3 to π_5 to *A*'s elaborating question π_2 reveals that he has adopted *A*'s preference. This follows his answer π_2 which specifies a non-empty extension for *what days*. Semantically, inferring π_3 to π_5 answers *A*'s question and inferring that the temporal expressions refer to next week are logically dependent.
Inferences about *B*'s preferences evolve as he gives his extended answer: from π_3 alone one would infer a preference for meeting any day next week other than Friday and its explanation π_4 would maintain this. But the continuation π_5 compels *A* to revise his inferences about *B*'s preference for meeting on Thursday. These inferences about preferences arise from both the content of *B*'s utterances and the semantic relations that connect them together. *A*'s response π_6 reveals he disprefers Tuesday, Wednesday and Thursday, thereby refining the preferences that he revealed last time he spoke. *A*'s follow-up proposal π_7 then reinforces the inference from π_6 that among Monday, Tuesday and Wednesday—the days that *B* prefers—*A* prefers Monday. This may not match

his preferred day when the dialogue started: perhaps that was Friday. Further dialogue may compel agents to revise their preferences as they learn about the domain and each other.

The dialogue moves exhibited in (1) are typical of the Verbmobil corpus, and we suspect typical also of task-oriented dialogues generally. [3] annotated 100 randomly chosen dialogues from the Verbmobil corpus with their discourse structure according to Segmented Discourse Representation Theory (SDRT, [2, 1])—these structures represent the types of (relational) speech acts that the agents perform. According to this labelled corpus, 40% of the discourse units are either questions or assertions that help to elaborate a plan to achieve the preferences revealed by a prior part of the dialogue—these are marked respectively with the discourse relations *Q-Elab* and *Plan-Elab* in SDRT, and the interpretations of utterances π_2, π_6 and π_7 and the segment π_3–π_5 in dialogue (1) invoke these relations (see Section 2)). Moreover, 10% of the moves revise or correct preferences from the context (like π_5 in (1)); and 15% of them explain prior content or prior moves (like π_4 in (1)). The remaining 35% are not pertinent to our modeling of preferences.

Inferring an agents' preferences from the speeh acts they perform is an important task because preferences are crucial for planning appropriate conversational moves, ensuring that responses in dialogue remain relevant and natural. We will model the interaction between dialogue content in dialogues of the Verbmobil corpus and preferences using (partial) CP-nets. These allow us to exploit dependencies between dialogue moves and mental states in a compact and intuitive way. But we start by motivating and describing the semantic representation of dialogue from which CP-nets will be constructed.

2 The Logical Form of Dialogue

Agents express *commitments* to beliefs and preferences through the speech acts they perform [7]. It is these commitments that concern us here, but in what follows we shall treat a commitment to a preference (or a belief) as an actual preference (or belief).

Our starting point is the aforementioned theory of discourse interpretation SDRT [1]. Like many theories [8, 10], it structures discourse into units that are linked together with *rhetorical relations* such as *Explanation*, *Question Answer Pair* (*QAP*), *Q-Elab*, *Plan-Elab*, and so on. Logical forms in SDRT consist of *Segmented Discourse Representation Structures* (SDRSs). As shown in Def. 1, an SDRS is a set of labels each representing a unit of discourse, and a mapping from each label to an SDRS-formula representing its content—these formulas are based on those for representing clauses or elementary discourse units (EDUs) plus rhetorical relation symbols between labels:

Def. 1 *An* SDRS *is a pair* $\langle \Pi, \mathcal{F} \rangle$,[4] *where* Π *is a set of labels; and* $\mathcal{F} : \Pi \longrightarrow$ SDRS-*formulas, where:*

- *If* ϕ *is an* EDU-*formula, then* ϕ *is an* SDRS-*formula.*
- *If* $\pi_1, \ldots, \pi_n$ *are labels and R is an n-ary rhetorical relation, then* $R(\pi_1, \ldots, \pi_n)$ *is an* SDRS-*formula.*
- *If* ϕ, ϕ' *are* SDRS-*formulas, then so are* $(\phi \wedge \phi')$, $\neg\phi$.

[9] represent a dialogue turn (where turn boundaries occur whenever the speaker changes) as a set of SDRSs—one for each agent representing all his current commitments, from the beginning of the dialogue to the end of that turn. The representation of the dialogue overall—a Dialogue SDRS or DSDRS—is that of each of its turns. Each agent constructs the SDRSs for all other agents as well as his own. For instance, (1) is assigned the DSDRS in Table 1, with the content of the EDUs omitted

[4] We omit the distinguished label Last from [1] as it plays no role here.

for reasons of space.[5] We adopt a convention of indexing the root label of the n^{th} turn, spoken by agent d, as nd; and $\pi : \phi$ means $\mathcal{F}(\pi) = \phi$.

