

Argumentation: Planning Other Agents' Plans

Katia P. Sycara
The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Most work in Distributed AI has dealt with agents that cooperate to achieve common high level goals. The assumption of cooperative behavior allows agents to use static models of one another to predict their future actions in order to promote coherent system behavior. Our research extends the domain of problems to include non-cooperative, multi-agent interactions where cooperation cannot be assumed but needs to be dynamically induced during problem solving. This necessitates the development of mechanisms to purposefully *modify* the plans, goals and behavior of other agents to increase agent cooperativeness in order to bring about convergence to a global solution. We advocate persuasive argumentation as such a mechanism for cohering the group problem solving of non fully cooperative agents. We present a model of persuasive argumentation that has been implemented as part of the PER-SUADER, a multi-agent computer program that operates in the domain of labor negotiations.

1. Introduction

Most work in Distributed AI has dealt with groups of agents pursuing common goals (e.g., [Durfee 87, Lesser 83, Cammarata 83]). Agents in such systems work as a team. Their interactions are guided by cooperation strategies meant to improve their common performance. In networks of cooperative agents, agent interactions take the form of communicating plans and goals at an appropriate level of abstraction. These communications enable the receiving agents to form expectations of the future behavior of the sending agent, thus improving agent predictability and network coherence [Durfee 87]. A consequence of the cooperativeness assumption is that the recipient agent will use the information in the plan to adjust *its own* local plan-

ning appropriately, so that the common planning goals (and planning effectiveness criteria) will be met. No attempt needs to be made to change the intentions or actions of other agents. Although the abstracted communicated information forms a basis for common representational regimes, in complex domains, such as scheduling and resource allocation, goal conflicts might still occur since the communicated information needs to be locally interpreted by the agents. Different interpretation of global goals, for example, can lead to conflict or competition among the agents. Thus, simply communicating planning information (even in cooperative situations) cannot deal with (a) disparate expertise, (b) inconsistent beliefs, and (c) goals conflicts.

In environments where each agent is pursuing his own goals (e.g., design experts responsible for different parts of a complex artifact, different departments in an organization), cooperative behavior cannot be taken for granted. It has to be planned for by each agent and has to emerge as a result of the agent interactions. The planning process for coordinating non fully cooperative agents is *negotiation*. In contrast to other work [Davis 83] that has not dealt with belief modification, in our work, negotiation is an iterative process that involves *identifying* potential interactions either through communication or by reasoning about the current states and intentions of other agents in the system and *modifying* the intentions of these agents so as to avoid harmful interactions or create cooperative situations.

In order to negotiate effectively, agents need the ability to (a) represent and maintain belief models, (b) reason about other agents' beliefs, and (c) influence other agents' beliefs and behavior. Since an agent cannot be sure of the full cooperativeness of another, the planning information communicated cannot simply be his own high level plans. It needs to take into consideration explicitly the model and plans of the recipient agent. The information communicated is intended to "convince" the recipient agent to cooperate, thus reducing conflict and promoting global coherence. To perform this task, an agent reasons about another agent using its own model of that other agent, finds as many ways as the model will allow to affect the other agent's outcomes (behavior), and uses them selectively to influence the other agent. This is the process of *persuasive argumentation*.

¹This research has been sponsored in part by the Army Research Office under contract No. DAAG 29-85-K0023 and in part by NSF Grant No. IST-8317711.

Persuasive argumentation is advocated as a general mechanism to promote coherence of non fully cooperative agents. Persuasive arguments are used by a persuader as a means to dynamically change the utilities associated with various plans and outcomes of another agent, the persuadee, so as to increase the willingness of the persuadee to cooperate. This, in turn, improves the efficiency of convergence to a global solution. By observing reactions to the arguments, the sending agent can update and correct its model of the recipient agent, thus refining its planning and argumentation knowledge. Persuasive argument generation is based on integration of goal graph search, use of multi-attribute utilities, and availability of a case memory of experiences with similar persuadees [Sycara 87].

In contrast to work [Appelt 85, Konolige 86, Cohen 87, Halpern 84]. that has primarily dealt with the development of formalisms for the representation of beliefs, belief deduction, or update of the belief base, our work concentrates on heuristic mechanisms (persuasive arguments) to *bring about changes* in other agents' beliefs and behavior. While others have worked on argumentation, none so far has worked on persuasive arguments. The work of Flowers, et al. [Rowers 82] was concerned with adversary arguments where no attempt is made to change the beliefs of the interlocutor, and Ashley and Rissland [Rissland 86] have examined the use of hypotheticals to test the strengths and weaknesses of case-based legal arguments.

