A GEOGRAPHICALLY DISTRIBUTED MULTI-REASONING MECHANISM FOR CHANGE NEGOTIATION MANAGEMENT OF LARGE SCALE ENGINEERING SYSTEMS

Peña-Mora F.[1] and Soibelman L.[2]

ABSTRACT: In large scale engineering projects, no design process can perfectly forecast every aspect of detailed project conditions. Sometimes oversights, mistakes or changes in the design must be corrected, and sometimes completely external factors also force changes in the situation. The basis for most changes is some form of conflict. In addition, the negotiation of such conflicts frequently results in suboptimal, inefficient agreements leading negotiators to believe that they could have reached other settlements.

This paper presents the preliminary results of a research that explores mechanisms to provide a model for this negotiation process and applies automated multi- reasoning and learning mechanisms to the same. The ultimate goal of this research is to provide negotiators with precedence information to help them reason in a co-operative way assisting in the effort to seek "all-gain" rather than "win-lose" solutions.

KEYWORDS: Artificial Intelligence, Multi-reasoning mechanisms, Computer Applications in Civil Engineering.

1. INTRODUCTION

According to Kitano and Handler[26], there are four limitations of the traditional Artificial Intelligence approach. First, most traditional models use rigid and formal knowledge representation schemes. Thus, all knowledge must be explicitly represented in order for the system to use that knowledge. Second, reasoning is generally driven by rules of principles which are abstract and which also generalize knowledge of how to manipulate specific information. Third, since the system depends on explicit knowledge and rules, the domain theory must be understood in order to build any system based on the traditional approach. Finally, knowledge and rules have been hand-coded at extensive labor cost. In many cases, coding has been carried out by experts in the field.

The fundamental assumption in the traditional approach is that experts possess the necessary knowledge regarding the problem domain, and that this expert knowledge can be explicitly written using formal representations. Most research has been carried out in such a way that researchers developed a highly intelligent system in a very restricted domain. Scholars believed that these systems could be increased in scope with larger funding and increased effort. However, experiences in expert systems, machine translation systems, and other knowledge-based systems indicate that scaling up is extremely difficult. As a consequence, poor results have been obtained. The few systems built show deficiencies such as brittleness, inflexibility, and no real time operationalism [34]. Thus, three factors prevent the application of the traditional Artificial Intelligence approach in

[1] Assistant Professor - Intelligent Engineering Systems Group - Department of Civil and Environmental Engineering - Massachusetts Institute of Technology - Cambridge - MA - USA -feniosky@mit.edu

[2] Ph.D. Candidate - Intelligent Engineering Systems Group - Department of Civil and Environmental Engineering - Massachusetts Institute of Technology - Cambridge - MA - USA -lucio@mit.edu

the real world: incompleteness - because it is almost impossible to obtain a detailed set of knowledge for a given problem domain; incorrectness - because there is no guarantee that expert knowledge is always correct, and inconsistency - because the set of knowledge may be inconsistent.

Subsequent research in Artificial Intelligence developed systems that assumed that data resources are inconsistent, incomplete, and inaccurate [10, 45, 8, 29, 46, 6, 17, 23, 18, 36, 25]. These were however, developed for very restricted domains. Furthermore, the studies lack feasibility analyses of scaling up the system.

Studies also exist on how to negotiate and manage the negotiation process [9, 3, 14, 21, 42, 5, 12, 40, 47, 20, 37, 13, 49, 28]. These however, focus only on the role of human relations without considering the process model. Existing works [41, 4, 33, 48, 27, 2, 1, 16], also fail to explore in-depth the multi-dimensional use of computers in managing this process. Whatever little exists, deals with the use of only one reasoning mechanism.

This research finds a connection among the different studies. It focuses on distributedness and parallel implementation as factors crucial to a fast and robust system. Instead of building general functional models, a novel approach [35] is adopted by developing competence models that provide expertise for particular and small task-oriented competence. Each module is responsible for doing all the representation, computation, reasoning, and execution, related to its particular competence. Furthermore, there is no central representation shared by several modules. The systems built are highly distributed. All of the competence modules operate in parallel. None of the modules is in control. Some simple arbitration method, however, is often included in order to select or fuse multiple conflicting solutions. Systems built with this framework tend to be less brittle because (1) none of the modules is more critical than the others; (2) they do not attempt to fully understand the current situation; (3) they incorporate redundant methods; and (4) they adapt over time.

The emphasis in the proposed architecture is on a more direct coupling of perception to action, distributedness and decentralisation, dynamic interaction with the environment and intrinsic mechanisms to cope with resource limitations and incomplete knowledge. Machine-learning algorithms, reasoning mechanisms, and negotiation models are being obtained from the literature and are not being developed during this research. The focus of this research is to examine the interconnection, collaboration and distribution of existing algorithms, mechanism, and models at the technical level, together with the application of this different expertise for negotiations.

