Generating Offers with Cosine Similarity in Multi-Attribute Negotiation

Aodah Diamah[1], Michael Wagner[1,2]
[1]Faculty of ESTeM, University of Canberra
[2]College of Engineering and Computer Science, Australian National University
Australia
{aodah.diamah , michael.wagner}@canberra.edu.au}

***Abstract*—Automated multi-attribute negotiation can be viewed as the search for a solution that satisfies two conflicting preferences. It is challenging not only because of the large multi-dimensional search space, but also when preferences of negotiating parties are kept private. Under such condition, negotiating agents must approximate their opponents' preferences. To this end, we employ cosine similarity to generate offers approximately close to the opponent preference. Offers received from the opponent are regarded as its partial preference. Hence, the more similar an offer to the opponent's most recent offer, the greater the chance that the offer will be accepted. To find an offer close to the opponent preference, the agent searches for the offer within its accepted utility range that has the highest cosine similarity to the opponent's last offer. Our experiments show that generating offers with cosine similarity leads to high agreement rate and mutual gain for the negotiating agents.**

Keywords—Multi-attribute negotiation, agents, cosine similarity

I. INTRODUCTION

Negotiation takes place every day, from as simple as deciding what movie to watch tonight to complex ones like recent China-US talks on their respective carbon emissions. Rarely does negotiation involve only one issue. Most of the time when trying to reach a deal, negotiating parties will look at several issues where they can work out a solution that satisfies, to some extent, all parties' needs. For example, a car dealer will offer a car sale along with extended guarantee and a regular service package. A potential buyer may feel more attracted paying a little extra money knowing he does not have to worry about the car maintenance for a few years. Because of the win-win solution potential inherent in multi-attribute negotiations, there have been studies on using agents to search for such solutions: parties get more on more important issues to compensate for getting less on less important issues [1]. Therefore negotiation is often treated as a multi-objective search for a solution satisfying two conflicting preferences [2], [3]. If both parties' preferences are known and a zone of agreement[1] exists, possible solutions can be explored and computed. Negotiation agents can sort and match issues based on their importance to guide the direction of conceding and demanding like in [3], [4].

In cases where the opponent does not reveal its preferences, finding a win-win solution is a challenging problem. Agents then need to find ways to approximate opponent preferences. One approach is by modelling the preferences such as in [5], [6] and [7]. However, modelling needs assumptions on the opponent preference structure (for example, the utility function of the opponent taking a triangle form) and its concession tactic. This approach works well when previous negotiation data is available or agents negotiate in a familiar domain or with a familiar opponent (a car dealer can, for example, recognize different opponents' needs from his experience), but may perform not so well when the assumptions made do not hold.

Another approach is using similarity measures to compare recent opponent offers with the agent's solution candidates [8], [9]. Solutions with higher similarity are expected to be more acceptable. Similarity measures are normally used in more general negotiation settings and do not require assumptions such as in [5], [6].

In this paper, we use cosine similarity to find offers that match the opponent partial preference, i.e. the opponent offers. Cosine similarity has been used before multi-attribute negotiation for constructing a utility graph of the opponent [7], but it is new for offer generation. We treat offers in negotiations as n-dimensional vectors, where n is the number of attributes/issues involved. In each round, an agent will generate offers with utility acceptable for that particular round, considering time and concession factors. Although the agent is indifferent to these offers since they have the same utility to the agent, that is not the case for the opponent because of the different preferences between the two agents. Using cosine similarity, the agent will calculate the similarity between all offer candidates and the last offer from the opponent. The chance of an offer being accepted can be assumed to be higher if that offer has a high similarity to the opponent offer.

The remainder of the paper is structured as follows: We describe how our work relates to the existing literature in section II. In section III we discuss the proposed negotiation model with cosine similarity. Section IV discusses our experimental results, and finally section V contains conclusions and an outlook to future work.

