Reducing complexity of an agent's utility space for negotiating interdependent issues

Koen Hindriks, Dmytro Tykhonov, Catholijn M. Jonker
Nijmegen Institute for Cognition and Information
Radboud University Nijmegen
{k.hindriks, d.tykhonov, c.jonker}@nici.ru.nl

1 Introduction

Negotiation is a process by which a joint decision is made by two or more parties [8]. The parties first express contradictory demands and then move towards agreement by a process of concession making. Negotiation is an important method for agents to achieve their own goals and to form cooperation agreements, see e.g. [2,3,12]. Raiffa [9,10] explains how to set up a preference profile for each negotiator that can be used during negotiation to determine the utility of exchanged bids.

The representation of an agent's preferences by mathematical functions, called *utility functions*, which map values of issues to the utility of bids (bundles of issue values), allows the development of software support for negotiations. The complexity of a utility function determines the computational complexity of the negotiation process. In case a utility function is a weighted sum of evaluation functions in a single issue variable negotiation is tractable. This model of issue variables influencing overall utility independent from other issues also seems to correspond to the way the average human tackles negotiation. Humans tend to simplify the structure of their preferences and prefer to negotiate one issue at a time [13]. Absence of issue dependencies allows for the use of efficient negotiation strategies.

However, in some domains issue dependencies influence the overall utility of a bid. In such cases it is no longer possible to negotiate one issue at a time and Klein at al in [6] show that there is no efficient method that an agent can use to negotiate multiple issues, even if the agent tries to guess the opponent's profile. They propose to use a mediator who uses a computationally expensive evolutionary algorithm that can solve non-linear optimization tasks of high dimensionality.

In this paper, we briefly discuss a new approximation approach to tackle the complexity problem of a utility space with interdependent issues and present experimental results concerning the outcome and the computational complexity of the method. For a more detailed discussion of the approach itself and an illustrative case study, the reader is referred to [4]. Based on two observations, a weighted averaging method is proposed to approximate complex utility functions involving non-linear dependencies with functions that do not involve such dependencies. The first observation is that not all bids are equally important for negotiation: some bids are not acceptable for the agent or too optimistic to be an outcome of the negotiation. In effect, it is possible to indicate an expected region of utility of the outcome. Second, the approximation method might work less ideal if the utility space is “wild” in the sense if it has a structure that is chaotic. However, we think that in reality utility functions are far from “wild”, thus enhancing in reality the experimental results presented here, since the experimental results are based on randomly chosen utility spaces.

The paper is organized as follows. The next section provides a formal definition of utility spaces containing interdependencies between issues. Section 3 describes the approximation method for eliminating such dependencies. In section 4 we present experimental results and an analysis of the approximation with respect to the original utility space in the same negotiation setup. Section 5 concludes the paper.

2 Utility of Interdependent Issues

The overall utility of a set of *independent* issues can be computed as a weighted sum of the values of each of the issues by associating an evaluation function with each issue variable (see e.g. [5, 10]). The properties of the utility function are derived from these evaluation functions which map issue values on a closed interval [0; 1]. This model, represented in equation (1), can be used for issue values that are numeric (e.g. price, time) as well as for issue values that are discrete (e.g. colors, brands).

$$u(x_1,\dots,x_n)=\sum_{i=1}^{n} w_i ev_i(x_i) \quad \textbf{(1)}$$

This setup cannot be used, however, for modeling dependencies between issues and equation (1) needs to be generalized to equation (2) (cf. also [1]). Of course, the value of an issue does not need to depend on all other issues and subsets of dependent issues will have to be considered to model individual examples.

$$u(x_1,\dots,x_n)=\sum_{i=1}^{n} w_i ev_i(x_1,\dots,x_n) \quad \textbf{(2)}$$

As an illustration of a utility space with two dependent issues consider figure (1).

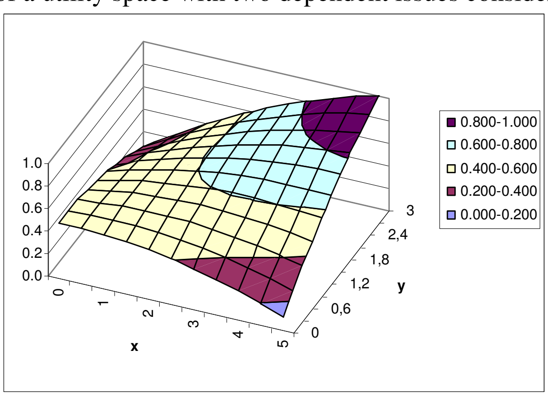

Figure 1. Utility space with issue dependencies.