Turn	A's SDRS	B's SDRS
1	$\pi_{1A} : \textit{Q-Elab}(\pi_1, \pi_2)$	$\emptyset$
2	$\pi_{1A} : \textit{Q-Elab}(\pi_1, \pi_2)$	$\pi_{2B} : \textit{Q-Elab}(\pi_1, \pi_2) \wedge \textit{QAP}(\pi_2, \pi) \wedge \textit{Plan-Elab}(\pi_2, \pi)$ $\pi : \textit{Plan-Correction}(\pi', \pi_5)$ $\pi' : \textit{Explanation}(\pi_3, \pi_4)$
3	$\pi_{3A} : \textit{Q-Elab}(\pi_1, \pi_2) \wedge \textit{QAP}(\pi_2, \pi) \wedge$ $\textit{Plan-Elab}(\pi_2, \pi) \wedge \textit{Plan-Elab}(\pi_1, \pi_6) \wedge$ $\textit{Plan-Elab}(\pi_1, \pi_7) \wedge \textit{Plan-Elab}(\pi_6, \pi_7)$	$\pi_{2B} : \textit{Q-Elab}(\pi_1, \pi_2) \wedge \textit{QAP}(\pi_2, \pi) \wedge \textit{Plan-Elab}(\pi_2, \pi)$ $\pi : \textit{Plan-Correction}(\pi', \pi_5)$ $\pi' : \textit{Explanation}(\pi_3, \pi_4)$

Table 1. The DSDRS for Dialogue (1).

A's SDRS for turn 1 in Table 1 commits him to 'caring' about the answer to the two questions π_1 and π_2 (because *Q-Elab* is veridical). We take π_1 to commit A to the implicature that he prefers to meet next week. And $\textit{Q-Elab}(\pi_1, \pi_2)$ entails that any answer to π_2 must elaborate a plan to achieve the preference revealed by π_1; this makes π_2 paraphrasable as "What days next week are good for you?", which doesn't add new preferences. B's contribution in the second turn attaches to π_2 with *QAP*; also *Plan-Elab* because of its non-empty extension for *what days*. [9] argue that this means that B is also committed to the illocutionary contribution of π_2, as shown in Table 1 by the addition of $\textit{Q-Elab}(\pi_1, \pi_2)$ to B's SDRS. This addition commits B also to the preference of meeting next week, with his answer making the preference more precise: π_3 and π_4 reveal that B prefers any day except Friday; but with π_5 he retracts the preference for Thursday. A's third turn exploits B's answer to identify a time to meet: his *Plan-Elab* move π_6 reveals he disprefers Tuesday through Friday; and the suggestion π_7 is a solution to the constraints imposed by his preferences, which have evolved through the dialogue.

3 CP-nets

We saw earlier that dialogue reveals information about preferences. These preferences influence subsequent utterances—people plan strategically so as to achieve outcomes that are most preferred. So in addition to a method for computing preferences from dialogue, we also need a method for computing which of all possible outcomes is the most preferred. We will use CP-nets [4, 5] for this. A CP-net offers a compact representation of preferences. This graphical model exploits conditional preferential independence so as to structure the decision maker's preferences under a *ceteris paribus* assumption. Representing dependencies among preferences while also exploiting their independence when appropriate is a major motivation for using CP-nets in our framework. As we shall demonstrate in Section 5, CP-nets have a major advantage for us in that it is relatively straightforward to build a CP-net *compositionally* from a DSDRS, exploiting recursion over SDRSs.
Although CP-nets generally consider variables with a finite range of values, for simplicity we consider here only propositional variables with binary values (think of each variable as the description of an action that an agent can choose to perform, or not). Moreover, we also introduce indifference

[5] We also ignore here how to construct this DSDRS from linguistic form and context; see [9] for details.

relations in these CP-nets, that is the possibility to be indifferent between both values of a variable. More formally, let V be a finite set of propositional variables and L_V the language built from V via Boolean connectives and the constants $\top$ (*true*) and $\bot$ (*false*). Formulas of L_V are denoted by ϕ, ψ, etc. 2^V is the set of interpretations for V, and as usual for $M \in 2^V$ and $x \in V$, M gives the value *true* to x if $x \in M$ and *false* otherwise. Let $X \subseteq V$. 2^X is the set of X-*interpretations*. X-interpretations are denoted by listing all variables of X, with a $\bar{}$ symbol when the variable is set to false: e.g., where $X = \{a, b, d\}$, the X-interpretation $M = \{a, d\}$ is denoted $a\bar{b}d$.
A *preference relation* $\succeq$ is a reflexive and transitive binary relation (not necessarily complete) on 2^V. Where $M, M' \in 2^V$, as usual, strict preference $M \succ M'$ holds iff $M \succeq M'$ and not $M' \succeq M$.
As we stated earlier, CP-nets exploit conditional preferential independence to compute a preferential ranking over outcomes:

