

Evaluation of Suggestions during Automated Negotiations

Sarit Kraus

Institute for Advanced Computer Studies and
Department of Computer Science
University of Maryland

Eithan Ephrati Daniel Lehmann
Department of Computer Science
The Hebrew University

Abstract

An automated agent that has to act in a multi-agent environment needs the capability to negotiate. In this paper we concentrate on problems that arise while evaluating suggestions during negotiations. We distinguish between different kind of suggestions and present methods and techniques for evaluating them. The suggestions are written using a formal Negotiation Language that we have developed. We show how our approach was successfully implemented in a specific environment: the Diplomacy game. As in other board games, playing Diplomacy involves a certain amount of technical skill but the capacity to negotiate, explain, convince, promise, keep promises or choose not to keep them, is an essential ingredient of good play. Diplomat was evaluated and consistently played better than well experienced players, and in games that were held, many players did not guess which player Diplomat was playing.

INTRODUCTION

Negotiations are part of everyday life. Negotiators try to reach an agreement or arrangement by discussion. Therefore, an automated agent acting in a multi-agent environment needs a capability to negotiate. The need for negotiations arises in an environment where cooperation is beneficial and becomes even more necessary where there also exist conflicts between the agents. Such an environment is common in the real world. Even in distributed cooperative systems, where all agents are designed to achieve the same goal, conflicts between the agents can arise, since every agent makes an effort to perform its mission in a way that may interfere with the other agents' activities. An agent will get even more benefits through negotiations if some or all of the following conditions exist:

1. The cooperation among negotiators requires precise cooperation or division of tasks between the participants.
2. The information the agents have is incomplete and the negotiations enable the agents to gain more information.
3. There are difficult problems to analyze and a single agent is unable to solve its problems alone, but its problem solving power may be multiplied by cooperation and exchange of ideas.
4. There is a multi-agent environment where the other human (or automated) agents have the capability to negotiate and an agent that does not have the capability is at a disadvantage.

In previous work, Rosenschein and Genesereth [Rosenstein & Genesereth, 1985], used certain game-theoretic techniques to model communication and promises in multi agents interaction. There, the process of negotiation was severely restricted (the agents could only make single, simultaneous offers), and it assumes that each agent knows the complete payoff matrix associated with the interaction. Also, for large games involving many agents and outcomes, the kind of environments in which we are interested, the size of a payoff matrix may quickly become intractable.

Davis and Smith [Davis & Smith, 1983] proposed an approach to cooperation using a contract-bid metaphor to model the assignment of tasks to processors. They used negotiation to match idle problem solvers to outstanding tasks as a basis for the transfer of control and as a way of viewing invocation as the matching of knowledge sources to tasks. Sathi, Morton, and Roth, ([Sathi, Morton & Roth 1986]), considered the problem of project management. In their approach the agents negotiate by relaxations so as to achieve a compromise. The constraints and their relaxations are statically known. No attempt is made to influence other agents' relaxation. Sycara [Sycara, 1987] presented a model of negotiation that combines case-based reasoning and optimization of the multi-attribute utilities of the agents. She implemented her ideas in a computer program called the PERSUADER that resolved adversarial conflicts in the domain of labor relations and tested her system using simulations of such domains.

Comparing this to our work, we examine negotiations in a more complex environment where a mediator is not available, where the agents may break their promises, where close cooperation between different agents is needed, and where possible coalitions between other agents must be taken into account. We have implemented our ideas by building a system that negotiates successfully with human partners.

ENVIRONMENTS AND CONCEPTS

Sending a suggestion is the way to propose a plan for acceptance or rejection. Actually, it is the main method used to try to and convince another agent in the environment to perform

actions, and it is the way to promise to the other agent, which actions sender will perform. Therefore a negotiator-agent must have the capability to evaluate suggestions and decide whether to accept or reject the suggestions.

