Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)

Preferences and Learning in Multi-Agent Negotiation

Reyhan Aydoğan
Department of Computer Engineering
Boğaziçi University
Bebek, 34342, Istanbul,Turkey
reyhan.aydogan@gmail.com

Abstract

In online, dynamic environments, the service requested by consumers may not be readily served by the producers. This requires the consumers and producers to negotiate on the content of the service. To automate this process, agents play a key role in e-commerce. As far as the agents' negotiation strategies are concerned, understanding and reasoning on their users' preferences are important to generate the right offers on behalf of their users. Besides taking other participant's needs into account is important to be able to negotiate effectively. However, preferences of participants are almost always private. The best that can happen is that participants may learn each other's preferences through interactions over time. As agents learn each other's preferences, they can provide better-targeted offers and thus enable faster negotiation. My research direction involves representing and reasoning on preferences, and learning preferences though interaction in automated negotiation.

Introduction

Automated negotiation is a key problem in agent-mediated e-commerce (Faratin, Sierra, and Jennings 2002). Under some circumstances the consumer's needs cannot be fulfilled by the producer but the producer may offer an alternative service instead of the requested one. In contrast to typical negotiation approaches based on service price (Maes, Guttman, and Moukas 1999), this thesis studies service oriented negotiation in which the participants negotiate on the content of the service.

In order to generate well-targeted offers, an agent needs to model its user's preferences, which dictate how the agent will act in negotiation on behalf of its user. Hence, representation and reasoning on preferences constitute an irreplaceable part of automated negotiation. There are a variety of ways to represent preferences. For instance, utility functions are commonly used in the literature to model the user's preferences. Alternatively, preferences can be represented by means of constraints. The user may express that the apartment should include three bedrooms and a parking area when renting an apartment. Further, a preference can be represented as an ordering of the alternative services. For example, someone may prefer a three-bedroom apartment over a one-bedroom. As far as the preference elicitation phase is concerned, eliciting user's preferences in a quantitative way would be arduous for the user especially when there exists preferential dependencies among issues.

Contrary to quantitative representations of preferences that are widely used in the literature, we advocate qualitative representations such as CP-nets (Boutilier et al. 2004) for preference representation and reasoning. A CP-net is a graphical model for representing partial preference ordering in an intuitive way mostly in the form of comparatives and conditionals. From the point of view of the user, it is relatively more natural to express her preferences in this way. However, CP-nets keep a partial preference ordering; thus some services cannot be compared under its semantics. The agent needs to compare two services in order to be able to negotiate. One challenge is how the agent negotiates with its user's partial preference information. This thesis pursues the ways of negotiation strategies work with partial preferences.

Moreover, an agent not only needs to understand its own user's preferences, but also other agents' preferences so that agreements can be reached. Since the agents do not know each other's preferences, they try to learn other participant's preferences during the negotiation. Learning other's preferences in negotiation is a challenging task. First, the participant does not know how the other agent represents its preferences. Further, the agent does not know whether preferential interdependencies among issues exist. This uncertainty leads the agent to make some assumptions about its opponent's preferences or negotiation strategy.

In open and dynamic environments, these assumptions may not work as expected. Consider that a producer agent tries to learn consumer's preferences and it assumes that the issues are independent. This assumption may be consistent with some consumer's preferences but it may fail in others. Furthermore, the number of training instances may be inadequate to learn the opponent's preferences especially if the agent meets the opponent for the first time. During this thesis, our aim is to learn a generic model which enables the negotiation faster rather than learning exact preferences.

Completed Research

We have proposed an automated negotiation framework in which a consumer and a producer negotiate on a service. The consumer agent uses constraints in the form of conjunctives and disjunctives to represent its user's preferences. For example, if the user prefers either an apartment at Etiler or an apartment having a parking area at Kadikoy, it is represented as *(Neighborhood="Etiler")* $\vee$ *(Neighborhood="Kadikoy"* $\wedge$ *Parking Area="Yes")*. In the proposed approach, the producer agent tries to learn the consumer's preferences by using the bids exchanged during the negotiation. To achieve this, we have developed an inductive learning algorithm, Revisable Candidate Elimination Algorithm (RCEA) (Aydoğan and Yolum 2009). RCEA is a version space algorithm based on Candidate Elimination Algorithm (CEA) (Mitchell 1997) but it supports learning disjunctive concept where CEA does not. It is important to learn disjunctives because there may be more than one service descriptions, which are acceptable by the consumer. To do this, we modify CEA algorithm to be able to learn disjunctives. Moreover, our learning algorithm also incorporates the idea of revision in a way that as the negotiation proceeds, a producer can revise its idea of the customer's preferences. Further, RCEA is capable of capturing domain knowledge and reasoning on the semantic similarities in performing specialization or generalization of a hypothesis.