We present a general theory of persuasive argumentation that we illustrate in the domain of labor mediation. Our argumentation model is part of a general multi-agent, multi-issue negotiation model [Sycara 87, Sycara 88a], that has been implemented in a computer program, the PERSUADER. Negotiation in the PERSUADER is performed through integration of Case-Based Reasoning [Kolodner et al. 85, Sycara 87] and use of multi-attribute utilities [Sycara 88b]. The PERSUADER system involves three agents: a company, its union and the mediator whose task is to help the other two agents reach an acceptable compromise. The mediator is engaged in parallel negotiations with the union and company agents. The PERSUADER'S input is the set of conflicting goals of a company and union and the dispute context. The final output is either a single plan in the form of an agreed upon settlement (contract) or an indication of failure if the parties to the dispute did not reach agreement within a particular number of proposals (to simulate the inability of parties in the real world to reach agreement before a strike deadline). A contract that is proposed by one agent to another is a plan that the proposing agent will follow, contingent on the second agent's agreement. Proposed contracts in the PERSUADER are subject to negotiation and possible modification, unlike bids in the contract net [Davis 83]. The negotiation process consists of iteration of three main tasks: generation of a proposal, generation of a counterproposal based on feedback from a dissenting party, and persuasive argumentation. The mediator generates an initial compromise proposal and presents it to both the union and

company who evaluate the proposal from their perspectives and give the mediator their reaction. If both accept the proposal, then it is the final compromise. If one of the agents rejects it, the mediator makes a decision whether to change the proposal or attempt to change the disagreeing party's position.

Belief and belief modification in the PERSUADER is based on the conjunctive goals of the agents and their interactions. A belief in the PERSUADER involves the correspondence between a state (a possible settlement) and the other agents' actions. Group knowledge [Genesereth 87] in the PERSUADER focuses on the facts of the case: proposals, counterproposals, negotiation context etc. If agreement on a compromise were obtainable by inference from these facts, negotiation would be unnecessary. Such is not the case, however, since the goals, plans and utilities of the agents are largely unknown and the evaluation function is distributed. The negotiation process itself is a search of a dynamic problem space where an agent's beliefs about other agents' beliefs and hence feasible solutions continuously changes the space being searched. What was not an acceptable solution at one point becomes a solution at a later point.² This occurs as a consequence of the agents' realization of the inevitability of partial goal satisfaction through acceptance of a compromise. The PERSUADER attempts to influence the process toward solution convergence by constructing arguments to bring about belief states of the agents that are necessary in achieving compromise states (global system solutions).

2. Representation of other agents' models in the PERSUADER

The PERSUADER integrates two representational vehicles to model agents: graphs and multi-attribute utilities. Both of these are representations at a high level of abstraction that reduces the overhead of constructing and updating agent models. As opposed to game theoretic approaches [Rosenschein 85] that assume common knowledge of the payoff matrix (an unrealistic assumption if the agents are not fully cooperative), models of agents based on approximation of the agents' utilities are more appropriate for negotiation, since they do not assume common knowledge of payoffs, incorporate reasoning about tradeoffs, and allow flexibility in combining individual utilities.

The main factor determining the effectiveness of arguments of persuasion is the attitudes and beliefs of the persuadee [Abelson 59]. A persuadee's *belief structure* in the PERSUADER involves his goals, the importance he attaches to them and relations between them. Different arguments change different parts of the belief structure. A persuadee's beliefs are represented in a directed acyclic graph, which is searched during argument generation and

² In a labor negotiation, for example, it is unlikely that either party would accept their eventual compromise, if it were presented at the inception of negotiations.

