

Eliminating Interdependencies Between Issues for Multi-issue Negotiation

Koen Hindriks, Catholijn M. Jonker, and Dmytro Tykhonov

Nijmegen Institute for Cognition and Information, Radboud University, Montessorilaan 3,
6525 HR Nijmegen, The Netherlands
{C.Jonker, K.Hindriks, D.Tykhonov}@nici.ru.nl}

Abstract. In multi-issue negotiations, issues may be negotiated independently or not. In the latter case, the utility associated with one issue depends on the value of another. These issue dependencies give rise to more complex, non-linear utility spaces. As a consequence, the computational cost and complexity of negotiating interdependent issues is increased significantly compared to the case of independent issues. Several techniques have been proposed to deal with this increased complexity, including, for example, introducing a mediator in the negotiation setting. In this paper, we propose an alternative approach based on a weighted approximation technique to simplify the utility space. We show that given certain natural assumptions about the outcome of negotiation the application of this technique results in an outcome that closely matches with the outcome based on the original, interdependent utility structure. Moreover, using the approximated utility structure, each of the issues can be negotiated independently which ensures that the negotiation is computationally tractable. The approach is illustrated by applying and testing it in a case study.

1 Introduction

Negotiation is a process by which a joint decision is made by two or more parties [7]. 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. [2,10,11]. Raiffa [8] explains how to set up a preference profile for each negotiator that can be used during negotiation to determine the utility of exchanged bids. For more information on utility and other game theoretic notions the reader is referred to e.g. [3]. Representing agent's preferences in terms of mathematical formulae expressing relationships between values of issues and the utility of bids allows the development of software support for negotiations. The complexity of these relationships determines the computational complexity of the negotiation process. One way to avoid such computational complexity is, as proposed in e.g. [4], to build up profiles as combinations of independent and simple evaluation functions per issue. This approach corresponds to the way the average human tackles negotiation. Humans tend to simplify the structure of their preferences ([12]) and prefer to negotiate one

issue at a time, which means that issues influence the utility of a bid independently from each other. Absence of issue dependencies allows for the use of efficient negotiation strategies. Until now this approach is only applicable if the values of the different attributes in the domain are independent from each other. However, in some domains the issues are interdependent.

In some domains, however, 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 et al in [5] argue 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. The authors propose to use a mediator who uses a computationally expensive evolutionary algorithm that can solve non-linear optimization tasks of high dimensionality. Bar-Yam [1] shows that in a multi-issue negotiation with issue dependencies the utility can only be described by non-linear functions of multiple issue variables.

In this paper, we present a new approach to tackle the complexity problem of a utility space with interdependent issues that is based on the following observations. First, not all bids are equally important for negotiation: there are some bids which are not acceptable for the agent or are 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, in real life cases a profile can be modeled by utility functions that are far from "wild"; they have a structure that is far from random. This paper proposes weighted averaging as a method to approximate complex utility functions with simpler functions that is based on these observations. Furthermore, the method provides a way to check the adequacy of the approximation by a measure of the introduced error.

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. A leading case study is used throughout the paper to illustrate the method. The theme of Section 4 is the analysis of the approximation with respect to the original utility space in the same negotiation setup. Section 5 summarizes the paper with conclusions about the proposed approximation method.

2 Utility of Interdependent Issues

The overall utility of a set of independent issues can be computed as a weighted sum of the values associated with each of the separate issues. As is common (see e.g. [4,8]), an evaluation function is associated with each issue variable and the utility of a bid then is computed by the following weighted sum of the issue evaluation functions:

$$u(x_1, x_2) = w_1 ev_1(x_1) + w_2 ev_2(x_2) \quad (1)$$

In equation (1), the (weighted) contribution of each issue to the overall utility only depends on the value associated with that issue and the contribution of a single issue can be modeled independently from any other issues. Evaluation functions for independent issues thus have the same properties as the utility function associated with the bids that consist of multiple issues: it maps issue values on a closed interval [0; 1].

This setup can be used for issue values that are numeric (e.g., price, time) as well as for issue values taken from ordered, discrete sets (e.g., colors, brands).

Bid utility functions that are weighted sums of the contribution of single issue values to the overall utility cannot be used, however, for modeling dependencies between issues. The value of one issue may depend on that of another, thus influencing the utility of a bid that includes both issues. For two issues, dependencies between these issues give rise to a generalization of equation (1) to:

$$u(x_1, x_2) = w_1 ev_1(x_1, x_2) + w_2 ev_2(x_1, x_2) \quad (2)$$

It is easy to generalize (2) to more than two issues. In that case, dependencies between selected subsets of issues instead of all issues may have to be considered.

