IDENTIFYING CLASSES OF USERS TO FACILITATE NEGOTIATION

Luís Nogueira¹, Eugénio Oliveira²

¹Instituto Superior de Engenharia do Porto
Instituto Politécnico do Porto, LIACC
luis@dei.isep.ipp.pt

²Faculdade de Engenharia
Universidade do Porto, LIACC
eco@fe.up.pt
PORTUGAL

As in many industries, e-commerce technologies and strategies are stimulating a rapid growth in the services that insurance companies can offer. This rapid growth brings with it several interesting issues such as improving the quality of customer service by ensuring that individual customer's needs are reflected in the products offered. However, finding an appropriate insurance product on the Web can be a difficult and time-consuming process.

In this paper we analyse the application of machine learning techniques in modeling customers of an electronic insurance market as an approach to better match customers and insurance product offers, reducing the needed effort in negotiating an insurance product, providing a valuable add-on to both customer's and insurers' sides.

1. INTRODUCTION

Almost all insurers now have a web site providing information on the company, its products and contact details. Some of them go one step further, allowing users to purchase insurance products online as well. However, the attraction of Internet distribution to customers is limited if they have to spend a fair amount of time familiarising themselves with each web site. Customers are usually not able to easily compare quotes from several insurers.

Most of the well-known negotiation protocols deal with a single dimension (usually the price). Moreover, (Wellman, 1998) also says that the preferences' reduction to the price single figure is what characterises the market. However, this seems to be a bit simplistic and unrealistic for the insurance domain. Real negotiation in the insurance domain implies taking into consideration not only one, but multiple attributes for defining the insurance product under discussion. For instance, although the policy's premium is an important attribute, coverage's length

and policy's renewability and convertibility, for instance, are complementary issues to include in the decision about whether or not to buy a specific insurance product.

Some authors (Vulkan, 1998)(Moore, 1992) advocate the use of auctions protocols for agent mediated electronic commerce, arguing that they are widely recognised by economists as the most efficient way of resolving one-to-many negotiations. Other authors point out the limitations of auctions protocols and look for more flexible negotiation models (Guttman, 1998), and propose more cooperative multi-attribute decision analysis tools and negotiation protocols using distributed constraint satisfaction policies to support it. However, both lines of research are aware of the need of enhancing simple price-based auction mechanisms, transforming them into new protocols encompassing multi-dimensional issues to be negotiated among the market participants.

In this paper we propose an agent-mediated insurance brokering system using a flexible negotiation model that includes multi-attribute and learning capabilities. Insurers use information automatically collected during negotiation to develop a more customer-directed kind of marketing strategy. Moreover, in the core of the provided brokering facility, we are using conceptual clustering procedures for grouping users in a small set of communities. Particular attention is paid to the interpretation and effective utilisation of those generated communities. For this purpose, we use an appropriated metric to identify the representative insurance product configuration of each community in place. This information will be used by the broker agent as a way to better match particular customers needs and insurance product offers, providing a valuable add-on to both customer's and insurers' sides and, therefore facilitating the needed interaction between all the players.

2. AN ELECTRONIC INSURANCE MARKET

It is our belief that intelligent agents are well suited to deal with the insurance brokering problem in a distributed manner. Our approach for an agent-based electronic insurance market includes an agent representing each of the insurers, an agent representing the customer and a broker agent for intelligent brokering services (Nogueira, 2003). Each insurer has full ownership of its agent, ensuring that all strategic information remains confidential.

Today, most insurers expect a customer to request a quotation by selecting a product configuration and providing all the parameters needed to drive the quotation engine. We propose that a customer can start his interaction with the system by sending his description and needs, and not product features that he may not know how to describe.

Rather than ask every customer whether they require some coverage, and to what level, the likely need for such coverage can be determined from other customers' information, which is collected over time. This results in a dialogue with the customer that talks about needs and requirements, and does not force the customer to answer questions relating to each and every product feature. In the core of the provided brokering facility, we are using conceptual clustering procedures for grouping potential customers in meaningful classes we call *communities*. Given a set of communities we need to associate customers with a limited set of common interests. For this reason we use a metric to decide which preferences are most

representative for each community, constructing a prototypical model, that we call *stereotype*. This approach allows the insurance company to target product configurations at specific market segments, and avoids the need to ask all customers the same typically large number of product specific questions.