II. RELATED WORK

Faratin et al. [8] introduced fuzzy similarity criteria to generate offers in multi attribute negotiation. The principle of this mechanism is to choose one or more offers from the agent's

[1]Zone of agreement is where the acceptable values of the issues of both parties overlap.

indifference curve using fuzzy similarity criteria. The indifference curve[2] contains all offers with the same utility to the proposing agent. The rationale behind this is that the more similar an offer is to the opponent's last offer, the higher is the chance that the offer will be accepted. For this mechanism to work better compared to a random selection of offers, agents need at least partial information on how the opponent orders the importance of attributes. Ros and Sierra [4], and Jazayeri et al. [6] extended this work by learning the opponent's order of attribute importance. We use the same principle of generating offers with high similarity to that of the opponent. However fuzzy similarity requires a similarity function that is defined depending on the negotiation domain. We use a similarity mechanism that is generic regardless of the negotiation domain.

In [7], Robu and Poutré retrieved the utility structure of buyers' preferences from past concluded negotiation data. Once constructed, the seller can use this structure in future negotiations. Cosine similarity was used in this work to discover interdependencies between items, i.e. if a buyer asked for an item or not. In our work, we do not model the utility structure of negotiating agents; rather we compute the cosine similarity between offers to generate bids while negotiation is in progress. Therefore we do not require past negotiation data.

The closest work to ours is that of Lai and Sycara [9] who use distance to the opponent's best offer to maximize the similarity of one (or more) counter offer(s). Lai and Sycara presented a numerical example of their proposed model, but did not evaluate the negotiation outcome in terms of joint utility. They showed how close the agreements were to the Pareto front in attribute– instead of utility– space. Like Lai and Sycara, we search for offers with the highest similarity to the opponent's last offer, however we use a different similarity measure: cosine distance. In addition, we evaluate the negotiation outcome in terms of joint utility and closeness to the Pareto front.

III. PROPOSED NEGOTIATION MODEL

A. Negotiation settings

We define the negotiation model between two agents, the buyer A and the seller B, as follows:

Utility function: In evaluating offers, agents need to calculate how closely the offer matches the preference set by their client using utility functions. In multi-attribute negotiation, each attribute value is assigned a score by an evaluation function that falls into a normalized interval of 0 to 1. Each attribute has a different weight, which corresponds to the importance the agent places on that attribute. In most of the literature, the common assumption is that the utilities of the attributes are positive and mutually independent [17]. For n attributes, the sum of the weights must be:

$$\sum_{i=1}^{n} w^i = 1. \tag{1}$$

The total utility score of the attribute vector $x = (x^1, x^2, \dots, x^n)$ is assumed to be linearly additive:

$$V_m(x) = \sum_{i=1}^{n} w_m^i u_m^i(x) \tag{2}$$

where $u_m^i(x)$ is the utility for agent m and attribute i of x, and $V_m(x)$ is the total utility of x for agent m.

Concession strategy: This defines how much utility an agent will concede in each round. We use the time-dependent strategy from [6] to determine the target utility:

$$V_m^k = 1 + (V_m^{min} - 1)(k/k_{max})^{\alpha} \tag{3}$$

where k is the current round, k_{max} is the maximum number of rounds, V_m^{min} is the minimum acceptable utility for agent m and α is the rate at which agent m is conceding.

Negotiation Protocol: The negotiation protocol defines how agents interact during the negotiation. We use the alternating offers protocol in [10]. Suppose at round k, agent B proposes offer x_B^k. In round $k+1$, A calculates its target utility V_A^{k+1} according to (3). Agent A's response to opponent offer B is formulated as:

$$R_A(k, x_B^k) = \begin{cases} \text{withdraw if } k > k_{max} \\ \text{accept } x_B^k \text{ if } V_A(x_B^k) \geq V_A^{k+1} \\ \text{offer } x_A^{k+1} \text{ otherwise} \end{cases} \tag{4}$$

Agent A will only accept the offer from B if offer B's utility is higher than its current acceptable utility V_m^{k+1}. If that is not the case, A has to find a counter offer x_A^{k+1} unless the negotiation has reached the maximum round.