As an illustrative example of dependent issues, in this paper, we consider the negotiation of an employment contract where two important issues are at stake: the number of days that have to be worked and the number of days that childcare will be provided by an employer. In the example, the candidate employee additionally has to take into account a dependency between these two issues: working time (issue variable x_1) needs to be balanced with the time s/he needs to spend with his/her child (issue variable x_2). Assuming that the partner of the candidate is working too and can take responsibility for only part of the childcare, the candidate has promised that s/he will take care of the child for at least 2 days, either by taking care in person, or by finding professional childcare. Thus the child care issue is really important and in case the employer proposes a contract for 5 days our candidate will try to negotiate a result which includes at least 2 days of childcare. In terms of utility, bids with 5 working days and less than 2 days of childcare have a low utility (e.g. $u(5,0) \approx 0.1$, $u(5,1) \approx 0.5$). In case the employer proposes a contract for only 4 days, the candidate will need to negotiate a result including only one day of childcare and a bid of 4 working days and one day of childcare has an acceptable utility value associated with it (e.g. $u(4,0) \approx 0.25$, $u(4,1) \approx 0.55$) though the candidate would prefer to work more. With respect to bids of the employer that require the candidate to work 3 days or less, there is no problem regarding the caretaking of the child. In that case, the childcare issue has much less influence on the value of the bid (e.g. $u(3,0) \approx 0.35$, $u(3,1) \approx 0.55$). Even in this relatively simple example, the values associated with each of the issues cannot be modelled independently and overall utility cannot be calculated using equation (1). The contribution of the childcare issue to overall utility depends on the number of working days associated with the other issue and vice versa in a way that introduces non-linear dependencies between the issues. Such non-linear dependencies can only be modelled by equation (2). To make the example concrete, the candidate's preferences are modelled using the following evaluation functions:

$$ev_1(x_1, x_2) = 0.01x_1^2 + 0.03x_1x_2 + 0.028x_2^2 \quad (3)$$

$$ev_2(x_1, x_2) = -0.04x_1^2 + 0.13x_1x_2 - 0.11x_2^2 + 1 \quad (4)$$

Figure 1 shows the utility space of the candidate employee defined by the evaluation functions (3) and (4) and weights $w_1 = w_2 = 0.5$.




Fig. 1. Utility space of the candidate employee with issue dependencies

The representation of a complex, interdependent utility space by evaluation functions as in equation (2) is similar to the model proposed in [5]. In contrast with Klein et al., who discuss binary issues only, however, we allow multi-valued, discrete, as well as continuous issues. Even so, they show that the computational complexity of searching through a utility space based on issue dependencies grows exponentially and cannot be handled efficiently by either agent when the opponent's utility function is unknown. Such complex negotiations are most efficiently handled by revealing the utility functions of the negotiating agents to a mediator that is trusted by both parties (cf. [5]). Computationally simple and efficient approaches covered in [6] mostly rely on the independence of issues to determine their next bid and are not applicable.¹

3 A Method for Approximating Complex Utility Spaces

Due to the inherent computational complexity and the limited number of negotiation strategies that can be used to handle issue dependencies in negotiations, it would be beneficial to have methods that simplify the negotiation process of dependent issues without using a mediator. One particularly interesting option is to investigate the complexity of the utility space itself and try to eliminate the dependencies between issues. In case issue dependencies can be eliminated, various alternatives for efficient negotiation become available: Searching through the utility space of multi-issue bids becomes feasible and negotiation strategies for independent issues can be applied.

In this section, a method based on weighted approximation is proposed to eliminate issue dependencies. It uses an averaging technique in which some general observations about negotiation have been integrated and which can take available knowledge about a negotiation domain into account. In particular, knowledge about the relative importance of bids and about outcomes which reasonably can be expected are part of the weighted averaging method.

¹ As we discuss below, however, the approach can be adapted by using exhaustive search through the utility space, but becomes intractable and in practice works only for small utility spaces.

Although elimination of issue dependencies implies a loss of information and accuracy with regard to utility, it is shown in this paper that if the influence of one issue on the associated value of another issue is “reasonable” (i.e., the utility space is not too wild) a good approximation of the complex utility space can be obtained.

The averaging technique proposed in this paper for eliminating dependencies is valid for utility spaces that have a certain “smooth” structure. The technique averages the values of bids close to each other. Therefore, utilities should not fluctuate too much from one bid to another within the proximity range set by the technique. In real life, common negotiations, this limitation on the applicability of the method is not seen as a problem considering that it is cognitively hard to make sense of wildly fluctuating utility spaces. As an indication, we think that the techniques are applicable to utility functions that can be modeled by polynomial functions of modest power. If the nature of the utility space is not clear, the applicability of the proposed techniques has to be tested for that case. A case study illustrates that the elimination of dependencies does not result in significant changes of the negotiation outcome. Additionally, a method for analyzing and assessing the difference between the original and approximated utility space is provided. This method analyze and assess the results can always be applied to arbitrary utility spaces.