The representation of a utility space with non-linear issue dependencies as in equation (2) is similar to the model proposed in [6]. The main difference is that instead of considering only binary issue values, we also allow multi-valued, discrete, as well as continuous issue ranges.

One of the main problems in dependent multi-issue negotiation is the computational complexity associated with searching for appropriate bids in the corresponding utility spaces which grows exponentially in the number of issues. One way to resolve this problem is to reveal the utility functions of both agents to a trusted mediator (cf. [6]). An alternative, interesting option is to investigate the complexity of the utility space itself and try to eliminate the dependencies between issues. In that case various alternatives for efficient negotiation become available. A new method based on *approximating* a utility space with dependent issues by a utility space without such dependencies is discussed in the following sections.

3 A Method for Approximating Complex Utility Spaces

The method to eliminate issue dependencies proposed here is based on a weighted averaging technique. It incorporates some general observations about negotiation and allows an agent to exploit knowledge about a negotiation domain, if available. In particular, knowledge about the relative importance of bids and about outcomes which reasonably can be expected is an integral part of our approximation method.

The main objective of the method is to transform a utility space $u(x_1,\ldots,x_n)$ based on dependent issues represented by equation (2) to a utility space $u'(x_1,\ldots,x_n)$ without such dependencies that can be represented by equation (1). Various techniques exist to transform complex (utility) spaces with non-linear functional dependencies between variables to spaces which are linear combinations of functions in a single variable [15].

The transformation consists of approximating each of the evaluation functions $ev_i(x_1,\ldots,x_n)$ by a function $ev'_i(x_i)$ in which the influence of the values of other issues x_j, $j\neq i$, on the associated value $ev_i(x_{1,\ldots,}x_n)$ have been eliminated. Mathematically, the idea is to "average out" in a specific way the influence of other issues on a particular issue. The averaging method first presented in [4], consists of four steps discussed next.

3.1 Estimate an Expected Outcome

In the first step, the utility of the expected outcome is estimated, using available knowledge, if any. This estimate is called the "m-point" and is used to define a region in the utility space where the actual outcome is expected to be. The m-point is used to feed information about the final goal of negotiation, i.e., the outcome, into the approximation technique used to transform the utility space. Any approach based on uniform averaging methods has the effect of discarding information uniformly and is indifferent to the fact that even before negotiation starts it can be assumed that certain regions of the utility space are more relevant to the negotiation than others.

For multi-issue negotiation in general we may assume that the expected outcome of the negotiation is located somewhere in the open utility interval (0.5; 1). Therefore, that region is of more importance than others and should be approximated with more care.

An experienced agent or one with additional knowledge about the domain can estimate an expected outcome. Using this expected outcome, the interval (0.5; 1) can be narrowed. But even if the agent has no additional information an m-point can be based on considerations of the agent's own utility space. In the latter case, we propose that the m-point can be identified with the average of the break-off point (an agent breaks off a negotiation in case any utility with a lower utility is proposed) and the maximum utility in the utility space.

3.2 Select a Weighting Function

The next step is to select a (type of) weighting function. The selection of a weighting function is based on the amount of uncertainty about the estimated m-point (expected outcome) in the previous step. As discussed above, not all points within the utility space are equally important for obtaining a good negotiation outcome. The approximation should be most accurate around the m-point, because those points are most important for getting a good negotiation outcome.