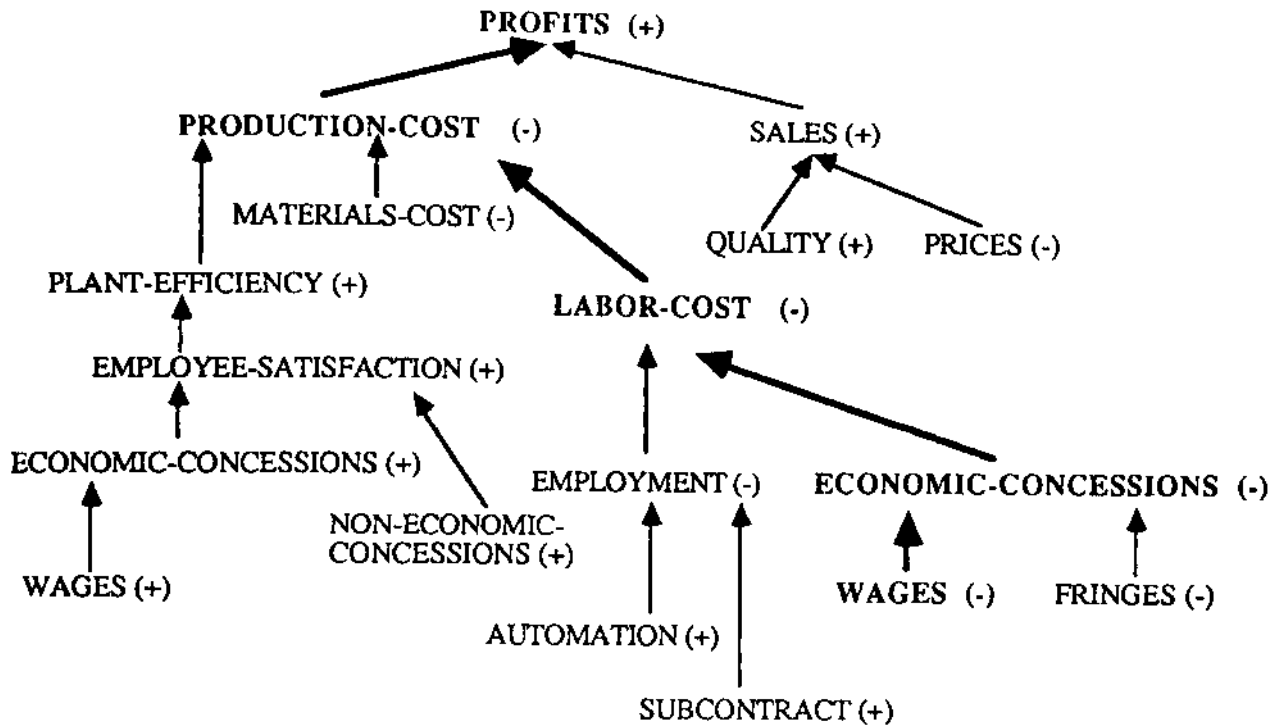


Figure 2-1: A company's partial belief structure

updated according to the persuader's reaction to an argument. The nodes represent goals with the associated *importance*, perceived *feasibility* on the part of the persuadee of achieving the goal, desired *direction of change* (increase, or decrease), and the *amount* by which the quantity should be changed. The arcs represent the *contribution* of a goal to the achievement of its parent goal. A negative contribution value has the interpretation that the subgoal not only does not contribute to the higher level goal but it is detrimental to it. For example, a negative public image would have a detrimental effect on a company's sales. A positive value means that the subgoal supports the achievement of the higher level goal. A contribution value of zero means that the subgoal is irrelevant to the achievement of the higher level goal. In terms of the graph, this is equivalent to removing the edge between the subgoal and the goal.

Figure 2-1 presents a partial view of a company's belief structure. In order not to clutter the figure, we are showing the graph as a tree depicting only the goals and the desired direction of change.

A path X to Y in a belief structure constitutes a causal chain that provides an explanation of the change in Y in terms of the change in X, assuming no other change has occurred in the rest of the tree. The path WAGES(-) to PRODUCTION-COST(-) in the company's goal tree can be interpreted as follows: "Other things being equal, diminishing the cost of wages results in decreasing the cost of the economic concessions, which causes a decrease in labor costs, leading to a decrease in production costs".

In addition to an agent's beliefs, the representation includes an estimation of his utilities.

The concept of utility is the basis for selecting among future alternatives and for evaluating past actions. Each alternative is evaluated in terms of a number of attributes that a decision maker considers important. Utilities express the *preference structure* of an agent. The utilities of the individual attributes are combined to give the overall utility (payoff) of an alternative. Being able to compare different alternatives enables a decision maker to choose the alternative that affords him maximal payoff. In labor negotiations, the pertinent attributes are the issues under discussion and different contract proposals are the alternative decisions. Utilities express the tradeoffs that a decision maker is willing to make among various attribute values. The payoff is expressed as a linear combination of the utilities associated with the issues.