Our main objective thus is to find and present a method for transforming a utility space $u(x_1, x_2)$ based on dependent issues that can be represented by equation (2) to a utility space $u'(x_1, x_2)$ without such dependencies that can be represented by equation (1). There exist various techniques to transform complex (utility) spaces with non-linear functional dependencies between variables to spaces which are linear combinations of functions in a single variable [13]. For our purposes, we are particularly interested in the linear separability of non-linear evaluation functions of dependent issues. The main idea is to transform a utility space $u(x_1, x_2)$ into an approximation $u'(x_1, x_2)$ of that space by approximating each of the evaluation functions $ev_i(x_1, x_2)$ 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 $ev_i(x_1, x_2)$ have been eliminated. Mathematically, the idea is to “average out” in a specific way the influence of other issues on a particular issue.

The weighted averaging method takes as input a utility space based on non-linear issue dependencies (i.e. issues cannot be linearly separated²) and transforms it into a utility space that can be defined as a weighted sum of evaluation functions of single issues (i.e. issues are independent). The weighted averaging method consists of the following steps:

1. As a first step, estimate the utility of an expected outcome that is reasonable (given available knowledge). This estimate is called the “m-point” and is used to define a region of utility space where the actual outcome is expected to be.
2. 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.

² In geometry, when two sets of points in a two-dimensional graph can be completely separated by a single line, they are said to be linearly separable. In general, two groups are linearly separable in n-dimensional space if they can be separated by an $n - 1$ dimensional hyperplane.

3. 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 step. The result of this step is a utility space that can be defined as a weighted sum of evaluations of independent issues (a function of the form of equation (1)).³
4. Perform an analysis of the difference of the original and approximated utility space by means of a δ -function to assess the range of the error for any given utility level. In this final step, based on the assessment, thresholds for breaking off the negotiation or accepting opponent's bids can be reconsidered.

Finally, the results of the approximation method can be used in combination with a particular negotiation strategy. In section 4, we study the results of using an approximated utility space for the child care example in a negotiation strategy and compare the results with an approach based on the original utility space. The sections below explain each of the steps in more detail and illustrate how these steps achieve the objective of eliminating issue dependencies.

3.1 Estimate an Expected Outcome

Any approach based on using uniform arithmetical averaging methods has the effect of discarding information uniformly. Such an approach does not take the final goal of negotiation into consideration: the negotiation outcome. A uniform averaging method 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. Some general observations about the structure of utility spaces that can be associated with negotiations taken from actual practice provide additional insight that can be used to increase the effectiveness of an approximation technique.

Consider, to make clear what we mean, a worst case scenario in which two agents A and B associate completely opposite utilities with bids. In other words, what is valuable for agent A is of no value for agent B. Formally, we can express this opposition in terms of utility functions as follows:

$$u_A(x_1, x_2) = 1 - u_B(x_1, x_2) \quad (5)$$

Given these utility functions, it is easy to see that the Nash product is 0.25 with associated utility values $u_A(x_1, x_2) = u_B(x_1, x_2) = 0.5$ and the same point within the utility space is an efficient negotiation outcome when using Kalai-Smorodinsky criteria, that is, a Pareto-optimal outcome with equal utilities for both parties. Assuming such opposite interests, none of the agents would accept a bid which has a utility below 0.5.

Typically, however, negotiations do not fit such worst case scenarios and there is something to gain for both parties. Formally, this means that there exist acceptable negotiation outcomes, i.e. bids, with associated utilities that are higher than 0.5. In such cases, the utility spaces of the negotiating opponents are not completely opposite as expressed by (11). This line of reasoning makes clear that in general we may as-

³ In the more general case of more than two issues, an evaluation function may depend on more than two issues and one of those issues has to be selected to be separated from the other issues.

sume that the expected outcome of the negotiation is located somewhere in the open utility interval (0.5; 1) and this region in the utility space is generally of more importance in a negotiation.

It follows from the previous considerations that some regions within the utility space are more important for obtaining a good negotiation outcome than others and in the approximation method proposed should be approximated as good as is possible. As a first step to identify these regions, an agent can estimate an expected outcome which would identify with some probability one of the more relevant points in the utility space. We call this point the “m-point”.

An agent will be able to estimate an expected outcome with reasonable exactness only if it has some knowledge about the opponent’s profile. In that case, as we illustrate below, the m-point can be computed in two steps. But even if an agent lacks any information whatsoever about its opponent 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. In the childcare example, the break-off point equals 0.37, which is equal to the minimum utility that still satisfies the candidate employee’s childcare constraint.