With learned stereotypes, the broker can advice new customers interacting with the system, to the typical insurance product configuration, selected by customers with similar characteristics. As such, customers can make requests, based on provided help, with constraints that can be satisfied by some of the insurers in the electronic market, reducing the needed effort to negotiate an insurance product.

3. ADVANCED NEGOTIATION

3.1 Negotiation Protocol

Figure 1 depicts the changes introduced in the Q-Negotiation protocol (Rocha, 2000)(Rocha, 2001), as a result of including a broker as a third participant in the negotiation process. The broker will negotiate an insurance product with participating insurers, on behalf of the customer, trying to fulfill customer's specific needs and taking into consideration customer's preferences.

Figure 1 – Negotiation protocol

The negotiation process starts when the customer sends his initial request to the broker (1). The broker sends an announcement, with customer's personal characteristics and needs, to all the insurers that can provide the policy requested by the customer (2).

Insurers reply with initial bids (3). The broker, using a multi-issue function, evaluates received bids and selects the most favourable ones. These are then sent to the customer (4).

The customer, with the help of a exploration tool provided by its customer agent, analyses received proposals and selects the insurers he wishes to continue to negotiate. He may also specify a counter-proposal (5). The broker, now taking into consideration customer's counter-proposal, again evaluates selected proposals. The broker then sends a qualitative feedback to those proposals (6). This feedback helps insurers in their task of formulating new proposals by giving them some hint about the direction they should follow in their negotiation space, in next negotiation rounds (7)(8).

This negotiation ends when a deadline, specified by the customer, is reached or a set of satisfactory proposals is received. The best-received proposals during negotiation are sent to the customer (9). If the customer selects one proposal (10), the broker informs insurers of the acceptance/rejection of their proposals (11).

3.2 Multi-attribute Bid Evaluation

Generally, an evaluation of a received proposal is a linear combination of the attributes' values weighted by their corresponding importances. In this way, a multi-attribute negotiation is simply converted in a single attribute negotiation, once there is a single result of the combined evaluation.

The multi-attribute function presented in the following equation encodes the attributes' and attributes values' preferences in a qualitative way and, at the same time, accommodates attributes intra-dependencies.

$$Evaluation = \frac{1}{Deviation}$$

$$Deviation = \frac{1}{n} * \sum_{i=1}^n \frac{i}{n} * |dif(prefV_i, V_i)|$$

$$dif(Pr efV_i, V_i) = \begin{cases} \frac{V_i - prefV_i}{\max_i - \min_i}, \text{if continuous domain} \\ \frac{Pos(V_i) - Pos(prefV_i)}{nvalues - 1}, \text{if discrete domain} \end{cases}$$

where n is the number of attributes that defines a specific insurance product component.

A proposal's evaluation value is calculated by the broker, as the inverse of the weighted sum of the differences between the optimal $prefV_i$ and the proposed value V_i of each one of the attributes. The proposal with the highest evaluation value so far is the winner, since it is the one that contains the attributes' values more closely related to the preferred ones, from customer's point of view.

The function $diff(prefV_i, V_i)$ quantifies for an issue i , the degree of acceptability of the current value V_i proposed by a specific insurer when compared to its preferable value $prefV_i$. If the insurance product component's values domain is of continuous values, this quantification simply is a normalised difference between the two values V_i and $prefV_i$. If the insurance product component's values domain is of discrete values, the difference is now calculated as a normalised difference between the preference attached to V_i and $prefV_i$. This difference can be calculated considering the relative position of the two values in the ontology's enumerated domain values specification.

The negotiation process is considered as a set of rounds where insurers concede, from round to round, a little bit more trying to approach customer's preferences, in order to be selected as the winning insurance company. The winner bid in the current round is selected as the one that presents the highest evaluation value, since it is the solution that contains attributes values the closest to the preferable ones. Then, current round's winner bid is compared with the bid in all past rounds, and the best on is selected.

The broker helps insurers in their task of formulating new proposals by giving them some hints about the direction they should follow in their negotiation space. The broker gives these hints as comments about attributes' values include in current proposals. This *qualitative feedback* reflects the distance between the values indicated in a specific proposal and the best one received so far, and is formulated as a qualitative comment on each of the proposal's attributes values, which can be classified in one of three categories: *sufficient*, *bad* or *very_bad*. The broker compares a particular proposal with, not customer's optimal values, but with the best proposal received so far because its more convincing to say to an insurer that there is a better proposal in the market than saying that its proposal is not the optimal one.