When a persuader is faced with an unknown persuadee, he needs a way to infer the belief and preference structure of that agent. One way is by using Case-Based Reasoning to access previously encountered and "similar" persuadees and transfer relevant characteristics [Sycara 87]. The information transfer could be influenced by the problem solving context. In the PERSUADER, this context is primarily economic. For example, in recession, a union's job security goal has higher importance than in boom. When competition is stiff in an industry, a company's automation goal assumes greater importance, and also has a greater contribution towards the goal of reducing production costs.

2.1. Connecting belief change to behavior change through utilities

The argumentation model that an agent/persuader in the PERSUADER system uses, makes provisions both for selecting and adapting previously used arguments through

Case-Based Reasoning, and for constructing arguments from scratch. Regardless of which method a persuader uses for argument generation, its reasoning is guided by argumentation goals and strategies. Argumentation goals are associated with the ways that a persuadee's beliefs and behavior can be affected by an argument. Argumentation strategies are used to achieve the argumentation goals. The application of an argumentation strategy may generate more than one argument. To be effective, a persuader must select the most convincing argument for the situation at hand. The strength of an argument's justification (support) is used as the selection criterion.

We claim that a party's satisfaction with a proposition expresses his willingness to accept the proposition. Hence, if a persuader could manipulate a party's utilities (resulting in manipulation of the party's payoff), he would be able to affect the behavior (outcomes) of the persuadee. Convincing a persuadee to change his evaluation and increase his cooperativeness regarding a proposition is modeled as producing an argument to increase the payoff of the proposition. Hence, the task of a persuader can be viewed as finding the most effective argument that will increase a persuadee's payoff. Since a persuadee's payoff can be approximated by a linear combination of his utilities, the payoff can be increased by (1) changing the importance (coefficient) the persuadee attaches to an issue, and (2) changing the utility value of an issue. These constitute a persuader's argumentation goals.

Changing the coefficient of an issue in the payoff function translates into changing the corresponding goal's importance in the belief structure. The change in value at a point on an individual utility curve for an issue can be interpreted as a change of a party's assessment of the value of that issue. This corresponds to changing the "amount" parameter A in a persuadee's belief structure. Consider, for instance the situation where a company assesses an increase of 20 cents an hour in wages as "too high". In the utility theory formulation, this can be translated as "payoff(20) is LOW", (i.e., the company's payoff if it were to give a wage increase of 20 cents per hour would be some low value). In the company's belief structure, this assessment is represented as $WAGES(+, A < 20)$ (i.e. the company's goal is to give a wage increase of at most some amount less than 20 cents). Convincing the company that this increase is not so high changes its wage goal to $WAGES(-f, A=20)$. Correspondingly, this change results in raising the company's payoff (payoff(20) greater than LOW).

Two argumentation strategies can be used to change the importance that a persuadee attaches to a proposition:

- (a) indicate a change (increase or decrease) in the contribution of the present goal to a higher level goal of the persuadee
- (b) indicate a change in the feasibility (effectiveness) of the proposed goal

The argument "A new three-level job structure will be established for the technical support forces at AT&T Information Systems. This structure will be more like that of

competitors"³, is intended to increase the feasibility of establishing the new job structure by pointing out that such an arrangement has worked for competitors.

The second argumentation goal that a persuader might select is to change the persuadee's assessment of the proposed value of an issue. This second goal can be effected using the following strategies:

- (c) recall a "counterexample" from the persuadee's past behavior
- (d) recall examples of similar peers that have accepted the same value for the issue.

To illustrate strategy (c), suppose an international union maintains during contract negotiations with company X that it never signs a managements rights clause. A persuader could remind the union of its contract with company Y where such a clause was indeed present. To illustrate the last strategy, consider a union's rejection of an increase of 10 cents per worker per hour in health benefits as unacceptably low. A persuader could present contracts signed by the same or similar local unions that incorporate an equal or lower increase. This argument is effective because perception of "low" or "high" values is determined by *prevailing practice*, namely what settlements peers of the persuadee have agreed to.