A second, more informed method to determine an expected outcome can be used when the agent does have some information, e.g. based on previous experience, concerning the opponent’s profile. In the childcare example, assuming that the employer will take the child care request seriously into consideration, but will try to minimize his contribution in this regard, bids with 1-2 child care days are reasonable to expect. Additionally, it may be more or less certain that the employer prefers the employee to work as much as possible and that these issues are independent from the other. Then, as an estimated model of the opponent’s profile, the following evaluation functions can be used, which, using equal weights of .5, specified by:

$$ev_1(x_1) = x_1/5 \quad (6)$$

$$ev_2(x_2) = (3 - x_2)/3 \quad (7)$$

An estimate of the expected outcome can now be computed from the agent’s own utility space and the educated guess of the opponent’s utility space using Kalai-Smorodinsky criteria, which ensures that a Pareto-optimal outcome is selected and the expected outcome is not strongly biased in favour of either one of the parties (see figure 8). Calculating the utility in our example yields $m=0.74$. This estimate may still be quite uncertain, but we will discuss this issue more extensively below. The estimated outcome only defines one parameter of the approach.

3.2 Select Weighting Function

As discussed above, not all points within the utility space are equally important for obtaining a good negotiation outcome. To take into account the relative importance of certain regions within the utility space, we introduce a weighting function associating a weight with each point (its “importance”) in the utility space. In general, there are

two useful considerations that can be made which provide clues for constructing an appropriate weighting function.

The first consideration is that a certain range of utility values are of particular interest in the negotiation. Also, certain bids may be more “appropriate” than others in a negotiation. As an example, bids with utility values below a break-off point are less significant than other bids and do not have to be approximated as well as others. In the childcare example, provided with the relevant domain knowledge, it is moreover unreasonable for our candidate employee to propose to do no work and at the same time to request 5 childcare days.

The first consideration concerning the approximation of the utility space can be given a formal interpretation by associating the highest weight with the expected outcome (the “ m -point” identified above, located within the $(0.5;1)$ interval).

The second consideration is the fact that 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 functions depending on the level of uncertainty that the agent has about the estimate of the m -parameter. In case the agent does not have information about the opponent, nor any past experience with the particular negotiation domain and 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. As a consequence, bids in a rather wide neighborhood of the m -point are equally important for the negotiation and only extreme points (with utilities close to one or zero) do not have to be approximated very accurately. Given a relatively large uncertainty, 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 (see figure 2a). The corresponding weighting function ψ then can be computed as follows:

$$\psi(x_1, x_2) = \frac{2}{m}u(x_1, x_2) - \frac{1}{m^2}u^2(x_1, x_2) \quad (8)$$

In the case the agent is reasonably certain about the estimate, for example, when the most probable region of the negotiation outcome is well defined on the basis of domain knowledge, knowledge about the opponent or experience gained in previous negotiations, a weighting function with a stronger differentiation of utilities values can be used. In that case, a Gaussian function 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 (see figure 2b):

$$\psi(x_1, x_2) = e^{-\frac{(u(x_1, x_2) - m)^2}{\sigma^2}} \quad (9)$$

The spread parameter σ provides an indication of the agent’s certainty about expected outcome. In both cases, the m -parameter represents the expected outcome and is a point in the interval $(0.5; 1)$; ψ assigns the m -point the maximal weight of 1.0.

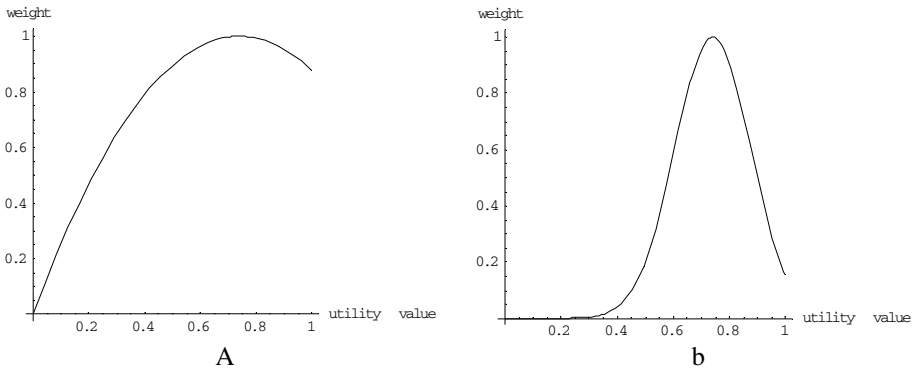


Fig. 2. Example of ψ function for $m=0.74$

In our example, an educated guess of the opponent’s profile could be made and therefore a Gaussian weighting function is selected and a value for the “spread” σ needs to be determined. To this end, we use the 3σ rule (or “Empirical rule”), which says that (most likely) 99,7% of all outcomes will be in the interval $(m - 3\sigma, m + 3\sigma)$, which gives us $\sigma=(0.37+0,74)/(2*3) = 0,19$.