Before we get into the details of the evaluation of suggestions, we will briefly describe, the environment and define some concepts that will be used later. In general, the negotiations are performed among a set of agents. Each agent has its own goals and tasks and searches for the best plan to achieve those goals. Each one of the agents negotiates in order to influence the behavior of the other agents' activities, in order to convince them to help him, or, at least, not to interfere with his plans.

We assume that a set \mathcal{S} of strategies is given. The negotiations are performed on the basis of the strategies for common activities of the parties that negotiate. A strategy includes a list of activities, the purpose of each activity, and the expected profit or loss for the agents from those activities. \mathcal{S} may be finite or infinite, but in the environments we are considering, \mathcal{S} is usually infinite (or at least very large).

In order to negotiate effectively, the negotiator-agent must keep and maintain general information about the environment and the current situation. It also has to keep information about the other agents, their personality and their relations and to keep the details of its agreements with the others. We will denote the negotiator-agent knowledge and beliefs base by KBB.

The personality of the human agents influences their behavior while bargaining. Especially, their willingness to take chances must be taken into consideration while evaluating their possible acts ([Bueno de Mesquita, 1981]). Their loyalty has to be considered when deciding whether the other agent will keep a signed agreement. Therefore the negotiator-agent has to try to estimate the personality of the other agents by examining their activities and by exchanging messages with the others. We allow the behavior of the negotiator-agent to be influenced by some "personality" traits such as aggressiveness, willingness to take chances and loyalty, that will be given to it at startup time or during the negotiations, and they will allow the negotiator-agent to change "personality" from time to time.

To make automated negotiations a bit easier, the negotiations need to be performed using a formal language. We have developed a Negotiation Language that includes four kinds of messages: declarations, questions, suggestions and answers ([Kraus, 1988], [Kraus & Lehmann, 1988a]). The building blocks for the messages are simple sentences that are specific to the subject of the negotiations.

Diplomacy

The environment we choose to deal with is a game named Diplomacy, marketed by Avalon Hill Company. Diplomacy is an environment of intense negotiating situation where we could really fit an automated agent against humans without the later knowing. We implemented and tested our ideas by building an automated Diplomacy player called Diplomat (general description in [Kraus, 1988], [Kraus & Lehmann, 1988a], and [Kraus & Lehmann, 1988b]).

Diplomacy is a board game played on a map of Europe during the years just prior to World War I. Each player plays one of seven European powers. The moves of the game are figured as two moves each year: a Spring move and a Fall move, beginning in year 1901. After negotiation, each player privately writes down the orders for all of his units. A unit may be ordered to do only one thing on each season: to hold, move, or to give support. A fleet may be also ordered to convoy another army from one coast to the other. The power that gains control over Europe wins the game. To be a good player one needs some technical skills in moving military units on the board according to the reasonable but complex rules of the game, but above all, one needs the ability to communicate and negotiate with the other players, to make agreements with the others and possibly to decide to break these agreements, since the rules do not bind a player to anything he says. Deciding whom to trust as situations arise is part of the game. Details of the rules of the game can be found in [Rules for Diplomacy, 1984]

Diplomacy satisfies the conditions that increase the need for negotiations that were mentioned in the introduction: it is a repeated game of incomplete information, certain moves require close cooperation between different allied powers (the units of a power may help the moves of another power but they must be explicitly ordered to do so), it is a very complicated game (average of 34^8 possible orders for a season), and the other players negotiate, since experience shows that a power can not last long without taking part in extensive negotiation.

Diplomat is implemented in Ylisp ([Levy & Dimitrovski 1982]) (a dialect of Franz-Lisp) on a Vax 11/785 running Unix, Berkeley 4.3. The current version is the product of three years work by three programmers. It includes over 10,000 lines of Ylisp code.