Def. 2 *Let V be a set of propositional variables and $\{X,Y,Z\}$ a partition of V. X is* conditionally preferentially independent *of Y given Z if and only if $\forall z \in 2^Z$, $\forall x_1, x_2 \in 2^X$ and $\forall y_1, y_2 \in 2^Y$ we have: $x_1 y_1 z \succeq x_2 y_1 z$ iff $x_1 y_2 z \succeq x_2 y_2 z$.*

For each variable X, the agent specifies a set of *parent variables* $Pa(X)$ that can affect his preferences over the values of X. Formally, X is conditionally preferentially independent of $V \setminus (\{X\} \cup Pa(X))$. This is then used to create the CP-net:

Def. 3 *Let V be a set of propositional variables. $\mathcal{N} = \langle \mathcal{G}, \mathcal{T} \rangle$ is a* CP-net *on V, where $\mathcal{G}$ is a directed graph over V, and $\mathcal{T}$ is a set of conditional preference tables with indifference $CPT(X_j)$ for each $X_j \in V$. $CPT(X_j)$ specifies for each instantiation $p \in 2^{Pa(X_j)}$ either $x_j \succ_p \overline{x}_j$, $\overline{x}_j \succ_p x_j$ or $x_j \sim_p \overline{x}_j$.*

Exploiting the CP-net formalism and semantics enables us to “flip” the value of a variable X within an outcome to obtain a different outcome, which the agent may prefer, disprefer or be indifferent to. An outcome o is better than another outcome o' iff there is a chain of flips from o' to o which yield either preferred or indifferent outcomes, and there is at least one *improving flip*. This definition induces a partial order over the outcomes.
Despite their many virtues, classical CP-nets won’t do for representing the preferences expressed in dialogue. Suppose an agent says “I want to go to the mall to eat something”. We can infer from this that he prefers to go to the mall given that he wants to eat, but we do not know his preferences over “go to the mall” if he does not want to eat. We thus need *partial* CP-nets. A partial CP-net, as introduced by [11], is a CP-net in which some features may not be ranked. Partiality forces us to relax the semantics:

– An *improving flip* in a partial CP-net changes the value of a variable X such that: if X is ranked, the flip is improving with respect to (wrt) the CPT of X; and if X is not ranked, it is improving wrt the CPT of all features that depend on X.
– An *indifferent flip* changes the value of a variable X such that: if X is ranked, the flip is indifferent in $CPT(X)$; otherwise wrt all CPT, the change in the value of X leaves the outcome in the same position.
– *Incomparable flips* are all those flips which are neither worsening, nor improving, nor indifferent.

As before, an outcome o is preferred to outcome o' ($o \succ o'$) iff there is a chain of flips from o' to o which are all improving or indifferent, with at least one improving one. An outcome o is indifferent wrt o' ($o \sim o'$) iff at least one chain of flips between them consists only of indifferent flips. o is incomparable to o' iff none of $o \succ o'$, $o' \succ o$ or $o \sim o'$ hold.

Unlike classical CP-nets, partial CP-nets with indifference can have more than one optimal outcome even if their dependency graph is acyclic. However, we can still easily determine a best outcome, using the *forward sweep* procedure [4] for outcome optimization (this procedure consists in instantiating variables following an order compatible with the graph, choosing for each variable (one of) its preferred value given the value of the parents).
Partial CP-nets are expressive enough for the examples we have studied in the Verbmobil corpus. Section 5 will show how discourse structure typically leads to a dependence among preferences that is similar to the one exploited in CP-nets.