3.3 Compute Approximation of Utility Space

Using the weighting function ψ a weighted approximation technique can be defined. The weighted approximation technique proposed here first multiplies each evaluation value with its corresponding weight and then averages the resulting space by integration. In the equation below, a function ω is introduced instead of ψ since the weighting must be normalized over the interval of integration. The range of integration is identical to the range of the integrated issue.⁴

$$ev'_i(x_1) = \int_{\xi_1}^{\xi_2} \omega_2(x_1, x_2) ev_i(x_1, x_2) dx_2 \tag{10}$$

Formally, the weighting function ω is defined by:

$$\omega_2(x_1, x_2) = \frac{\psi(x_1, x_2)}{\int_{\xi_1}^{\xi_2} \psi(x_1, x_2) dx_2} \tag{11}$$

So far we have been assuming a negotiation with only two issues. It is not difficult, however, to generalize the approximation technique to arbitrary numbers of issues. In case a negotiation involves N issues with interdependencies between these issues, and

⁴ If the issue has discrete values, integration simply means summation over all these values.

evaluation functions $ev_i(x_1, x_2, \dots, x_N)$ for the i^{th} issue are given, equation 14 generalizes to the equation below:

$$ev'_i(x_i) = \int_V \omega_i(x_1, x_2, \dots, x_N) ev(x_1, x_2, \dots, x_N) dV \tag{12}$$

Here V is a volume of $N-1$ dimensionality build on the dimensions $\{x_1, x_2, \dots, x_{i-1}, x_{i+1}, \dots, x_N\}$. Of course, not all issues have to depend on all others. The approximation technique can be applied sequentially for each evaluation function in the negotiation setup, which involves dependencies between issues.

As an illustration, we apply the weighted averaging technique to our employment contract negotiation. Figure 4 shows the ψ -functions for the original utility space using a polynomial function (8) for the left chart and a Gaussian function (9) for the right one. The flat section in the middle of the left chart represents a rather wide neighborhood of the m -point: this corresponds to the expected outcome and weights in its neighborhood are high. Outside this region the weighting function slowly declines to zero. For the Gaussian function (right chart) we obtain a different picture: the function has high values (close to 1) for the small band of bids with utility values close to the m -point and declines rapidly for the remainder of the utility space.

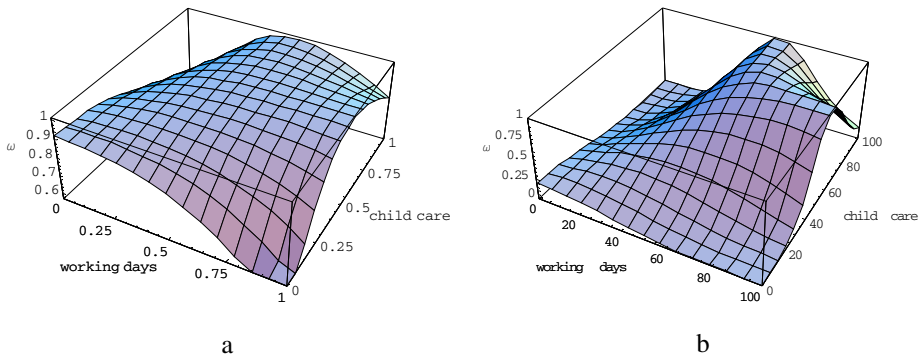


Fig. 3. Examples of ψ -functions for the employee’s utility space: (a) polynomial function with $m=0.74$; (b) Gaussian function with $m=0.74$ and $\sigma=0.19$

We apply expressions (10) and (11) to the evaluation functions of our employment contract negotiation example to derive an approximated utility space without interdependencies from the original utility space. Figure 5 shows the utility spaces obtained by approximation with a polynomial weighting function (a) and a Gaussian weighting function (b).

The utility spaces obtained by approximation with the polynomial and Gaussian weighting functions have a similar structure. However, the Gaussian weighting function due to its stronger utility discrimination power makes it more precise in the vicinity of the m -point. This is explained in more detail in the next section.

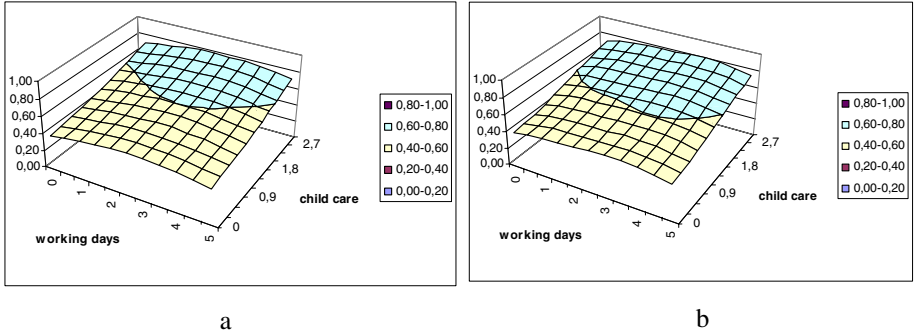


Fig. 4. The approximation by (a) the weighted averaging method using a polynomial weighting function with $m=0.74$ and (b) the Gaussian weighting function with $m=0.74$, $\sigma=0.19$