3, Generating arguments

The PERSUADER generates a variety of arguments based on different argumentation goals and strategies. The argument generating process may result in producing multiple potential arguments. The most effective of the arguments produced must be chosen. The PERSUADER uses the strategy of presenting the "weakest" (less convincing) argument first, presenting "strong" arguments only when the weak ones have been rejected. We have developed a hierarchy of argument types according to their convincing power. In general, the position of an argument type in this hierarchy is domain dependent. For the labor domain, the hierarchy (from weakest to strongest) is: (1) Appeal to universal principle, (2) Appeal to a theme, (3) Appeal to authority, (4) Appeal to "status quo", (5) Appeal to "minor standards", (6) Appeal to "prevailing practice", (7) Appeal to precedents as counterexamples, (8) Appeal to self-interest, (9) Threats.

In the rest of the paper, we will present mechanisms used in the PERSUADER to generate arguments that change the importance of an issue, arguments to change an issue's perceived value, and threatening arguments.

3.1. Generating arguments to change the importance of an issue

In this section, we concentrate on arguments meant to change the contribution of a goal to a higher level goal of a

³This text is a verbatim quoted excerpt from a full-page communique by AT&T entitled "A message from AT&T to our people represented by the Communications Workers of America" that appeared in the AUanta Constitution, June 7, 1986. At that time, the CWA was on strike against AT&T.

persuadee (strategy (a), section 2.1). Consider the following example:

Wages will increase 8 percent over a three-year period⁴. Three percent the first year, three percent the second year, and two percent the third year. However, no Cost of Living Adjustment (COLA) is included in the contract. As you may know, the rate of inflation is currently running at less than one percent per year. Obviously, guaranteed increases are more meaningful.

The above argument illustrates the situation where a persuader's argumentation goal is to decrease the importance of an issue by showing decreased contribution of this issue towards a higher level goal. The union's higher level goal is WAGES(+). Subgoals to this goal are "COLA clause" and "straight percent increase". AT&T's argument is intended to show that the contribution of the "COLA clause" subgoal is less than the union had thought. Moreover, the contribution of the "straight percent increase" goal is more than the union had thought. As a consequence of the argument, the importance to the union of a COLA clause diminishes, whereas a straight wage increase becomes more important. To generate the COLA argument, a persuader needs to know what a COLA is useful for, namely that it is there to safeguard workers' income during future periods of high inflation.

After examining a great number of arguments in the labor domain, we observed that very often the reasons for the existence of a goal, denoted by g , could be cast in the following general form: A contingency C threatens a high level goal G ; a subgoal g is created to protect G from C . In the COLA example C is high inflation and G is a worker's income. A persuader can infer the high level goal toward which a goal g contributes from knowing what g is useful for. To decrease the contribution that g makes toward the achievement of the higher level goal G , a persuader simply denies the contingency C . The strategy is as follows:

To decrease the contribution of a goal to a higher level goal, assert the negation (in a qualified or unqualified manner) of the contingency C .

If the negation is unqualified, then the contribution value goes to zero. The above is a very simple (and simplistic) technique and it will produce crude arguments. Surprisingly enough, a lot of arguments in the real world, as illustrated by the AT&T message are crude and simplistic too. The negation of the contingency C acts as a justification for asserting the decreased contribution of g towards G . It is, however, a very weak justification. A persuader can strengthen the convincing power of his argument by justifying his justification.

In section 3, we have presented nine argument types that use various sources of justification (e.g., universal principle, status quo, authority). Which justification a persuader will choose for each situation is governed in our work by a set of

heuristics that are linked to the semantics of the issues/goals at hand. An example of such a heuristic is:

IF the negated contingency C deals with a possible change of a state to a future state S_f ,
THEN
(a) IF the state's present value is NOT(S_f),
then use "appeal to status quo" as justification
(b) IF the state's present value is already S_f ,
then use "appeal to authority" as justification

The above heuristic has been used by AT&T in the COLA argument. It justifies its assertion that high inflation will not happen because the *current inflation is low*. If, on the other hand, current inflation were high, AT&T could have appealed to authority and said "Noted economists are forecasting that, though the inflation is high at present, it will be drastically reduced for the next three years". (Not such an uncommon argument, either).

Persuasive arguments that *increase* the importance of the issue under discussion can be generated in the same framework as the ones to decrease the issue's importance. In this case, the persuader *asserts rather than negates* the contingency C . As an example, consider an argument to increase the importance of a COLA clause by the CWA union. "A COLA clause is important because inflation will increase in the next three years". The union can further justify its argument by appealing to authority.