An agent may be more or less uncertain about its estimate of the utility of the negotiation outcome. To take this into account, we propose to use two different weighting functions. In case the agent has no additional information and, therefore, is quite uncertain about the most probable outcome, a relatively broad range of utility values around the expected outcome should be assigned a high weight. We propose to use a polynomial function of the second order, which is rather flat near the m-point and declines closer to the extreme utilities. The corresponding weighting function ψ then can be computed as follows by 3(a).

$$\psi(x_1, x_2 ..., x_n) = \frac{2}{m} u(x_1, x_2,..., x_n) - \frac{1}{m^2} u^2(x_1, x_2,..., x_n) \qquad \psi(x_1, x_2,..., x_n) = e^{-\frac{(u(x_1, x_2,..., x_n) - m)^2}{\sigma^2}} \qquad (3)$$

(a) Polynomial weighting function (b) Gaussian weighting function

In the case the agent is reasonably certain about the estimate, a Gaussian function as in 3(b) that is defined in terms of a maximum point *m* and spread σ can be used that assigns high weights only to bids with a utility close to the expected outcome *m*.

3.3 Compute Approximation of Utility Space

The third step is to calculate an approximation of the original utility space based on non-linear issue dependencies using the m-point and the weighting function determined in the previous steps. The result of this step is a utility space that can be defined as a weighted sum of evaluations of independent issues, i.e. of the form of equation (1). The weighted approximation technique proposed here multiplies each evaluation value with its corresponding weight and then averages the resulting space by integration. Additionally the weighting is normalized over the interval of integration.

$$ev'_i(x_i) = \int_V \psi(x_1, x_2,.., x_n) ev_i(x_1,.., x_n) dV \Big/ \int_V \psi(x_1,..., x_n) dV \qquad (4)$$

V denotes the range of integration and is a volume of *n-1* dimensionality build from the issue dimensions $\{x_1,x_2,\ldots,x_{i-1},x_{i+1},\ldots,x_n\}$. Of course, not all issues have to depend on all others and some issue variables may be dropped from the equation in that case. The approximation technique can be applied sequentially for each issue variable which involves dependencies between issues. We refer the reader to [4] for some illustrations of original and approximated spaces.

3.4 Analyze Difference δ with Original Utility Space

In the last step an analysis of the difference of the original and approximated utility space is performed. Equation (8) is used to assess the range of error for any given utility level. Based on the assessment, thresholds for breaking off the negotiation or accepting opponent's bids can be reconsidered.

$$\delta(x_1, x_2) = |u(x_1, x_2) - u'(x_1, x_2)| \quad (8)$$

If the error of the bid under consideration is too big, then another bid is proposed. This is iterated, until a proposal is found with acceptable error, which is then proposed to the negotiation partner.

4 Experimental Results

In this section, we present experimental results for utility spaces of various dimensions with non-linear issue dependencies. We want to compare the outcomes obtained by our approximation method with those obtained using the original utility space, as well as assess the efficiency of using an approximation instead of the original space. The experimental results presented in this paper are obtained from utility spaces modeled by multivariate quadratic polynomials.

Negotiation takes place between agent A and B. We assume that agent B approximates a utility space *U* with non-linear dependencies by a space *U'* to more efficiently find bids with a particular utility value. The negotiation strategy used by both agents is based on the ABMP-strategy proposed by Jonker and Treur [5] and summarized in figure 2 for agent B. Any other strategy could have been used for this purpose, since it is not the strategy that is under evaluation, but the utility space approximation method. Note that the approximated utility space *U'* is only used in the initialization and step 3 in line with its purpose of speeding up the negotiation. Agent B uses the same strategy with steps 1 and 4 exchanged and uses exhaustive search in step 3 to determine a next bid. To compare outcomes, the negotiation is performed also with agent B using exhaustive search in step 3 for each generated utility space.

Negotiation Strategy Used by Agent B
Initialization: set initial utility to maximum of U'_B.

1 **Evaluate bid $bid_A(i)$ received from opponent A:**
Accept and end negotiation if $U_B(bid_A(i)) > U_B(bid_B(i))$

2 **Compute concession and target utility:**
Concession $\gamma = \beta * (1 - \mu / U_B(bid_B(i))) * (U_B(bid_A(i)) - U_B(bid_B(i)))$
Target Utility $\tau = U_B(bid_B(i)) + \gamma$

3 **Determine a next bid:**
Find a bid $bid_B(i+1)$ such that $U'_B(bid_B(i+1)) = \tau$
and allow for maximum δ, i.e. $|U'_B(bid_B(i+1)) - U_B(bid_B(i+1))| < \delta$

4 **Send bid to opponent.**

Figure 2. Outline of Negotiation Strategy used by Agent B.