3.4 Analyze Difference δ with Original Utility Space

The technique presented approximates the original utility space and consequently, introduces an error in the utility associated with bids. To obtain a measure for the distance of the values of bids in the original utility space compared to the bids in the approximated utility space, a difference function δ can be defined as follows:

$$\delta(x_1, x_2) = |u(x_1, x_2) - u'(x_1, x_2)| \tag{13}$$

As is to be expected, the δ -values for the approximation using the Gaussian weighting function shift the utility considerably for some bids. For certain bids in the childcare example, the difference is almost 0.5. However, this only is the case for bids that are unreasonable and are not relevant for reaching a negotiation outcome. In particular, this shift in utility occurs for bids that involve more days of child care than working days. Approximations of the utility of bids that are close to the m -point are very good and close to zero.

To see the effect of the weighted averaging method near the m -point we take a section in the original utility space for the m -point ($m=0.74$ for our negotiation example). By fixing the utility to 0.74, an expression can be obtained for the value of one of the issues as a function of another one:

$$u(x_1, x_2) = 0.74 \Rightarrow x_1 = f(x_2) \tag{14}$$

The function thus obtained can be substituted into the expression of the delta function (10). This provides us with the values of δ for a fixed utility as a function of only one of the issues, and can be obtained for other utility values in a similar way.

The δ -values obtained by weighted averaging with the polynomial weighting function and the Gaussian weighting function for utility equal to 0.74 are rather small for both (see Figure 6b), but weighted averaging with a Gaussian function produces smaller approximation errors: it is almost twice as good. For bids with utilities of 0.9 the δ -values (see Figure 6c) rise in comparison with that of 0.7, however, the Gaussian weighting function still gives a better result. For bids with a utility of 0.5 (see Figure 6a) the δ -values are quite similar.




Fig. 5. Graphs depicting values of the δ functions for utility equal to (a) - 0.5, (b) - 0.7 in the original space based on a polynomial weighting function (solid line), and a Gaussian weighting function (dashed line).

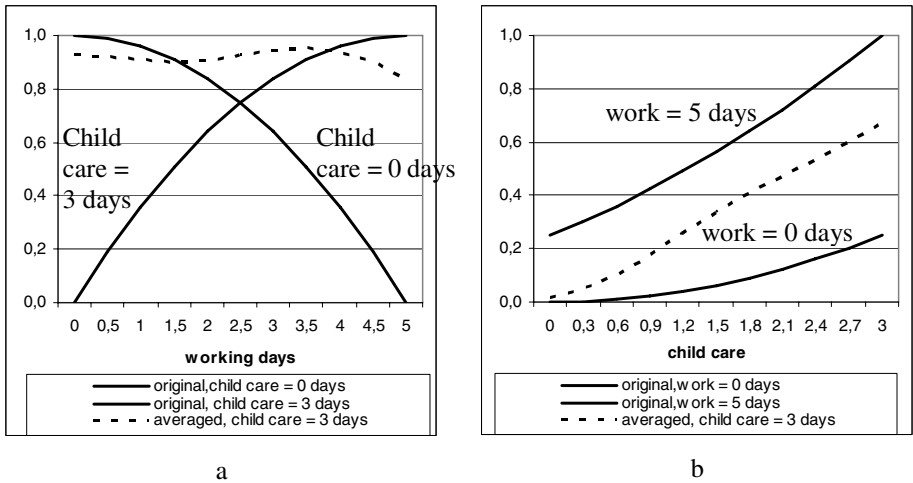


Fig. 6. Original and averaged utility values running through maximum δ -point

In figure 7, a worst case analysis is illustrated. It presents the utilities for extreme values of childcare (figure 7a) and for the number of working days (figure 7b) that run through the maximum δ -value, corresponding to the bid with 0 working days and 3 days of childcare. It shows that the evaluation function associated with 0 days of childcare (0 working days) is almost mirrored with respect to the evaluation function associated with 3 days of childcare (5 working days). In effect, this shows that our childcare example presents a serious test for our approximation method that somehow has to average these differences.

4 Case Study

In this section, a particular negotiation strategy is used to study the bids that an agent will offer during a negotiation using the original as well as the approximated utility

space. The negotiation strategy that an agent decides to use should not only fit the agent's personality profile and culture, its experience in general and the current domain and negotiation partner, but it also has to be applicable given the utility space.