4 From EDUs to Preferences

Speech acts are relations between sets of commitments, just as factual statements in dynamic semantics are relations between information states. While some speech acts, like greetings, don't affect preference commitments, many speech acts do affect them, as we have seen. We must therefore extract (commitments to) preferences from speech acts. We will compute preferences in two stages: we extract them from EDUs; and modify them recursively via the discourse structure (see Section 5).
EDUs include what we call *atomic* preference statements (e.g., *I want X* or *We need X*). They can be complex, expressing boolean combinations of preferences (e.g. *I want X and Y*); they can also express preferences in an indirect way (e.g., interrogatives like *Shouldn't we go home now?* or expressions of sentiment or politeness). We regiment such complexities via a function P that recursively exploits the logical structure of an EDU's logical form to produce a *boolean preference representation* (BPR), expressed as a propositional formula. For the purposes of this paper, we define the BPR output of P manually, although in principle it is possible to learn this mapping from labelled corpus data. This BPR will then affect preferences expressed as partial CP-nets (see Section 5).
SDRT's description logic (*glue logic* or GL) is designed to express statements about the logical structure of SDRS-formulae, and so we use it here to define the function P. Formulae in GL partially describe DSDRSs in general, and the formulae associated with EDUs in particular. For instance, $\pi : \text{Not}(\pi_1)$ means that the label π in the DSDRS being described is associated with a formula $\neg\phi_{\pi_1}$, where $\neg$ is the constructor from the SDRS language that's denoted by Not, and ϕ_{π_1} is the SDRS-formula associated with π_1. We define P recursively over these GL-formulae.
We treat disjunction non-exclusively: i.e., *I want X or Y* means I prefer one of the literals or both. If the preference is exclusive, we rely on model constraints to rule out states where X and Y are satisfied. Conjunctions are ambiguous with respect to preferences, but in certain cases we can resolve the ambiguity. *I want X and Y* can mean that my most preferred state is one where both X and Y are satisfied, but I would still prefer to satisfy one of them to neither being satisfied. This disambiguation for *and* will be represented with the GL predicate &. On the other hand, this EDU could mean that I prefer the "fusion" of X and Y while not preferring either X or Y separately; we mark this in GL with $\wedge$. A final case has to do with 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. This yields the following axioms in GL for mapping EDUs to a BPR:

1. $P(\pi) = X_\pi$ for atomic π
2. $\pi : \text{Not}(\pi_1) \rightarrow P(\pi) = \neg P(\pi_1)$
3. $\pi : \text{Or}(\pi_1, \pi_2) \rightarrow P(\pi) = P(\pi_1) \vee P(\pi_2)$

4. $\pi : \&(\pi_1, \pi_2) \rightarrow P(\pi) = P(\pi_1)\&P(\pi_2)$
5. $\pi : \wedge(\pi_1, \pi_2) \rightarrow P(\pi) = P(\pi_1) \wedge F(\pi_2)$
6. $\pi : ?(\pi_1) \rightarrow P(\pi) = P(\pi_1)$
7. $\pi : ?(\neg\pi_1) \rightarrow P(\pi) = P(\pi_1)$

5 From Discourse Structure to Preferences

We now define how to update CP-nets representing an agent's preferences with the BPRs of EDUs and by discourse structure. More formally, we define a function *Commit* from a label π or discourse relation $R(\pi_1, \pi_2)$ and a contextually given CP-net $\mathcal{N}$ to an updated CP-net. We focus here on the relations that are prevalent in the Verbmobil corpus (see Section 1).
Below, X denotes a propositional variable and ϕ a propositional formula from BPR. $Var(\phi)$ are the variables in ϕ, and $\succ_X$ the preference relation associated with $CPT(X)$. $Sat(\phi)$ is a conjunction of literals from $Var(\phi)$ that satisfy ϕ, while $non\text{-}Sat(\phi)$ is a conjunction of literals from $Var(\phi)$ that do not satisfy ϕ. $Sat(\phi) - X$ is the formula that results from removing the conjunct with X from $Sat(\phi)$.

1. Where $P(\pi) = X$ (e.g., *I want X*), $Commit(\pi, \mathcal{N})$ updates $\mathcal{N}$ by adding $X \succ \overline{X}$.
2. Where $P(\pi) = \phi \wedge \psi$ (the agent prefers both ϕ and ψ, but is indifferent if he can't have both), $Commit(\pi, \mathcal{N})$ updates $\mathcal{N}$ as follows:
   - For each $X \in Var(\phi)$, add $Var(\psi)$ to $Pa(X)$ and modify $CPT(X)$ as follows:
     a. If $Sat_i(\psi), Sat_j(\phi) \vdash X$ (resp. $\overline{X}$), then $Sat_i(\psi), Sat_j(\phi) - X : X \succ \overline{X}$ (resp. $\overline{X} \succ X$), for all satisfiers i and j.
     b. If $Sat_i(\psi), Sat_j(\phi) \not\vdash X$ and $\not\vdash \overline{X}$, then $Sat_i(\psi), Sat_j(\phi) - X : X \sim \overline{X}$, for all satisfiers i and j
     c. $non\text{-}Sat_i(\psi), Sat_j(\phi) - X : X \sim \overline{X}$ and $Sat_i(\psi), non\text{-}Sat_j(\phi) - X : X \sim \overline{X}$ for all satisfiers i and j
   - Similarly for each $Y \in Var(\psi)$.

   Where ϕ and ψ are literals X and Y, this rule yields the following: $X : Y \succ \overline{Y}$, $\overline{X} : Y \sim \overline{Y}$. $Y : X \succ \overline{X}$, $\overline{Y} : X \sim \overline{X}$.