As a result, the producer that uses RCEA to learn the preferences generates well-targeted offers, which yields faster negotiation because the producer agent does not offer a service, which is possibly rejected by the consumer according to the learned preference information. The reason for why version spaces are used in learning preferences is that they do not require any assumptions about preference representation and it is a generic model independent from the consumer's preference representation. For instance, preferential dependencies among issues does not affect our learning process, which is also capable of handling dependencies.

Ongoing Research

We are currently studying to develop negotiation strategies for the consumer agent using CP-nets (Boutilier et al. 2004) to represent its user's preferences. Compared to quantitative representation such as utility function, elicitation of CP-nets is more natural and intuitive for the user. The user expresses a preference ordering for each issue such as "I prefer *Etiler* rather than *Kartal* for the neighborhood". These statements are interpreted under ceteris paribus semantics ("all else being equal"). For example, the previous preference statement means that an apartment at *Etiler* is better than that at *Kartal* for the user when all other issues such as number of bedrooms, price and so on are the same. Moreover, we can also express conditional preferences such as "I prefer a two-bedroom apartment if the apartment is at *Etiler*". In this example, the preference on number of bedrooms depends on neighborhood. If the apartment is at *Etiler* and all other issues such as price are the same, the user prefers a two-bedroom apartment. However, CP-nets represent a partial preference ordering; thus we cannot compare some services. There is a trade of between the simplicity of the elicitation of the preferences and the amount of the information gained.

The consumer agent should know which service is better/worse than other in order to negotiate effectively. Since CP-nets keep partial preference information, we need to develop some heuristics to obtain an estimated total ordering and accordingly develop negotiation strategies based on these heuristics. To achieve this, we induce a preference graph from a given CP-nets. The service nodes having path from each others are comparable under ceteris paribus semantics but we cannot compare services having no path between them. We are studying developing heuristics to be able to compare the services and negotiate over the service (Aydoğan and Yolum 2010). One of our heuristics for CP-nets is based on the idea of capturing the depth of a service node in the preference graph (Aydoğan, Taşdemir, and Yolum 2008). Our aim is to compare negotiation strategies with the partial preference information with those having a total preference ordering. If the agents can be developed so that they negotiate successfully with this partial preference information when compared to a total preference ordering, the usability of the negotiation system would improve drastically. Our current experimental results show that with effective use of CP-nets, agents can negotiate comparably well to other agents having a total preference ordering.

Future Work

Since generating the entire graph from a given CP-net may be costly, we plan to generate a partial graph, which is capable of generating the right offers as well as a complete graph is. To achieve this, we need to develop some techniques for pruning and generating the necessary part of graph. We may use domain knowledge in inducing the preference graph by reasoning on a given service ontology. Moreover, reasoning on CP-nets may be performed without inducing a preference graph. We may find a way to obtain suitable requests without constructing a preference graph.

References

Aydoğan, R., and Yolum, P. 2009. Ontology-based learning for negotiation. In *IEEE/WIC/ACM International Conference on Intelligent Agent Technology*, 177–184.

Aydoğan, R., and Yolum, P. 2010. Effective negotiation with partial preference information. In *9th International Joint Conference on Autonomous Agents and Multiagent Systems*.

Aydoğan, R.; Taşdemir, N.; and Yolum, P. 2008. Reasoning and negotiating with complex preferences using CP-nets. In *The Tenth International Workshop on Agent-Mediated Electronic Commerce*.

Boutilier, C.; Brafman, R. I.; Domshlak, C.; Hoos, H. H.; and Poole, D. 2004. Cp-nets: A tool for representing and reasoning with conditional ceteris paribus preference statements. *J. Artif. Intell. Res. (JAIR)* 21:135–191.

Faratin, P.; Sierra, C.; and Jennings, N. R. 2002. Using similarity criteria to make issue trade-offs in automated negotiations. *Artificial Intelligence* 142:205–237.

Maes, P.; Guttman, R. H.; and Moukas, A. G. 1999. Agents that buy and sell. *Comm. ACM* 42(3):81–91.

Mitchell, T. M. 1997. *Machine Learning*. New York: McGraw Hill.