B. Offer generation with cosine similarity

In each round, once an agent has computed the acceptable utility for that round, it needs to generate an offer commensurate with that utility. With multiple attributes, there may be several offers corresponding to a given utility. The offers with the same utility are referred to as iso-utility offers. Suppose for a given round, using the concession function, the agent has computed its target utility to be θ for that round. Then, the set of iso-utility offers with utility θ for agent A can be defined as [8]

$$X_A(\theta) = \{x \mid V_A(x) = \theta\} \tag{5}$$

Although these offers are of identical utility value θ for the proposing agent, that is not necessarily the case for the opposing agent. The proposing agent needs to find an offer with higher utility to the opponent to increase the chance of the offer being accepted. Because the agent does not know the opponent's utility function, the agent can approximate the opponent's partial preference from its last offer using cosine similarity. First, the agent calculates the cosine distance of each offer in the iso-utility set to the opponent's last offer. The offer with the highest similarity is considered the best match to the opponent's preference. Thus, the chance of the offer being accepted is higher. We express our offer generating mechanism as:

$$x_A^{k+1} = arg\ max_{x \in X_A(\theta)} \left\{ \frac{x \cdot x_B^k}{|x||x_B^k|} \right\} \tag{6}$$

[2] Indifference curve is also referred to as iso-utility curve.

where x_A^{k+1} is the offer agent A is about to propose at round k+1, and x_B^k is the opponent's last offer. Here, the similarity of two offers x and x_B^k is expressed as the cosine distance of the two vectors

$$\frac{x \cdot x_B^k}{|x||x_B^k|} \tag{7}$$

Inevitably, in negotiations there will be attributes of different scales. For example, the attribute price may have values ranging from $1000 to $5000 while the attribute days of delivery may range from 1 to 15 days. In such cases, the attribute price will likely dominate the similarity measure. In this paper, we also investigate if normalization has any effect on the negotiation outcome.

IV. EXPERIMENT AND RESULT

A. Experimental settings

In this paper we employ two negotiation settings in [4], which have appeared before in [6] (negotiation setting 1) and [10] (negotiation setting 2). Because of space limitation, we ask readers to refer to [4] for explanation on the domain and utility functions of the negotiations. In negotiation setting 1, we define the weights attached to each attribute as follows:

w_A = (0.35, 0.15, 0.45, 0.05); w_B = (0.10, 0.15, 0.40, 0.35),

and for negotiation setting 2:

w_A = (0.20, 0.55, 0.25); w_B = (0.05, 0.30, 0.65).

For each negotiation setting, we ran 560 dyad negotiations, with the concession rate α systematically changed from 0.7 to 1.5, and the maximum number of negotiation rounds varied from 8 to 15.We used random offer generation as a baseline, i.e. offers were made on a random basis from the iso-utility set. Generating bids randomly is a common method when agents are indifferent to a set of offers [11], [12]. Because it is not always possible to find offers exactly matching the target utility, we modify (5) to be:

$$X_A(\theta) = \{x \mid \theta \le V_A(x) \le \theta + \delta\} \tag{8}$$

where we set δ=0.01 and the minimum acceptable utility for both agents to 0.5.

To evaluate the outcome of a negotiation, these metrics are used:

- Joint utility, expressed as the sum and the product of individual utilities. Individual utility of a deal is calculated using (2). If a negotiation does not reach an agreement, each agent receives zero utility.
- Percentage of successful negotiation, defined as:

$$\%\, agreement = \frac{\text{numbers of agreements reached}}{\text{number of negotiations}} \tag{9}$$

- Distance to largest sum and product of joint utility on the Pareto front, which contains all optimum solutions of a negotiation. A solution (x, y) is Pareto optimal when it is not possible to increase the value of x without lowering y.

B. Results and discussions

We hypothesized that negotiations will have better – if not Pareto optimal – outcomes if agents generate offers using cosine similarity compared to negotiation outcomes of randomly generated offers. As Table 1 shows, the average joint sum utility achieved by cosine distance agents is 1.1597 in negotiation setting 1, well over the average joint utility achieved by random agents (0.7990). The joint utility is even higher for agents using normalized cosine distance: 1.2630, very close to the maximum joint utility possible: 1.3. Such result is consistent for negotiation setting 2. The agents produce higher gain closer to the maximum joint utility on the Pareto front when they normalize offer attributes prior to applying cosine similarity. The average joint sum utility for negotiation setting 2 is 1.3204, which is 0.0796 from the maximum joint utility (1.4) on the Pareto front. The complete results of average joint utility and the distance to maximum joint utility for random and cosine similarity agents are shown in Table 2.