The paper is organised as follows. section 2 presents a detailed statement of the work being undertaken; section 3 - the research phases; section 4 - the system prototype; and section 5 - the conclusions.

2. RESEARCH DESCRIPTION

The main research goals are: (1) The definition of a meta-representation of the negotiation problem and solution; (2) the application of a multi-reasoning distributed mechanism to solve negotiations in large scale engineering problems; (3) the reuse of past experience for negotiators support; and (4) the application of intelligent autonomous resolution agents for negotiation. Boston's Central Artery Tunnel (CAT) Project, a $7.7 billion, and decade-long project that will replace Boston's elevated Central Artery (I-93) is being used as a case study where data is being obtained and different technical solutions are being tested to evaluate the research hypothesis.

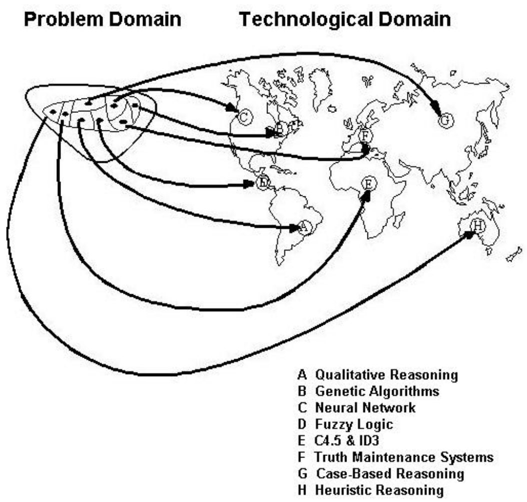

Figure 2 - Research Domain

Eight machine-learning algorithms and reasoning mechanisms are being evaluated. The list of reasoning mechanisms being evaluated are: Qualitative Reasoning, Neural Networks, Genetic Algorithms, Fuzzy Logic, Expert Systems, C4.5, Case-Base Reasoning and Truth Maintenance Systems.

The Internet is being used as a communication backbone among the different systems that implement the reasoning mechanisms being evaluated. Each different reasoning mechanism considered as an autonomous module is being developed in different servers that can be dispersed in diverse locations around the world (figure 1). Each distributed server acts as an intelligent agent (figure 2). An agent is a collection of modules each of which has its own specific competence.These modules operate autonomously and are solely responsible for the modelling, computation, reasoning, and control that are necessary to achieve their specific competence. Communication among modules is reduced to a minimum and happens on an information-low level where modules communicate to assure a consistent solution for the problem as a whole. A World Wide Web (WWW) interface receives the expected data as input and launches a manager program that broadcast the problem to all existent modules. Each module solves a part of the problem and sends back this partial solution to the manager program. This program sends the solution or possible solutions for the change negotiation problem back to the WWW interface that displays it to the user.

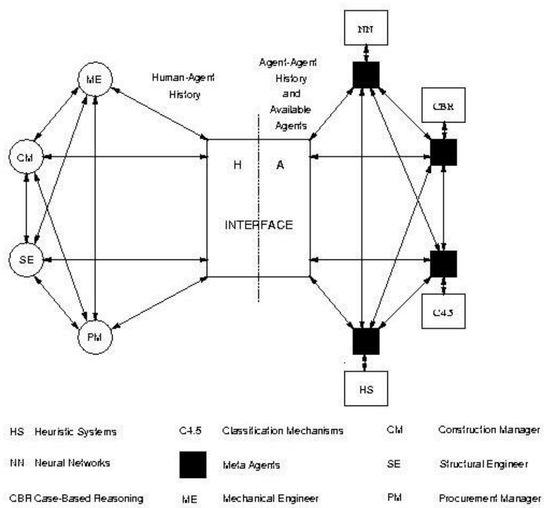

Figure 1 - System Interfaces

3. RESEARCH PHASES

The research is divided into four different phases: Conceptualization, Modelling, Mapping, and Simulation. During the Conceptualization Phase, the knowledge required to develop the research was summarized. In the Modelling phase, the model of the problem was developed and the model of the system is being elaborated. During the Mapping phase, each part of the defined problem is being mapped for the most suitable technical solution, and in the Simulation Phase, several change negotiation problems will be simulated for the case study.

3.1 Conceptualization

The Conceptualization Phase involved a bibliographical research during which knowledge about negotiation theories, machine-learning algorithms, reasoning mechanisms, support algorithms, and Knowledge Discovery in Databases methodologies were summarized.