   And we obtain the following preference relation:
   $$\begin{tikzcd} & XY \arrow[dl] \arrow[d] \arrow[dr] & \\ \overline{X}Y \arrow[r, leftrightarrow] & X\overline{Y} \arrow[r, leftrightarrow] & \overline{X}\overline{Y} \end{tikzcd}$$
   Even though the dependencies are cyclic here, the use of indifference allows us to find the best outcome XY easily.
3. $P(\pi) = \phi\&\psi$ (the agent prefers to have both ϕ and ψ and prefers either one if he can't have both). We use a similar definition to that for $\wedge$, where if ϕ and ψ are literals X and Y we get $Y \succ \overline{Y}$ and $X \succ \overline{X}$.
   We obtain the following preference relation:
   $$\begin{tikzcd} & XY \arrow[dl] \arrow[dr] & \\ \overline{X}Y \arrow[r] & \overline{X}\overline{Y} & X\overline{Y} \arrow[l] \end{tikzcd}$$
4. $P(\pi) = \phi \vee \psi$ (the agent prefers to have at least one of ϕ and ψ satisfied). The definition is similar to that for $\wedge$, where if ϕ and ψ are X and Y, we get:
   - $Var(X) \in Pa(Var(Y))$ and $X : Y \sim \overline{Y}$, $\overline{X} : Y \succ \overline{Y}$.
   - $Var(Y) \in Pa(Var(X))$ and $Y : X \sim \overline{X}$, $\overline{Y} : X \succ \overline{X}$.

   We have the following preference relation:
   $$\begin{tikzcd} \overline{X}Y \arrow[r, leftrightarrow] \arrow[dr] & X\overline{Y} \arrow[r, leftrightarrow] \arrow[d] & XY \arrow[dl] \\ & \overline{X}\overline{Y} & \end{tikzcd}$$
   As before, the use of indifference allows us to find the best outcomes (XY, $X\overline{Y}$ and $\overline{X}Y$) easily.

Due to lack of space, we won't describe rule for $P(\pi) = \neg\phi$.
Iexplanation. *Iexplanation*(π_1, π_2), as illustrated with example (2), means that $P(\pi_1)$ (here, going to the mall) is causally dependent upon $P(\pi_2)$ (eating something).

(2) π_1 I want to go to the mall
π_2 to eat something

Being a veridical relation (and assuming that a commitment to content implies a commitment also to the preferences expressed by it), $Commit(Iexplanation(\pi_1,\pi_2),\mathcal{N})$ starts by applying $Commit(\pi_2, Commit(\pi_1,\mathcal{N}))$ to the contextually given CP-net $\mathcal{N}$. Then, the preferences arising from the illocutionary effects of *Iexplanation*, given its semantics, must ensure that CPTs are modified so that each variable in $P(\pi_1)$ depends on each variable in $P(\pi_2)$: i.e., $\forall X \in Var(P(\pi_1)), \forall Y \in Var(P(\pi_2)), Y \in Pa(X)$. So, $\forall X \in Var(P(\pi_1)), CPT(X)$ is constructed by simply adding all conjunctions $Sat(P(\pi_2))$ to the conditional part of $CPT(X)$. On the other hand, $\succ_X$ when the condition includes $non\text{-}Sat(P(\pi_2))$ is undefined (i.e., we don't know preferences on X if $P(\pi_2)$ is false).
For example, let $P(\pi_1) = X \vee Z$ and $P(\pi_2) = Y$. That is, the agent explains his preferences on $X \vee Z$ by Y: he wants either X or Z if Y is satisfied. We first apply $Commit(Y, Commit(X \vee Z, \langle\emptyset,\emptyset\rangle))$. By rules 4 and 1, we obtain:

- $X \in Pa(Z)$ and X: $Z \sim \overline{Z}$, $\overline{X}$: $Z \succ \overline{Z}$.
- $Z \in Pa(X)$ and Z: $X \sim \overline{X}$, $\overline{Z}$: $X \succ \overline{X}$.
- $Y \succ \overline{Y}$

Then, the rule for *Iexplanation* modifies $CPT(X)$ and $CPT(Z)$:

- $Y \in Pa(X)$ and $Z \wedge Y$:$X \sim \overline{X}$, $\overline{Z} \wedge Y$: $X \succ \overline{X}$.
- $Y \in Pa(Z)$ and $X \wedge Y$:$Z \sim \overline{Z}$, $\overline{X} \wedge Y$: $Z \succ \overline{Z}$.