3.2. Generating arguments to change a persuadee's perception of an issue's value

When the argumentation goal of a persuader is to change the persuadee's perception of an issue's value, two argumentation strategies can be used. One is to find a counterexample from the persuadee's past behavior. The second is to find evidence from the behavior of the persuadee's peers regarding the value of the issue. For both strategies, the argument generation algorithm involves search of the Case Memory for past experiences (cases) that the persuader might have had with the persuadee or his peers. For example, if a company rejects a proposed increase in wages as "too high", then a persuader can produce evidence showing that the company's competitors have given even higher wage increases. The justification for this type of argument is "appeal to prevailing practice". A persuader in the above example, could use the first argumentation strategy, if evidence can be found that in the past the company has always given big wage increases. As a matter of fact, this type of argument is used very frequently by unions seeking higher wages during negotiations. The justification is "appeal to counterexample".

The heuristic that is used to generate a persuasive argument using the first strategy is:

Retrieve past experiences of the behavior of the persuadee with respect to the present issue.
IF a past experience is found with the persuadee's behavior

⁴This text is another quoted excerpt from the AT&T communique, Atlanta Constitution, June 7, 1986.

contrary to his present claim of the issue's value,
THEN point out this discrepancy
ELSE use the second argumentation strategy

The heuristic that is used to generate a persuasive argument using the second strategy is:

1. Retrieve past experiences of similar persuadees with respect to the same issue.
2. Collect the ones whose value for the issue is appropriate
3. Present them to the persuadee

The second step needs some clarification. By "appropriate" we mean having a greater (lesser) or same value in case the objection was "too much" ("too little").

3.3. Generating threatening arguments

A threatening argument is one where goal abandonment on the part of the persuadee is sought. In the utility theory framework, goal abandonment corresponds to the coefficient (importance) of a goal becoming zero. The strategy used to produce a threatening argument is: Discover an action that endangers a goal of the persuadee that is more important than the goal targeted for abandonment and point out to the persuadee that this action will be carried out if the targeted goal is not abandoned.

Consider a negotiation between a company and its union, where the PERSUADER has suggested a compromise. The company has agreed saying that the increase in fringe benefits (fringes) is the highest it can afford. The union wants a higher increase. The PERSUADER'S argumentation goal becomes to convince the union to accept the proposed increase thus abandoning the goal of higher increases. By examining the company's goal graph, depicted in Figure 2-1, one observes that since an increase in fringes contributes to an increase in economic concessions, labor costs and production costs, and a decrease in company profits, one of the company's goals is to decrease fringes. The sub-goal of decreasing employment contributes to a decrease in economic concessions, labor costs and production costs. Decreasing employment violates the goal of a union of increased employment. This can be checked by examining a union's belief structure. The argument, addressed to a union that has refused a proposed increase in fringes, "If the company is forced to grant higher increases in fringes, then it will decrease employment" is meant to make the union abandon the goal of further increases in fringes by pointing out unpleasant consequences for the union of forcing an unwanted by the company increase.

To generate the above argument, the PERSUADER finds out which company goals are violated by the union's refusal. Then, it finds out what compensating actions the company might use in retaliation (for full details on the algorithm, see [Sycara 85]). To do this, the PERSUADER

matches the fringes goal in the company's belief graph. It propagates the increase in fringes that the union wants to force to the ancestors of the fringes goal. Children of these ancestors might indicate subgoals that the company can fulfill to counteract the increase. Such a counteracting action that violates a union goal that is more important than the increase in fringes constitutes an argument that is aimed at making the union abandon the goal of further increases in fringes.

```
Importance of fringes-goall is 5 for unionl
Searching companyl goal-graph...
A increase in fringes-goall by companyl
will result in a increase in
economic-concessions1, labor-cost1,
product!on-cost!
```

```
A increase in fringes-goall by companyl
will result in a decrease in profitsl
```

```
To compensate, companyl can
decrease wage-goall, decrease employment1,
increase plant-efficiencyl, increase salesl
```

```
Only decrease wage-goall, decrease
employmentl violate goals of unionl
```

```
Importance of wage-goall is 4 for unionl
Importance of employmentl is 7 for unionl
```

```
Since importance of employmentl >
importance of fringes-goall
threat to decrease employmentl can be used
to abandon fringes-goall
One possible argument found
```