SUGGESTIONS EVALUATION

One of the main technical problem in evaluating suggestions by a negotiator-agent is the large number of possible suggestions and the fact that different kinds of suggestions need to be evaluated in different ways. In order to be able to solve this problem and to make the negotiations more modular a negotiator-agent can distinguish mainly between two kinds of suggestions: general suggestions and detailed suggestions. A general suggestion discusses the general purpose of the negotiations, as where a detailed suggestion discusses the specific common activities and the ways to achieve those common activities. A general suggestion can include the kind of agreement to be achieved (in Diplomacy, a cooperation or non-aggression

<MESSAGE 1 FROM England TO France> :
I would like to suggest to you a Cooperation Agreement
between England and France against Germany now.
END OF MESSAGE.

Figure 1: Spring 1901: General Suggestion

< MESSAGE 2 FROM Turkey TO Russia > :
I would like you to know the following facts:
Russia seems to be strong, (1)
and Turkey will attack Serbia now, (2)
and Russia will help Turkey's attempt to enter Serbia, (3)
and Russia will move from Warsaw (inland) to Galicia (inland). (4)
END OF MESSAGE.

Figure 2: Spring 1902: Detailed Suggestion

agreement), relations with other agents (in Diplomacy, possible common enemies), or general directions of possible common activities (in Diplomacy, the directions of possible common attacks). For an example see the message in Figure 1 which provides an example of a general suggestion. The detailed suggestions can also be divided as follows:

1. Suggestions concerning general purposes of an agreement (In Diplomacy: spaces on the board to attack, spaces to defend, areas to leave or to enter.)

For example, see sentences (2) and (3) in Figure 2.

2. Suggestions about the specific movements in order to achieve the purposes of 1. (in Diplomacy which unit has to attack or to give support, which fleet can convoy, and which unit has to move in order to cut support). For example, see sentence (4) in Figure 2. We have developed different methods for evaluating different kind of suggestions. The General Suggestion Evaluator (GSE) (see Figure 3) evaluates general suggestions using very fast heuristics, because it is not worthwhile to spend a lot of time searching for good detailed strategies in this step of the negotiations. The principle directing the GSE is that if there is some hope of profiting from the suggested agreement, it is worth continuing the negotiations.

The Detailed Suggestion Evaluator (DSE) (see Figure 3) evaluates detailed suggestions by finding fitted strategies to them. Before sending the suggestion to one of the above modules, we move it through a Pre Analyzer (PA) that will fill in gaps of missing information.

Suppose a negotiator-agent gets a message which includes a detailed suggestion and must decide how to respond. After passing the message through the PA and filling in the gaps, the next step, which is done by the DSE, is to try to evaluate the expected profit (or losses) for the negotiator-agent and its partners from this suggestion. The DSE translates any suggestion into a set of strategies that *fit* the given suggestion, taking into account the current situation and the beliefs of the agent about the other agents and the environment. This translation changes each suggestion received or sent by the negotiator-agent into a unique strategy format, which allows it to be compared with other suggestions or strategies. The translation is done using the Strategies Finder of the negotiator-agent, and the strategy found is used as a basis for further negotiations when needed. The next step is to examine the expected profit from the fitted strategy and compare it with the expected profits from other possible strategies.

A strategy *fits* a suggestion when all the specific activities that are precisely mentioned in the suggestion appear in the strategy and the strategy does not include activities whose negations appear in the suggestion. If general activities are mentioned in the suggestion, the strategy must include at least one order that implements every such general activity. If the parties had agreed upon some details during the previous steps of the negotiations, where messages were exchanged between the parties, and the current suggestion does not contradict those details, then the *fitted* strategy has to fit those details too. So the negotiator-agent assumes that a detail of an agreement is valid until the other party says the opposite or does something that contradicts this assumption. We found out from human negotiators that they make the same assumption, and an automated agent that negotiates with humans should make it, too.

When the set of strategies, \mathcal{S} , is finite and small the implementation of the DSE is easy. The agent may check all the strategies and find which of them fit the suggestion. However, in the environments on which we concentrate \mathcal{S} is usually too big to be computed and stored, so the negotiator-agent must use heuristic methods to find strategies that fit a suggestion.

CONCLUSION

We have proposed methods for an automated negotiator to evaluate suggestions in a complex environment, where the set of possible strategies is very large, a mediator is not available, the agents may break their promises, close cooperation between different agents is needed and possible coalitions between other agents must be taken into account.