This yields the following, partial, preference relation. As we do not have any information on the preference on X and Z if Y is false, the states in which Y is false are incomparable, as required.

$$\begin{tikzcd}[column sep=small, row sep=small]
 & XYZ \arrow[d] \arrow[ddl] \arrow[ddr] & \\
 & X\overline{Y}Z & \\
\overline{X}YZ \arrow[r, leftrightarrow] \arrow[d] & \overline{X}Y\overline{Z} \arrow[r, leftrightarrow] \arrow[d] & XY\overline{Z} \arrow[d] \\
\overline{XY}Z & \overline{XYZ} & X\overline{YZ}
\end{tikzcd}$$

The causal dependence in *Iexplanation* is very close to the logical dependence exhibited in an *Elab*:

(3) π_1 I want wine
π_2 I want white wine

That is, a preference for white wine depends on a preference for wine. This leads us to the following **Elab** rule: $Commit(Elab(\pi_1,\pi_2),\mathcal{N}) = Commit(Iexplanation(\pi_2,\pi_1),\mathcal{N})$ when π_1 and π_2 express a preference (i.e., $P(\pi_1)$ and $P(\pi_2)$ are defined); otherwise there is no modification of the given CP-net.
Plan-Elab marks those cases where the second term of the relation details a plan to achieve the preferences expressed in the first term (see Table 1). So $Commit(Plan\text{-}Elab(\pi_1,\pi_2),\mathcal{N}) = Commit(Elab(\pi_1,\pi_2),\mathcal{N})$.

We now turn to questions.
Q-Elab $Q\text{-}Elab_A(\pi_1,\pi_2)$ implies that the speaker A who utters the question π_2 takes over the preferences expressed in π_1 (in future, we may often identify the agent who's committed to the speech act as a subscript on the relation, as done here). More formally, $Q\text{-}Elab_A(\pi_1,\pi_2)$ implies that we update A's CP-net $\mathcal{N}$ by applying the rule for $Elab(\pi_1,\pi_2)$, where if π_2 expresses no preferences on their own, we simply set $P(\pi_2) = P(\pi_1)$. Note that this means that A's CP-net is updated with the preferences expressed by utterance π_1, regardless of who said π_1.
QAP Answers to questions affect preferences in complex ways. The first case concerns yes/no questions and there are two cases, depending on whether B replies *yes* or *no*:
Yes $QAP_B(\pi_1,\pi_2)$ where π_2 is *yes*. B's preferences $\mathcal{N}$ are updated by applying $Commit(Elab_B(\pi_1,\pi_2),\mathcal{N})$ (and so B's preferences include those expressed by π_1 and π_2).

No $QAP_B(\pi_1,\pi_2)$ where π_2 is *no*. If $P(\pi_1)$ and $P(\pi_2)$ are consistent, then B's preferences $\mathcal{N}$ are updated by applying $Commit_B(Elab_B(\pi_1,\pi_2),\mathcal{N})$; if they are not consistent, B's preferences are updated by applying $Commit(Plan\text{-}Correction(\pi_1,\pi_2),\mathcal{N})$ (see below).

Now consider $QAP_B(\pi_1,\pi_2)$, where π_1 is a *wh*-question. Then B's preferences over variables in π_1 and π_2 are exactly the same as the ones defined for a yes/no question where the answer is *yes*: variables in π_2 will refine preferences over variables in π_1. So, B's preferences $\mathcal{N}$ are updated by applying $Commit_B(Elab_B(\pi_1,\pi_2),\mathcal{N})$.

Alternatives The last and most complex sort of question and answer pair involves so called *alternative* questions such as *would you like fish or pizza?* Suppose agent A asks B an alternative question π_1 involving n variables. Then B's answer $QAP_B(\pi_1,\pi_2)$ provides information about B's preferences. Suppose $\pi_2 : \&(X_i,\dots X_n)$. Intuitively, this response provides several answers as good as any other: for $i \leq j \leq n$, B wants to satisfy the literal X_j. Therefore, we add the following preferences for each X_j, or we change the existing preferences if appropriate: $Pa(X_j) = \emptyset$ and $X_j \succ \overline{X_j}$.

Plan-Correction may affect preferences in several ways. For example, it can correct what variables are operative. That is, given $Plan\text{-}Correction(\pi_1,\pi_2)$, some variables in $P(\pi_1)$ are replaced by variables in $P(\pi_2)$. We have a set of rules of the form $X \leftarrow \{Y_1,\dots,Y_m\}$, which means that the variable $X \in Var(P(\pi_1))$ is replaced by the set of variables $\{Y_1,\dots,Y_m\} \subseteq Var(P(\pi_2))$. We assume that X cannot depend on $\{Y_1,\dots,Y_m\}$ before the *Plan-Correction* is performed. Then replacement proceeds as follows:

6. If $Pa(X) = \emptyset$, we add $Y_k \succ \overline{Y}_k$ for all $k \in \{1,\dots,m\}$ and remove $X \succ \overline{X}$ (or $\overline{X} \succ X$). Otherwise, we replace every preference statement in $CPT(X)$ with an equivalent statement using Y_k (to create $CPT(Y_k)$), for all $k \in \{1,\dots m\}$.
7. For all W such that $Var(X) \in Pa(W)$, we re-define $CPT(W)$ so that every occurrence of X and $\overline{X}$ is replaced by a set of k statements where each statement replaces replaces X with $\overline{X}$ respectively with $\bigwedge_{1\leq k\leq m} Y_k$ and $\bigvee_{1\leq k\leq m} \overline{Y_k}$.