Insurers will use this feedback information about its past proposals, in order to formulate, in the next negotiation rounds, new proposals trying to follow the hints included in the feedback comments. The negotiation ends when a deadline is reached or a set of satisfactory proposals is received. The winner bid is the one that presents the highest evaluation value among all bids received until then.

3.3 Learning in Proposals Formulation

The Q-Negotiation algorithm uses a reinforcement learning strategy based in Q-learning (Watkins, 1992) for the formulation of new proposals. Q-learning enables on-line learning, which is a very important capability in our insurance domain, enabling insurers to learn in a continuous way during the negotiation process, with information extracted from each negotiation round, and not only when the negotiation process is over.

The adaptation of the Q-learning algorithm, included in the Q-Negotiation algorithm, has lead to the inclusion of two important features detailed in (Rocha, 2000). The reward value for a particular state is calculated according to the received

qualitative feedback. This reward's value will be smaller for proposals with less favourable classifications.

The exploration space, which can become very large and thus implies a long time to learn, is reduced in order to include those actions that can be applied to a previous proposal following the hints included in the feedback formulated by the broker.

4. CLUSTERING TO LEARN CUSTOMERS' PREFERENCES

The customer wishes to buy an insurance policy to cover certain risks. His requirement will usually be incomplete and uncertain and, possibly, the customer will not be aware of all the options available and may be prepared to compromise on certain aspects. This difficulty in defining an initial request has led us to try to learn customers' preferences in order to help and advise them and to appropriately deal with insurance companies always taking customers' preferences into consideration.

However, we are not only interested in acquiring models of single customers interacting with the brokering system. We pretend to identify different communities in the total population of observed customers. Also, in order to effectively use this communities, we need to construct a model of each community, representative of its users and significantly different from other customer communities. For that, two issues must be addressed: (i) how to determine meaningful subsets (*communities*) and (ii) how to determine meaningful concepts for each subset (*stereotype*).

The main distinction in machine learning algorithm paradigm is between *supervised* and *unsupervised* learning. We propose that customer communities can be automatically constructed using an unsupervised learning method, namely with *conceptual clustering* techniques. A conceptual clustering algorithm accepts a set of object descriptions (in our case, customer data) and produces a classification scheme over the observations. These algorithms do not require a teacher (i.e., they are unsupervised) but they use an evaluation function to discover classes with "good" conceptual descriptions.

Our insurance brokering system is then using COBWEB (Fisher, 1987) as a tool for grouping customers in meaningful classes we call communities. COBWEB is an incremental clustering algorithm that employs the concept of *category utility* (Gluck, 1985) to create a clustering that maximises inter-cluster dissimilarity and intra-cluster similarity. The algorithm incorporates objects into the concept hierarchy using four clustering operators: placing the object in an existing cluster, creating a new cluster, combining two clusters into a new one (merging) and dividing a cluster (splitting). Given a new object, the algorithm applies each of the previous operators and selects the hierarchy that maximises category utility. The clusters generated by COBWEB represent well-defined customer communities.

However, besides the customer data used for incrementally build up customer communities, a customer is characterised by its own preferences and by the insurance configuration chosen in the negotiation process. In order to help customers defining their needs, the broker must associate a "typical" insurance configuration to each constructed community. Thus, the natural way to define meaningful stereotypes associated to constructed communities is by trying to identify patterns in the insurance configurations chosen by such users, that are representative of the

participating users' preferences and significantly different from other communities of users.

Our system is using a metric to measure the increase in the frequency of a specific preference or negotiation result within a given community, as compared to the default frequency in the whole number of available observations (Paliouras, 1998). In (Paliouras, 1999) and (Weinberg, 1998) the increase in frequency was used as an indication of the increase in the predictability of a feature (a given preference, for example) within the community. Given a component c (a user preference), with the default frequency f_c , if the frequency of this component within a community i is f_i , the frequency increase is defined as a simple difference of the squares of the two frequencies:

$$FI_c = f_i^2 - f_c^2$$

When FI_c is negative there is a decrease in frequency and the corresponding component is not representative of that community. A community's representative characteristic is found through $FI_c > \alpha$, where α is pre-established as the required threshold for considering that frequency increase enough relevant.