Also observe that in step 3 an additional check has been incorporated into the strategy to avoid the risk of proposing bids with (very) low utilities in the original space that have (much) higher utilities in the approximated space. The concessions made in step 3 thus are controlled by a parameter δ to ensure they are not too big. This additional check itself is computationally cheap, since it involves only a simple calculation using the original utility equations. Of course, the computational costs will increase since an agent has to search for a bid that is acceptable. The probability of finding an appropriate bid, however, is high in regions close to the m-point. Adding a check thus still results in significant reduction of the computational costs in comparison with exhaustive search.

Utility spaces have been randomly generated. Since we also want to compare outcomes by using exhaustive search, we obviously need to restrict the dimensionality of these spaces. We have performed experiments with 2, 3, 4 and 5 issues. In total, we performed 3715 experiments in which outcomes were compared with the original space: 150 for 2 issues, 300 for 3 issues, 2965 for 4 issues, and 300 for 5 issues. The results are depicted in figure 3, which shows a distribution of the deviation from the original outcome as a result of using the approximation technique.

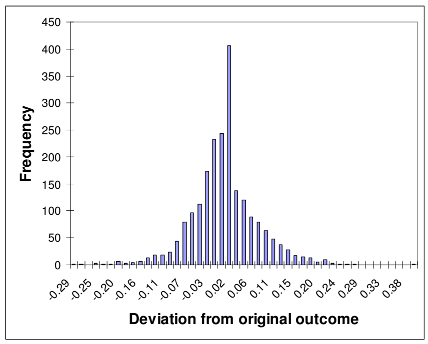

Figure 3. The distribution of deviations from original outcomes.

A linear regression model was built to analyze the relationship between negotiation outcomes in the original utility space and those obtained by approximation. In our experiments, 47% of the outcomes obtained by means of the approximated space (agent B) show an improvement over the original outcome whereas 53% do not. The graph confirms that the bids computed by using the approximated space preserve their meaning for the negotiator, although there is a small chance that the outcome is significantly worse than that obtained (at a significantly higher cost) by using the original space.

The number of computational steps required using exhaustive search versus the approximated space for 3, 4 and 5 issues is depicted in figure 4. These figures show a significant reduction in number of steps that are needed to reach an agreement when the number of issues grows. The computational gain for 3 issues is smaller than that for 4 and 5 issue negotiations, as is to be expected. Since averaging of utilities lowers the spread of utility values in the space (both minimum and maximum utilities are raised respectively lowered) it is easier to find an appropriate bid in step 3 in the strategy of figure 2 for higher dimensional spaces.

The average computational step in figure 4 for higher number of bids strongly decreases for 3 and 5 issues, indicating that for negotiations with high number of bids a significant reduction in computation time is obtained. Averages for high number of bids in these graphs, however, do not provide the same accuracy as that for 4 issues, since negotiations with very high number of bids only rarely occur and less 3 and 5 issue than 4 issue experiments have been run.