We used our methods to develop the system Diplomat, which plays Diplomacy as one of the players. Diplomat successfully evaluates more than 70% of the suggestions it gets, and answers its partner properly. In the rest of the cases it asks for more details. Also, in its other missions, Diplomat performs well. For example, we examined 63 agreements that were

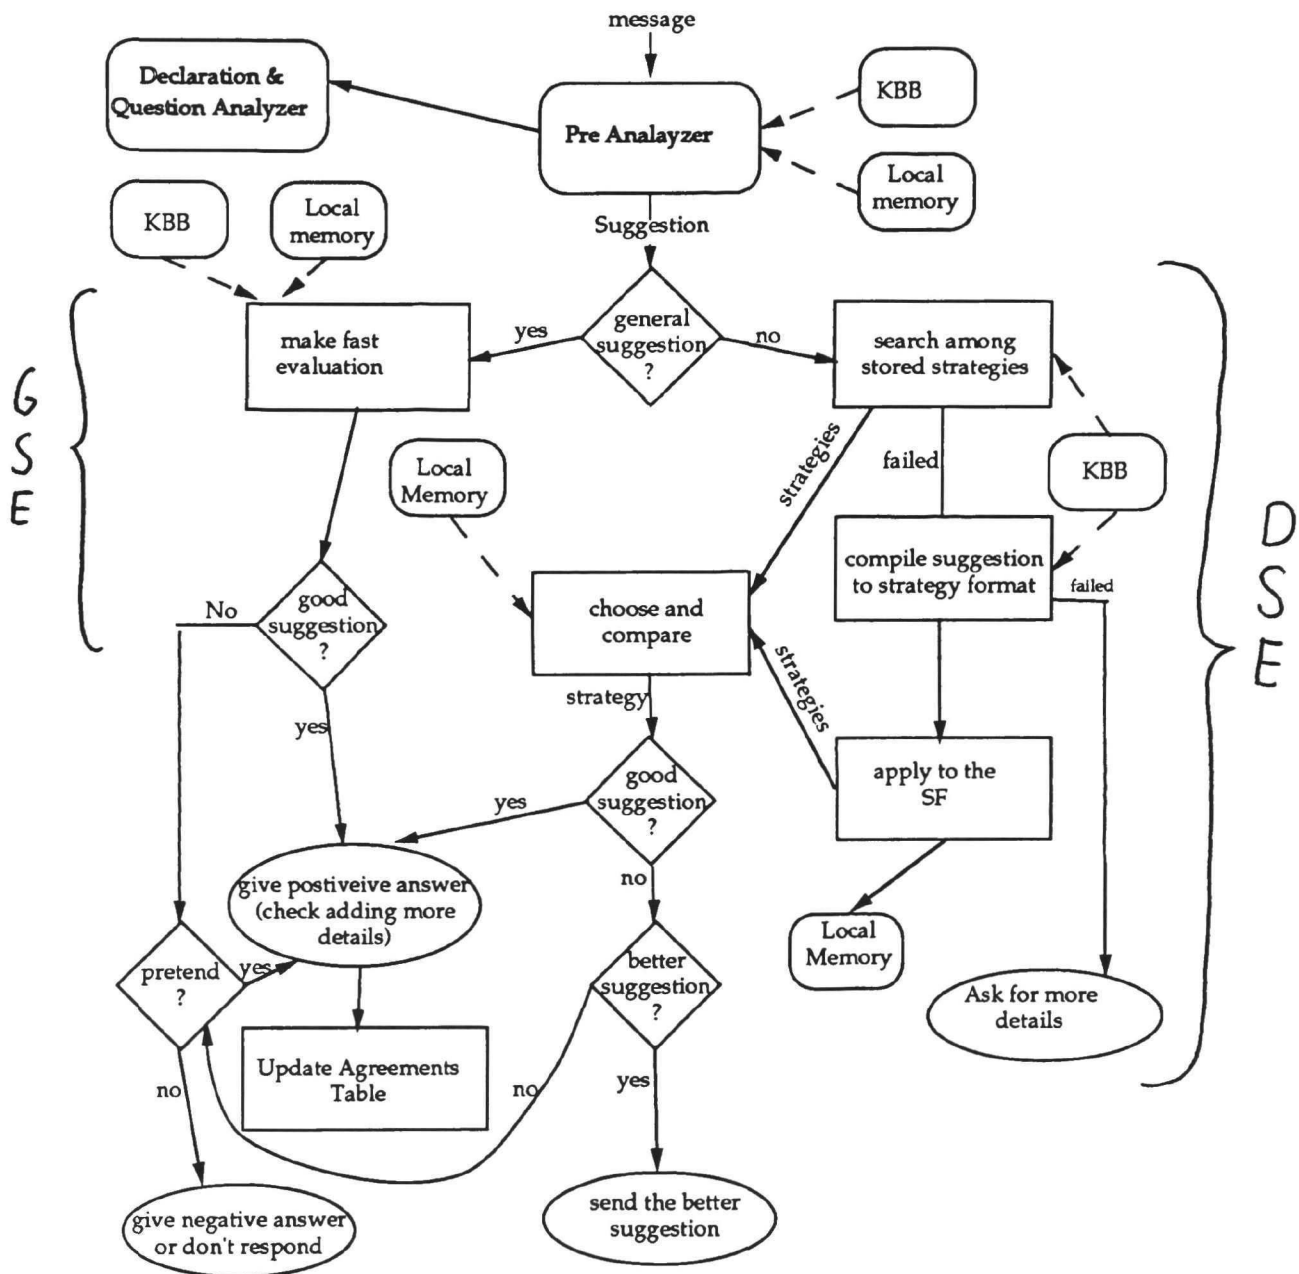


Figure 3: General Description of the Suggestion's Evaluation

KRAUS, EPHRATI, LEHMANN

signed between Diplomat and another power, and Diplomat predicted successfully the ally's intention to keep an agreement in 92% of all the agreements it had signed. We actually tested Diplomat in 100 Diplomacy seasons and, determined that Diplomat plays better than a well experienced player in a way that is difficult to distinguish from a human player. This is because Diplomat negotiates like human players, break agreements as they do and shows good strategic skills.

Further work will be to extend the used of the tools and the ideas of this paper to other domains.

REFERENCES

- Rules for Diplomacy* (1984). The Avalon Hill Game co.
- Bueno de Mesquita, B. (1981) *The War Trap*. Yale University.