The frequency increase metric has the ability to filter out either components chosen by too many customers or components chosen only by an insignificant number of customers. Filtering out this to types of preferences is a positive feature of the used metric (Paliouras, 1998), greatly improving the quality of the generated stereotypes.

5. CONCLUSIONS

This paper proposes a distributed multi-agent system where customers are grouped together, exploiting user modelling and machine learning techniques. These communities can be used to facilitate the negotiation of an insurance product. However, this can be done effectively, only when the generated communities are meaningful. That's why our system is using a specific metric to characterise the generated communities of users.

Several other researchers have also addressed learning algorithms for agents in electronic commerce negotiation. Reinforcement learning algorithms have been used in the work reported in (Oliveira, 1999), where simulation results showed that learning agents performed better than their non-learning competitors agents in different market situations. However the algorithm only addresses single-issue type of negotiations. Learning in multi-issue negotiation is addressed in (Zeng, 1996), where negotiation occurs in a one-to-one basis. Our work combines both one-to-many and multi-issue negotiation characteristics. Learning is an important characteristic that should be available in automated e-commerce negotiation. Past proposals can, and should, constraint the value of the next insurer proposal. Moreover, in open systems, as it is the case of e-commerce, learning becomes a powerful capability for deal with the environment dynamics. Therefore, we consider both multi-issue and adaptive negotiation features as being of great importance for agent-mediated electronic commerce.

6. REFERENCES

1. J. B. Weinberg, G. Biswas and D. Fisher. Iterate: A conceptual clustering algorithm for data mining. *IEEE Transactions on Systems, Man and Cybernetics*, 28:100-111, 1998.
2. M.A. Gluck and J.E. Corter. Information, uncertainty and the utility of categories. In *Proceedings of the 7th Conference of the Cognitive Science Society*, pages 283-287, 1985.
3. R. Guttman, A. Moukas and P. Maes. Agent-mediated electronic commerce: a survey. *Knowledge Engineering Review*, 13(2):147-159, 1998.
4. J. Moore. Implementation, contracts and negotiation in environments with complete information. *Advances in Economic Theory*, 1, 1992.
5. Luís Nogueira and Eugénio Oliveira. A multi-agent system for e-insurance brokering. *Agent Technologies, Infrastructures, Tools and Applications for e-Services*, LNAI 2592, pages 263-282, 2003.
6. Eugénio Oliveira, José Fonseca and Nick Jennings. Learning to be competitive in the market. In *American Association of Artificial Intelligence Workshop on Negotiation*. 1999.
7. G. Paliouras, V. Karkaletsis, C. Papatheodorou and C. Spyropoulos. Exploiting Learning Techniques for the Acquisition of User Stereotypes and Communities. In *Proceedings of the International Conference on User Modelling (UM '99)*. 1999.
8. G. Paliouras, C. Papatheodorou, V. Kakaletsis, C. Spyropoulos and V. Malaveta. Learning User Communities for Improving the Services of Information Providers. LNAI 1513, pages 367-384. 1998.
9. Ana Paula Rocha and Eugénio Oliveira. Adaptive Multi-issue Negotiation Protocol for Electronic Commerce. *Proceedings of The Fifth International Conference on The Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM 2000)*. 2000.
10. Ana Paula Rocha and Eugénio Oliveira. Agents Advanced Features for Negotiation in Electronic Commerce and Virtual Organisations Formation Process. LNAI 1991, pages 77-96. 2001.
11. Nir Vulkan and Nick Jennings. Efficient Mechanisms for the Supply of Services in Multi-Agent Environments. In *Proceedings of the First International Conference on Information and Computation Economics*. 1998.
12. Christopher J.C.H. Watkins and Peter Dayan. Q-Learning. *Machine Learning*, 8(3-4): 279-292. 1992.
13. M. Wellman and P. Wurman. Market-aware Agents for a Multiagent World. *Robotics and Autonomous Systems Journal*, 24: 115-125. 1998.
14. Dajun Zeng and Katia Sycara. How can an Agent Learn to Negotiate. *Proceedings of the Third International Workshop on Agent Theories, Architectures and Languages*. 1996.
15. Fisher, D. H., 1987. Knowledge Acquisition via Incremental Conceptual Clustering. *Machine Learning*, Vol. 2, pp 139-172.