These results confirm our hypothesis that by maximising the cosine similarity of the proposer's offer to that of the opponent, negotiation agents can obtain better deals (i.e., higher joint utility). This is because agents incorporate opponent preference when selecting a counter offer as expressed in the opponent's most recent offer. When cosine similarity between the counter offer and the opponent's last offer is maximised, the counter offer will be of higher utility to the opponent, although not necessarily a Pareto optimal one.

Table 1. Results of negotiation setting 1.

Bid generation Method	Average Joint sum utility	Average Joint product utility	% agree-ment reached	Distance to max joint sum on Pareto front	Distance to max joint product
Random	0.7990	0.2296	69	0.50	0.18
Cos distance	1.1597	0.3512	100	0.14	0.06
Normalized cos distance	1.2630	0.3986	100	0.04	0.01

Table 2. Results of negotiation setting 2.

Bid generation Method	Average Joint sum utility	Average Joint product utility	% agree-ment reached	Distance to max joint sum on Pareto front	Distance to max joint product
Random	0.8851	0.2726	72	0.5149	0.2149
Cos distance	1.1670	0.3402	100	0.2330	0.1473
Normalized cos distance	1.3204	0.4353	100	0.0796	0.0522

In contrast, random selection of offers does not take into account the opponent preference. The effect of using cosine similarity to increase joint utility is more pronounced when attributes are normalised. This is because attributes of larger scale no longer dominate the similarity calculation, i.e. each attribute has the same weight when applying cosine similarity to two offers. As a result, joint payoffs achieved by normalised cosine distance are closer to the Pareto front.

Fig. 1 depicts an example of offers and counter offers using different offer generation algorithm: random, cosine distance and normalized cosine distance with all other parameters are set identical. It can be seen that offers made by random agents exhibit random pattern despite both agents are employing conceding strategy. With random algorithm, current offer proposed by an agent may be of worst utility to the opponent compared to its previous offer. Therefore, such concession made by the agent is an unfortunate step (see [13] for a complete classification of steps) , since both agents are worse off after the concession. This also explains the lower rate of agreement reached for negotiations performed by random agents (69% and 72% for negotiation setting 1 and setting 2, respectively). An ideal step would be for the proposer agent to lower its utility in order to offer a higher gain for its opponent, so that each time an agent concedes, the opponent receives a better offer, closer to its target utility. However such monotonicity cannot be guaranteed by random offers regardless the concession tactic used.

V. CONCLUSION AND FUTURE WORK

A negotiation can be regarded as the act of exchanging offers and counter offers until an offer with a mutual gain for the negotiating agents can be achieved. When preferences or utility functions of the negotiating agents are known, it is easy to search for or compute offers compatible for both, so these offers have higher chance of being accepted. However, when the agents keep their preferences private, they must approximate the preferences of their opponent.In this paper we

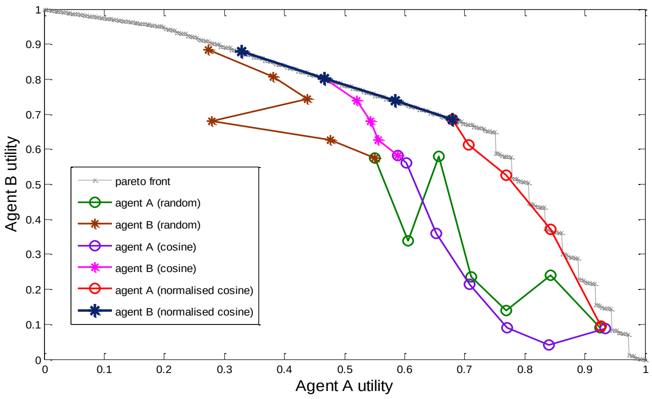

Fig. 1. *Examples of offers generated by agent A and agent B with three different methods: random, cosine similarity and normalized cosine similarity. Almost every offer made by cosine similarity agents improves opponent's gain. In addition, normalized cosine similarity offers are either on or very close to the Pareto optimal front.*

have evaluated the use of cosine similarity to generate offers in multi attribute negotiations with the preferences of the agents unknown. We have shown that offers from the opponent can be regarded as its partial preference. Offers with high cosine similarity to opponent last offer have higher utilities compared to offers made in random fashion, thus increasing both joint utility and negotiation success rate.