- Davis, R. & Smith R.G. (1983) Negotiation as a metaphor for distributed problems solving. *Artificial Intelligence*, 20:63–109.
- Kraus, S. (1988), *Planning and Communication in a Multi-Agent Environment*. PhD thesis, Hebrew University, Jerusalem, 1988. (Written largely in Hebrew).
- Kraus, S. and Lehmann, D. (1988a), *Automated Negotiator*. Technical Report 88-7, Leibniz Center for Computer Science, Hebrew University, Jerusalem.
- Kraus, S. and Lehmann, D. (1988b), Diplomat, an agent in a multi-agent environment: an overview. In *Proc. of the Seventh Annual IEEE Phoenix Conference on Computers and Communications*, pages 434–438, Arizona.
- Levy, J. & Dimitrovski, Y. (1982) *The Ylisp 2 Manual*.
- Rosenschein J. R. & Genesereth M. R., (1985) Deals among rational agents. In *Proc. of the Ninth International Joint Conference on Artificial Intelligence*, pages 91–99, California.
- Sathi, A., Morton, T.E.& Roth, S.F., (1986) Callisto: an intelligent project management system. *The AI Magazine*, 7(5):34–52.
- Sycara, K. P. (1987) *Resolving Adversarial Conflicts: An Approach Integrating Case-Based and Analytic Methods*. PhD thesis, School of Information and Computer Science, Georgia Institute of Technology.