A transformation of the utility space will have an effect on the negotiation process as well as on the negotiation outcome. To assess the impact of the weighted averaging approximation method, a negotiation strategy is applied to the employment contract example. Here, we use the ABMP-strategy proposed by Jonker and Treur [4].

The ABMP-strategy determines a bid in two steps: the strategy first (a) determines the target utility for the next bid, and then (b) determines a bid that has that target utility. The (b) part of the strategy is very efficient for independent utility spaces. For the purpose of comparison, however, we can use exhaustive search through the complete utility space to find a bid in the second step, provided that the space is discretized in a suitable manner (using small enough steps). In this way, the first step (a) in the ABMP-strategy followed by the second step (b) using exhaustive search can be applied to the original utility space whereas the original ABMP strategy can be applied to its approximation.

An additional check is incorporated into the strategy when the approximated utility space is used to avoid the risk of accepting bids with low utilities in the original space that have much higher utilities in the approximated space. The bids with high δ -values, that have shifted significantly due to application of the averaging method, can be filtered out in this additional step. The check applies both to a received bid as well as to the computation of a proposal for a new bid. When the agent receives a bid from its opponent, the agent has to calculate the associated original utility as well and compare it with the bid acceptance threshold. When a new bid is send to the opponent, the agent also has to check the associated utility in the original space to ensure that the bid is not worse than the current utility acceptance level. If the bid does not satisfy this condition, then agent has to find an alternative bid with the same utility value but with different issue values. These new values can be selected by systematically going through the bid space using (variants of) equation (12). This procedure guarantees that the agent will never propose or accept a bid which has a very low utility in the original utility space.

This additional check is computationally cheap, involving only a simple calculation using the original utility equations. Still, the computational costs may increase again since an agent may repeatedly need to find new bids that are 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 compared with exhaustive search.

In our experiments, the same profile of the employer was used in the original as well as in the approximated case.

Figure 8a shows the outcome space build up out of the utilities of the employer and employee per bid. Each point on the chart represents one bid. The coordinates of the bid are the utilities of the opponents (x-coordinate is the employer's utility of the bid, y-coordinate is the employee's utility of the bid). The Nash product, see e.g., [8], representing a bid with the highest utilities simultaneously for both opponents of the original utility space equals 0.53 and corresponds to a bid of 5 working days with 2.5 days of childcare, which satisfies the employee's constraints. The Kalai-Smorodinsky solution is 1.5 days of child care and 5 working days. This bid is found by locating a

bid on the Pareto-optimal frontier, which is closest to the line drawn from points with utilities of (0; 0) to points with utilities (1; 1), see e.g., [8]. This bid represents a negotiation outcome where both parties get the same utility. Using the ABMP strategy with exhaustive search for both parties, the negotiation lasts 4 rounds (4 bids from each side, the employer starts) and finishes when the employee accepts a bid of 2 days of childcare with 4.5 working days.

Figure 8b presents the result using the original ABMP strategy for both parties, where the profile of the employee has been approximated. The bids in the utility space are now concentrated around the employees original and approximated utility level of 0.7 (the m-point) with some spread towards lower utilities. The Nash product shifts to the bid of 5 working days and 1.5 days of childcare and the Kalai-Smorodinsky solution now is 4 working days and 1.5 days of childcare.

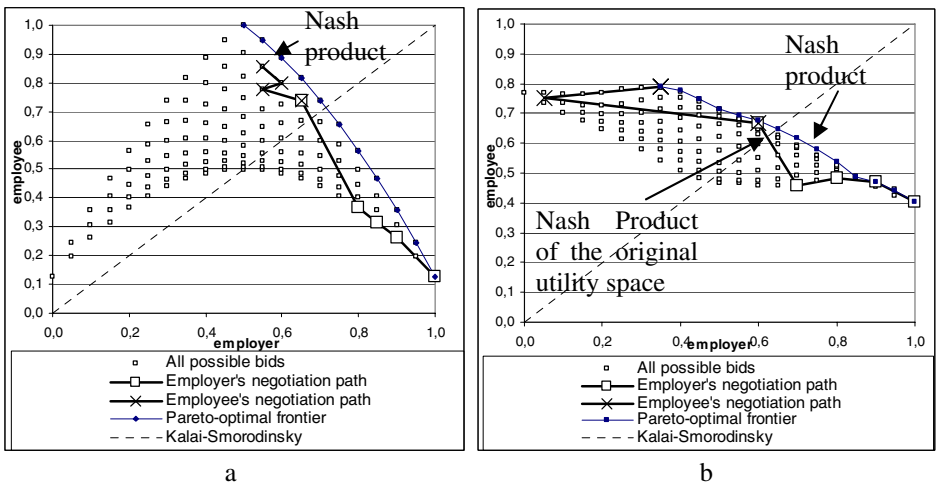


Fig. 7. Outcome space, optimality criteria, and negotiation paths (a) for the original utility space of the employee, and (b) for the approximated utility space of the employee

The original outcome space and the approximated one are significantly different. However, the difference is not critical for the negotiation itself due to the fact that most of the bids for which the difference is significant will not be used in a negotiation and we basically aim for the efficient solutions (Kalai-Somorindinsky point, and Nash Product). Also note that the bids are shifted only on the vertical axis (employee’s utility), because the employer’s profile remains the same.