Plan-Corrections, like the one in (1), can also remove certain options from consideration in realizing a particular plan or it can put certain options into play that were previously excluded. In particular, suppose, π_1 countenances k options $X_1,\dots,X_k$ and rules out n options $Y_1,\dots,Y_n$; thus, $P(\pi_i) = \bigvee_{1\leq i\leq k} X_i \wedge \bigwedge_{1\leq r\leq n} \neg Y_r$. Suppose π_2 removes an option X_m from π_1. Then we must replace $P(\pi_1)$ with $\bigvee_{\{1\leq i\leq k\setminus\{m\}\}} X_i \wedge (\bigwedge_{1\leq r\leq n} \neg Y_r) \wedge \neg X_m$. The rule for putting an option into play that was previously excluded is similar; one removes one of the conjuncts in $P(\pi_i)$ and adds to the disjunction. It seems impossible to state the effects of *Plan-Correction* without the level of boolean preference representations afforded by the function P; we have not found a way to modify CP-nets directly.

6 Treatment of our example

Dialogue (1) illustrates how our rules work to refine preferences as conversation proceeds. While this dialogue doesn't feature all of our rules, other examples in the Verbmobil corpus verify the other rules.

π_1 A: Shall we meet sometime in the next week?
$Commit_A(\pi_1,\langle\emptyset,\emptyset\rangle) = P(\pi_1) = M$, where M means Meet.

Fig. 1. A's preferences

π_2 *A*: What days are good for you?
Q-Elab(π_1, π_2). *A* continues to commit to M on π_2 and no new preferences are introduced by π_2 (i.e. $P(\pi_2) = P(\pi_1)$).
π_3 *B*: Well, I have some free time on almost every day except Fridays.
π_4 *B*: Fridays are bad.
π_4 is linked to π_3 with explanation, but this has no effect on preferences. In π_3, *B* says he has some free time on Monday, Tuesday, Wednesday and Thursday (and so can meet on these days); he does not want to meet on Friday. So, we update *B*'s CP-net $\langle \emptyset, \emptyset \rangle$ with *Q-Elab*(π_1, π_2) and then *QAP*(π_2, π_3), where $P(\pi_3) = (J_1 \vee J_2 \vee J_3 \vee J_4) \wedge \neg J_5$, with J_1 being Monday, J_2 Tuesday, J_3 Wednesday, J_4 Thursday and J_5 Friday. Where $I = \{1,2,3,4,5\}$ this update yields:

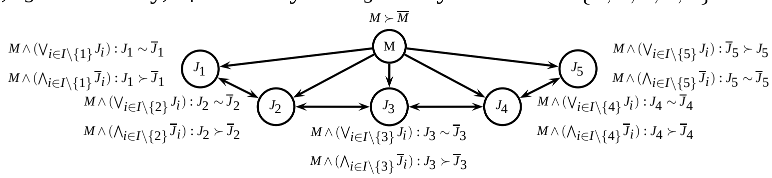

Fig. 2. *B*'s preferences

π_5 *B*: In fact, I'm busy on Thursday too.
This is a *Plan-Correction* with $P(\pi_5) = \neg J_4$, and thus $J_4 \leftarrow \neg J_4$. Thus J_4 is no longer an option. The above rule for updating a CP-net $\mathcal{N}$ with this dialogue move *Plan-Correction*(π, π_5) (where π outscopes π_3 and π_4) therefore removes the disjunct J_4 from the BPR for the first argument π, and adds the conjunct $\neg J_4$. The effect of the resulting BPR is this update to *B*'s CP-net:

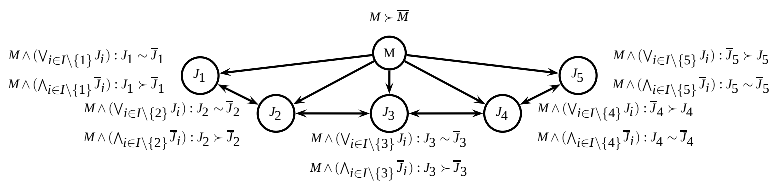

Fig. 3. *B*'s preferences

π_6 *A*: Well next week I am out of town Tuesday, Wednesday and Thursday.
The above rule for updating *A*'s prior CP-net (see Figure 1) with *Plan-Elab*(π_1, π_6), where $P(\pi_6) = \neg J_2 \wedge \neg J_3 \wedge \neg J_4$, yields the following CP-net.