4. Concluding Remarks

In non fully cooperative multi-agent problem solving, where cooperativeness cannot be taken for granted, the agents need to be induced to cooperate. We have advocated persuasive argumentation as a general mechanism for planning how to influence agents' intentions in order to increase their cooperativeness and avoid harmful interactions. A persuader uses its model of another agent in order to generate persuasive arguments. In our work, an agent's model is represented in terms of the agent's beliefs and preferences which can be used to make inferences about the agent's behavior. We have presented a model of persuasive argumentation and strategies for argument generation. Construction of arguments is performed using integration of Case-Based Reasoning, graph search and approximate estimation of agents' utilities. The argumentation model is part of a negotiation model, implemented in the PERSUADER program, that resolves multi-agent conflicts in the labor domain.

References

- [Abelson 59] Abelson, H.
Persuasion.
Springer Publishing Company, Inc, New
York, N.Y., 1959.

- [Appelt 85] Appelt D.
Planning English Sentences.
PhD thesis, Stanford University, 1985.
- [Cammarata 83] S. Cammarata, D. Mc Arthur, R. Steeb.
Strategies of Cooperation in Distributed
Problem Solving.
Tech Report N-2031-ARPA, The Rand
Corporation, 1983.
- [Cohen 87] Cohen, P.R. and Levesque, H.J.
Intention=Choice + Commitment.
In Proceedings of AAAI-87, pages
410-415. AAAI, Seattle, WA., 1987.
- [Davis 83] Davis, R., and Smith, R.G.
Negotiation as a metaphor for distributed
problem solving.
Artificial Intelligence 20:63-100, 1983.
- [Durfee 87] Durfee, E.H.
A Unified Approach to Dynamic Coordination: Plannign Actions and Interactions in a Distributed Problem Solving Network.
PhD thesis, COINS, University of Massachusetts, 1987.
- [Flowers 82] Flowers, M., McGuire, R., and Birnbaum, L.
Adversary arguments and the logic of personal attacks.
In W.G. Lehnert and M.H. Ringle (editor), Strategies for Natural Language Processing. Lawrence Erlbaum Associates, Hillsdale, N.J., 1982.
- [Genesereth 87] Michael R. Gcnscsreth and Nils J, Nilsson.
Logical Foundations of Artificial Intelligence.
Morgan Kaufmann, 1987.
- [Halpern 84] Joseph Halpern and Yoram Moses.
Knowledge and Common Knowledge in a Distributed Environment.
IBM RJ 4421, IBM Research Laboratory, San Jose, CA 95193, October, 1984.
- [Kolodner et al. 85] Kolodner, J.L., Simpson, R.L., and Sycara-Cyranski, K.
A process model of case-based reasoning in problem solving.
In IJCAI-85, pages 284-290. IJCAI, Los Angeles, CA, 1985.
- [Konolige 86] Konolige, K.
A Deduction Model of Belief
Pitman Publishers/Morgan Kaufman Publishers, San Mateo, CA., 1986.
- [Lesser 83] Victor R. Lesser and Daniel Corkill.
The Distributed Vehicle Monitoring Testbed: A tool for investigating distributed problem solving networks.
The AI Magazine 4(3):15-33, Fall, 1983.
- [Rissland 86] Rissland E. and Ashley K.D.
Hypothetical as Heuristic Device.
In Proceedings of AAAI-86, pages 289-297. AAAI, Philadelphia, PA., 1986.
- [Rosenschein 85] Rosenschein, J.S.
Rational Interaction: Cooperation Among Intelligent Agents.
PhD thesis, Stanford University, 1985.
- [Sycara 85] Sycara-Cyranski, K.
Arguments of persuasion in labor mediation.
In IJUCAI-85, pages 294-296. IJCAI, Los Angeles, CA, 1985.
- [Sycara 87] Sycara, K.
Resolving Adversarial Conflicts: An Approach Integrating Case-Based and Analytic Methods.
PhD thesis, School of Information and Computer Science Georgia Institute of Technology, 1987.
- [Sycara 88a] Sycara, K.
Resolving Goal Conflicts via Negotiation.
In Proceedings of AAAI-88. St Paul, MN, 1988.
- [Sycara 88b] Sycara, K.
Utility Theory in Conflict Resolution.
Annals of Operations Research 12:65-84, 1988.