We are aware that both agents in our experiment concede at the same rate. In real negotiations this may not always be the case. However we intended to see if agents can benefit from using cosine similarity in negotiation. To this end – with such setting of identical conceding rates – cosine similarity can help agents achieve higher joint utility.

For our future work, we intend to evaluate if generating offers with cosine similarity can reduce the number of negotiation rounds. We will also investigate if higher joint utilities obtained in this workwill also apply to negotiations with non-linear utility functions. In addition, we will explore if combination of cosine similarity with other techniques can give better negotiation results.

REFERENCE

[1] G. Lai, C. Li, K. Sycara, and J. Giampapa, "Literature review on multi-attribute negotiations," *Robot. Inst Carnegie Mellon Univ Pittsburgh PA Tech Rep CMU-RI-TR-04-66*, 2004.

[2] A. Wierzbicki and M. Makowski, *Multi-objective optimization in negotiation support*. International Institute for Applied Systems Analysis, 1992.

[3] B. Rubenstein-Montano and R. A. Malaga, "A weighted sum genetic algorithm to support multiple-party multiple-objective negotiations," *IEEE Trans. Evol. Comput.*, vol. 6, no. 4, pp. 366–377, Aug. 2002.

[4] R. Ros and C. Sierra, "A negotiation meta strategy combining trade-off and concession moves," *Auton. Agents Multi-Agent Syst.*, vol. 12, no. 2, pp. 163–181, 2006.

[5] K. Hindriks and D. Tykhonov, "Opponent modelling in automated multi-issue negotiation using bayesian learning," 2008, pp. 331–338.

[6] H. Jazayeriy, M. Azmi-Murad, N. Sulaiman, and N. Izura Udizir, "The learning of an opponent's approximate preferences in bilateral automated negotiation," *J. Theor. Appl. Electron. Commer. Res.*, vol. 6, no. 3, pp. 65–84, 2011.

[7] V. Robu and H. L. Poutre, "Retrieving the structure of utility graphs used in multi-item negotiation through collaborative filtering of aggregate buyer preferences," *Proc 2nd Int Wk Ration. Robust Secure Negot. MAS Hakodate Jpn*, 2006.

[8] P. Faratin, C. Sierra, and N. R. Jennings, "Using similarity criteria to make issue trade-offs in automated negotiations," *Artif. Intell.*, vol. 142, no. 2, pp. 205–237, 2002.

[9] G. Lai and K. Sycara, "A Generic Framework for Automated Multi-attribute Negotiation," *Group Decis. Negot.*, vol. 18, no. 2, pp. 169–187, Mar. 2009.

[10] C.-B. Cheng, C.-C. H. Chan, and K.-C. Lin, "Intelligent agents for e-marketplace: Negotiation with issue trade-offs by fuzzy inference systems," *Decis. Support Syst.*, vol. 42, no. 2, pp. 626–638, 2006.

[11] T. Baarslag, K. Hindriks, C. Jonker, S. Kraus, and R. Lin, "The first automated negotiating agents competition (ANAC 2010)," in *New Trends in Agent-Based Complex Automated Negotiations*, Springer, 2012, pp. 113–135.

[12] S. Chen and G. Weiss, "An Efficient and Adaptive Approach to Negotiation in Complex Environments," in *ECAI*, 2012, vol. 242, pp. 228–233.

[13] K. Hindriks, C. M. Jonker, and D. Tykhonov, "Let's dans! An analytic framework of negotiation dynamics and strategies," *Web Intell. Agent Syst.*, vol. 9, no. 4, pp. 319–335, Jan. 2011.