The negotiation performed for the same setup but using the approximated employee’s utility space is also finished in 4 rounds as in the previous experiment and also results in a deal of 4.5 working days and 2 days of childcare.

This example shows that the approximation procedure leads to some shifts in the efficient outcomes of the negotiation with respect to Nash and Kalai-Smorodinski. However, it also confirms that these bids and those around them preserve their meaning for the negotiator. Negotiation outcomes for both utility spaces are rather close even though the negotiation paths are different.

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 using the observation that in common negotiation settings the expected negotiation outcome is approximately known and the insight that the nature of utility spaces for such common negotiation settings has enough structure to make our approach applicable. The method provides a means to analyze the impact of the approximation on a particular utility space, thereby making it possible to determine up front, whether or not the approximation is useful in any particular domain.

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. The additional check that compares the utility of exchanged bids with the utility of the original utility space during a negotiation prevents an agent from accepting low-utility bids in the original space with a high δ (error) in the approximated space. This check in itself is computationally cheap and ensures reasonable negotiation performance.

Robu et al. in [9] propose a graph-based technique to learn complex opponent's profiles. The authors propose an algorithm of exponential computational complexity for searching through a learned utility space of the opponent. The main interest in [9], however, is the scalability of a model for representing an opponent's profile which is different from the approach proposed here to simplify an agent's profile.

In future research, we want to identify in more detail which classes of utility functions can be approximated by weighted averaging sufficiently accurate. 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*. Reading, Mass., Addison-Wesley.
- [2] Barbuceanu, M., Lo, W.K., (2001), Multi-attribute Utility Theoretic Negotiation for Electronic Commerce. In: Dignum, F, and Cortés, U., (eds.), *Agent-Mediated Electronic Commerce III*, Lecture Notes in Computer Science, Vol. 2003, Springer – Verlag, pp.15-30.
- [3] Binmore, K., (1992), *Fun and Games: A Text on Game Theory*, D.C. Heath and Company, Lexington.
- [4] Jonker, C.M., Treur, J., An Agent Architecture for Multi-Attribute Negotiation. In: B. Nebel (ed.), *Proceedings of the 17th International Joint Conference on AI, IJCAI'01*, 2001, pp. 1195 – 1201.
- [5] Klein, M., Faratin, P., Sayama, H., and Bar-Yam, Y., (2002), Negotiating Complex Contracts. In: *Autonomous Agents and Multi-Agent Systems*, Bologna, Italy: AAAI Press. Paper 125 of the Center for eBusiness@MIT. <http://ebusiness.mit.edu>.

- [6] Lai, G., Li, C., Sycara, K., and Giampapa, J., (2004). *Literature Review on Multi-attribute Negotiations*. Technical Report CMU-RI-TR-04-66, Cornege Mellon University, Robotics Institute.
- [7] Pruitt, D.G., (1981), *Negotiation Behavior*, Academic Press.
- [8] Raiffa, H., Richardson, J., and Metcalfe, D., (2002). *Negotiation Analysis: The Science and Art of Collaborative Decision Making*, Cambridge, MA: Belknap Press of Harvard University Press.
- [9] Robu, V., Somefun, D.J.A., La Poutre, J.A., (2005). Complex Multi-Issue Negotiations Using Utility Graphs, *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'05)*, Utrecht, The Netherlands, July 2005, pp. 280-287.
- [10] Rosenschein, J.S., and Zlotkin, G., (1994). *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. MIT Press.
- [11] Sierra, C., Faratin, P., and Jennings, N.R., (1997), A service-oriented negotiation model between autonomous agents. In: Boman, M., and Velde, W. van de (eds), *Multi-Agent Rationality: Proceedings of the 8th European Workshop on Modelling Autonomous Agents in Multi-Agent World, MAAMAW'97*, Lecture Notes in Artificial Intelligence, Vol.1237, Springer-Verlag, pp. 17 – 35.
- [12] Thompson, Leigh, (2000), *The Mind and Heart of the Negotiator*, Prentice-Hall.
- [13] Wang, I-J., Chong, E. K. P., Kulkarni, S. R., (1996), Weighted Averaging and Stochastic Approximation, *Proceeding of the 35th Conference on Decision and Control*, Kobe, Japan, December 1996, pp. 1071-1076.