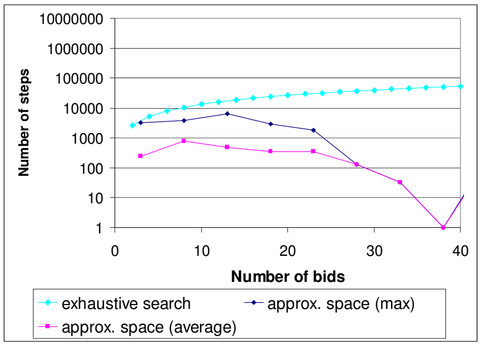

(a) 3 issues

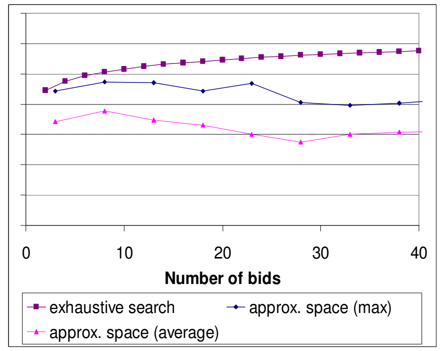

(b) 4 issues

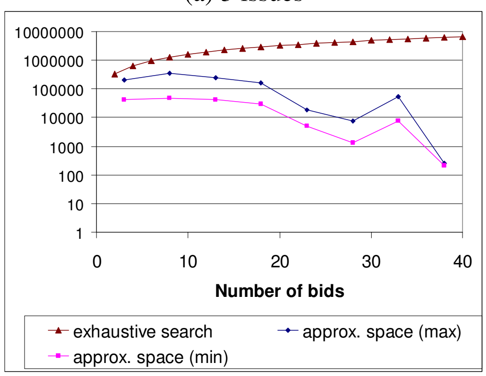

(c) 5 issues

Figure 4. Number of steps required to reach agreement.

5 Conclusion

In this paper we introduced a new approach that allows agents to deal with complex utility functions in a negotiation environment with interdependent issues. Instead of representing the negotiation task as an optimization task for interdependent issues we propose an approximation method to simplify the agent's utility space.

The main advantage of the proposed method is that it enables applicability of a wider range of computational negotiation strategies without introducing a mediator into the negotiation. Available information about the domain and the most probable negotiation outcome can be used to increase the accuracy of the method in the utility area around the expected outcome, which is most important for the negotiation. As the results show, for utility spaces involving multiple issues approximation significantly speeds up the process whereas the likelihood that the resulting outcomes remain acceptable is high.

In future research, we want to identify in more detail which classes of utility functions (instead of random spaces) can be approximated accurately by weighted averaging. Another interesting direction for research would be a modeling experiment with humans, to gain a better understanding of the nature of the complexity of human preferences and the ways in which humans simplify the negotiation task.

References

[1] Bar-Yam, Y., 1997. *Dynamics of complex systems,* Addison-Wesley (Reading).

[2] Barbuceanu, and M., Lo, W.K., 2001, Multi-attribute Utility Theoretic Negotiation for Electronic Commerce, in *Agent-Mediated Electronic Commerce III*, edited by F. Dignum and U. Cortés, Lecture Notes in Computer Science, **2003**, Springer–Verlag, pp.15-30.

[3] Davies, R., and Smith, R.G., 1983, Negotiation as a metaphor for distributed problem solving, in *Artificial Intelligence*, **20**, 1, pp. 63 – 109.

[4] Hindriks, K.V., Jonker, C.M. and Tykhonov, D., 2006, Eliminating Interdependencies between Issues for Multi-issue Negotiation, submitted.

[5] Jonker, C.M., and Treur, J., 2001, An Agent Architecture for Multi-Attribute Negotiation, in *Proceedings of the 17th International Joint Conference on AI, IJCAI'01*, edited by B. Nebel, pp. 1195 – 1201.

[6] Klein, M., Faratin, P., Sayama, H., and Bar-Yam, Y., 2002, Negotiating Complex Contracts, in *Autonomous Agents and Multi-Agent Systems*, AAAI Press (Bologna).

[7] Lai, G., Li, C., Sycara, K., and Giampapa, J., 2004, *Literature Review on Multi-attribute Negotiations,* Technical Report CMU-RI-TR-04-66, Carnegie Mellon University, Robotics Institute.

[8] Pruitt, D.G., 1981, *Negotiation Behavior*, Academic Press.

[9] Raiffa, H., 1982, *The art and science of negotiation, how to resolve conflicts and get the best out of bargaining,* Belknap Press of Harvard University Press (Cambridge).

[10] Raiffa, H., 1996, *Lectures on Negotiation Analysis*, PON Books, Program on Negotiation at Harvard Law School, 513 Pound Hall, Harvard Law School (Cambridge).

[11] Robu, V., Somefun, D.J.A., La Poutre, J.A., 2005, Complex Multi-Issue Negotiations Using Utility Graphs, in *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'05),* Utrecht, pp. 280-287.

[12] Rosenschein, J.S., and Zlotkin, G., 1994, *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*, MIT Press.

[13] Thompson, Leigh, 2000, *The Mind and Heart of the Negotiator*, Prentice-Hall.

[14] Wang, I-J., Chong, E. K. P., and Kulkarni, S. R., 1996, Weighted Averaging and Stochastic Approximation, in *Proceeding of the 35th Conference on Decision and Control,* Kobe, Japan, December 1996, pp. 1071-1076.