Automated Negotiation Under User Preference Uncertainty

Adel Magra, Peter Spreij[1], Tim Baarslag[2], and Michael Kaisers[2]

[1] Korteweg-de Vries Institute for Mathematics, University of Amsterdam
[2] Centrum Wiskunde & Informatica, Amsterdam

1 Introduction

We are concerned with automated agents representing humans in negotiations. To negotiate effectively and obtain a favorable outcome, the agent must know the preferences of the human user it is representing. These preferences are often represented by a utility function. When the agent does not know these preferences, we say that it is negotiating under uncertainty (Fig 1). To gather information about these preferences, the agent can interact with the user by asking questions or queries. The whole point of automating a negotiation is to make it more convenient for the user, we therefore do not want the agent to ask too many queries. Optimal queries were previously considered as ones with high expected value of information, which is the prospected gain in utility that a query can add to the final outcome of the negotiation (i.e. the one agreed upon) [1]. We bring forward another perspective: We consider queries as optimal based on their inherent ability to reduce uncertainty on the user's preferences.

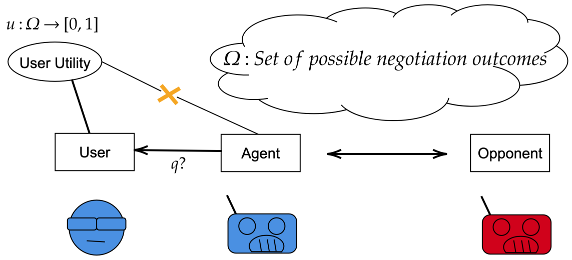

Fig. 1. Negotiating Under User Preference Uncertainty

2 A Framework for Optimal Reduction of Uncertainty

We build a general framework to formally deal with the problem of reducing uncertainty (Fig 2). We assume that the true user utility is parametrizable by $\theta^* \in \Theta$ and that queries can be answered by the user according to an answer function a. We introduce the notion of information potential of a query, which is the minimal amount of information that the agent can extract on the user's utility by asking it. It quantifies the worst possible reduction of uncertainty that is obtained when asking a query. Formally, we denote it by $I(q)$ for a query q and define it as such:

$$\begin{aligned} I(q) &:= \min_{r \in \mathcal{A}} -\log \Pi(\Theta_{(q,r)}) \\ &\textit{where } \forall r \in \mathcal{A},\ \Theta_{(q,r)} := \{\theta \in \Theta : a(\theta^*, q) = r\} \end{aligned} \tag{1}$$

Based on its current belief on the user preferences, the agent's objective therefore becomes to ask a query that maximizes the information potential. After observing the answer to a query $a(\theta^*, q)$, the agent's belief is narrowed down to the set $\Theta_{(q,a(\theta^*,q))}$, which we call the posterior set. We thus provide an objective approach of reducing uncertainty through a sequential optimization problem: the agent must find a sequence of queries maximizing the information potential.

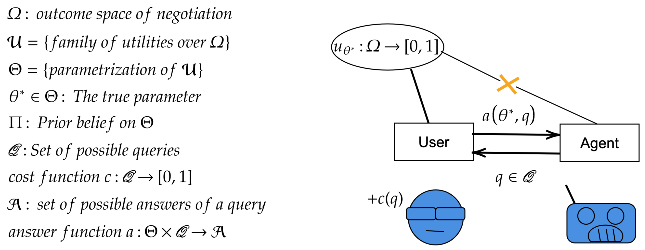
Ω : outcome space of negotiation
$\mathcal{U}$ = {family of utilities over Ω}
Θ = {parametrization of $\mathcal{U}$}
$\theta^* \in \Theta$: The true parameter
Π : Prior belief on Θ
$\mathcal{Q}$: Set of possible queries
cost function $c : \mathcal{Q} \to [0,1]$
$\mathcal{A}$: set of possible answers of a query
answer function $a : \Theta \times \mathcal{Q} \to \mathcal{A}$
$u_{\theta^*} : \Omega \to [0,1]$
$a(\theta^*, q)$
User
Agent
$q \in \mathcal{Q}$
$+c(q)$

Fig. 2. Our Framework to Deal with User Preference Uncertainty

3 Application to Multi-Issue Negotiation

A multi-issue negotiation domain is divided into different issues whose combinations form possible deals for the negotiation: $\Omega = I_1 \times \cdots \times I_n$. A popular way of representing a user's preferences on a multi-issue domain is through a linear additive utility function, which gives relative importance weights to issues and valuations for the possible values of issues. The unknown user utility is: $u_{\theta^*}(\omega) = \sum_{i=1}^{n} \theta_i^* \cdot \mathrm{val}(\omega_i)$, with θ^* being in the standard $n-1$ simplex Δ^{n-1}, where $\Delta^{n-1} = \{\theta \in \mathbb{R}^+, \sum_{i=1}^{n} \theta_i^* = 1\}$. The queries we consider are pairwise outcome comparisons: asking the user to compare two given outcomes in Ω.

We use our framework to derive an optimal querying algorithm to reduce uncertainty on the issue weights vector θ^*. We assume a uniform prior on Δ^{n-1}. Because u_{θ^*} is linear additive, pairwise outcome comparisons correspond to hyperplanes. Our Optimal Query Sequence (OQS-n) algorithm (Algorithm 1) exploits that by successively finding a query that bisects the current posterior set. Under some assumptions on the valuation functions, we show that OQS-n generates query sequences of absolutely maximal information potential of arbitrary length T (Theorem 1).

Algorithm 1: OQS-n

```
Input: $\Omega, \mathrm{val}_1, \cdots, \mathrm{val}_n, T$
$P \leftarrow V(\Delta^{n-1})$                                  // Store the $n$ vertices of $\Delta^{n-1}$
for $t \in \{1, \cdots, T\}$ do
    $(p, q) \leftarrow_{(p_i, p_j) \in P^2} d(p_i, p_j)$         // Find longest edge of $P$
    $m \leftarrow \frac{1}{2}(p + q)$                            // mid point of longest edge
    $\ell \leftarrow \mathcal{HP}(P \backslash \{p, q\}, m)$     // $\ell$ is the hyperplane bisecting $P$
    $q \leftarrow Query(\ell)$                                   // Find a query corresponding to $\ell$
    $a \leftarrow Ask(q)$
    $P \leftarrow Update(q, a)$
```

Theorem 1. *For any given $\theta^* \in \Delta^{n-1}$, and any length $T \in \mathbb{N}$, OQS-n (Algorithm 1) produces a query sequence of length T of maximal information potential.*

References

1. Baarslag, Tim and Kaisers, Michael: The value of information in automated negotiation: a decision model for eliciting user preferences. In: *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, May 2017, pp 391-400.