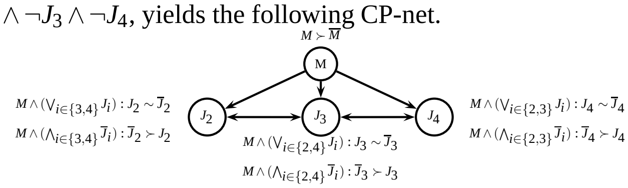

Fig. 4. *A*'s preferences

π_7 *A*: So perhaps Monday?
Commit updates the CP-net in Figure 4 with the move *Q-Elab*(π_1, π_7), where $P(\pi_7) = J_1$. Using the same rules as before this yields:

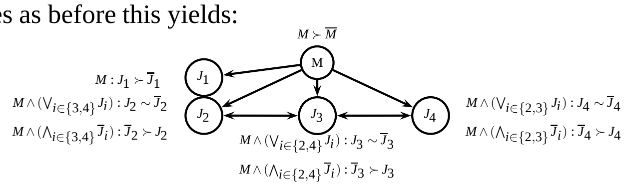

Fig. 5. *A*'s preferences

Our rules suffice to analyse the dialogues we have examined from the Verbmobil corpus. We have also analyzed examples from a tourism corpus where our rules suffice to extract preferences.

7 Conclusion

Computing preferences expressed in texts is important for many NLP applications. We have shown how to use CP-nets and models of discourse structure, together with the intermetiate level BPR, to investigate this task formally. Our rules for preference modelling are straightforward, intuitive and of low complexity. While CP-nets can loose their polynomial time complexity for computing best outcomes, if conjunctive ($\wedge$) or disjunctive ($\vee$) preferences occur, on the whole the formalism remains tractable. Once we can extract preferences, we are in a position to broaden current analyses of dialogue beyond the usual Gricean cooperative settings [6], in which agents' preferences are assumed to be aligned, and to use game-theoretic techniques to analyze strategic conversations, in which preferences are not aligned or not known to be aligned. Thus, our work here opens a way to attack the complex interaction between what agents say, what their preferences are, and what they take the preferences of other dialogue agents to be. Of course, all this depends on extracting discourse structure from text, which has proved to be a difficult task. Nevertheless [3, 12] show how one can begin to extract discourse structure automatically from texts like those found in the Verbmobil corpus. So we hope that our proposal will eventually find its way into automatic systems. In any case, our formal approach serves as a model for what such systems should aim to accomplish with respect to preference modeling.

References

1. N. Asher and A. Lascarides. *Logics of Conversation*. Cambridge University Press, 2003.
2. N. Asher. *Reference to Abstract Objects in Discourse*. Kluwer Academic Publishers, 1993.
3. J. Baldridge and A. Lascarides. Probabilistic head-driven parsing for discourse structure. In *CoNLL'05*, 2005.
4. C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. CP-nets: A Tool for Representing and Reasoning with Conditional *Ceteris Paribus* Preference Statements. *Journal of Artificial Intelligence Research*, 21:135–191, 2004.
5. C. Boutilier, R. Brafman, C. Domshlak, H. Hoos, and D. Poole. Preference-Based Constrained Optimization with CP-nets. *Computational Intelligence*, 20(2):137–157, 2004.
6. H. P. Grice. Logic and conversation. In *Sytnax and Semantics Volume 3: Speech Acts*, p. 41–58, 1975.
7. C. Hamblin. *Imperatives*. Blackwells, 1987.
8. J. R. Hobbs, M. Stickel, D. Appelt, and P. Martin. Interpretation as abduction. *Artificial Intelligence*, 63(1–2):69–142, 1993.
9. A. Lascarides and N. Asher. Grounding and correcting commitments in dialogue. In *Proceedings to SIGDIAL*, 29–36, 2008.
10. W. C. Mann and S. A. Thompson. Rhetorical structure theory: A framework for the analysis of texts. *International Pragmatics Association Papers in Pragmatics*, 1:79–105, 1987.
11. F. Rossi, B. Venable, and T. Walsh. mCP nets: representing and reasoning with preferences of multiple agents. In *AAAI'04*, p. 729–734, 2004.
12. D. Schlangen and A. Lascarides. Resolving fragments using discourse information. In *Edilog'02*, 2002.
13. W. Wahlster, editor. *Verbmobil: Foundations of Speech-to-Speech Translation*. Springer, 2000.