3.1.1 Negotiation

The negotiation process has been a research topic for several decades. Negotiation is a discussion between two or more parties with a goal of reaching an agreement on issues separating the parties when neither side has the power or the desire to use its power, to get its own way. According to Mintzberg[39], negotiation is one of the four major aspects of a manager's decision-making roles. Managers can spend as much as 20 percent of their time dealing with potential and actual conflict and its resolution, primarily, through the negotiation process[39]. The actual tendency of the research in the field is toward models like “the negotiated approaches to consensus building” [47] and “principled negotiation” [13]. Both models are efforts to seek “all-gain” rather than “win-lose” solutions looking for the true interests that lie behind postures, trying to discuss these interests objectively, sharing information frankly, and siting down as fellow explorers of facts rather than, antagonists. Negotiation is usually a contest of preparation and the system proposed will allow negotiators to gather information in an innovative manner helping them simulate different possible solutions for the problems being negotiated.

3.1.2 Machine-Learning Algorithms

These tools allow raw data to be processed into significant insights, to glean important trends, to search for patterns, to describe past trends, to predict future trends, and to sift through a vast amount of information in order to extract meaning and discover new knowledge. The algorithms investigated were: Neural Networks, C4.5 and Genetic Algorithms.

Neural Network is a machine learning paradigm directly and self-consciously modelled on the form and functioning of the biological brain. The explicit goal of researchers in this field is to mimic the brain's ability to process information to learn. Since much more is unknown than known about exactly how the brain learns, the field of neural network research is admittedly, in a primitive stage of development. Neural Network can be defined as a massively parallel distributed processor which resembles the brain in two respects: 1) Knowledge is acquired by the network through a learning process; and 2) Interneuron connection strengths known as synaptic weights are used to store the knowledge. It is not clear whether the Neural Nets indeed resemble the biological brain. The point is far from being settled, as to how exactly the brain stores knowledge.

It matters little, however, for neural nets may be explored as a useful machine-learning algorithm independently of what is presently known about the functional biology of the brain. Neural Networks can be very helpful as one of the mechanisms to be employed in the development of a multi-reasoning mechanism but it is important to realise that they are not a solution to all kinds of problems. Fundamental research to understand these systems is underway. Partly because of this incomplete understanding, and due to the very nature of neural nets themselves, they cannot be used as “black boxes”. In addition, it is known that the learning phase of the implementation can be significant, possibly to the point of impeding practical applications.

In C4.5 numerous recorded classifications are examined and a model is constructed inductively by generalization from specific examples. This machine-learning algorithm can be defined as a set of computer programs that construct classification models by discovering and analyzing patterns found in given records. The algorithm that has fundamental importance in C4.5 is the one that generates the initial decision tree from a set of training cases. Since the cases do not all belong to the same class, a divide-and-conquer algorithm attempts to split them into subsets using statistical evaluations tests, and an information gain ratio criterion to find the best possible tree structure.

C4.5 may be useful in the current problem domain of change negotiation in the construction process but first a definition of the most significant attributes of the negotiation process must be found. The biggest advantage of C4.5 is its simplicity of use. The disadvantages are that, unlike Neural Network that can be used to predict real numerical values, C4.5 just allows us to obtain categorical results, and that it lacks incremental induction. Each time a new case is added to the collection of data a new classifier from all the accumulated data needs to be constructed.

Genetic Algorithms (GA) provide a robust yet efficient search methodology, explicitly modelled upon the biological “survival of the fittest" reproductive model. Computer scientists still debate the precise definition of a GA. In the broadest sense, a GA creates a set of solutions that reproduce based on their fitness in a given environment. The process follows the following pattern: (1) An initial population of random solutions is created; (2) Each member of the population is assigned a fitness value based on its evaluation against the current problem; (3) Solutions with a higher fitness value are most likely to parent new solutions during reproduction; and (4) The new solution set replaces the old, a generation is complete, and the process continues from step 2. That sequence implements, in a most simplistic way, the concept of the survival of the fittest. The outcome of a genetic algorithm is based on probabilities just as biological success is grounded in chance. The standard model for a GA solution is an anonymous bit string called chromosome after its biological counterpart. During reproduction, the chromosomes of parent solutions combine and undergo mutation in creating the next generation. It should be noted that searching for an optimum value in a domain space can be called “learning” in the sense that one searches for this value because it is unknown, and after the search, the information represented by the value, is known. The method is not an adaptive learning mechanism like Neural Network but given a criterion function it does provide the means to learn the value of a target. Thus the method has found a place among machine learning applications such as a classifier system.

3.1.3 Reasoning Mechanisms and Supporting Algorithms

These mechanisms are being applied as tools for the analysis of previous similar situations, to model expert's knowledge, and to give justification to the decisions adopted by different algorithms. The mechanisms investigated are: Case-Based Reasoning, Fuzzy Systems, Heuristic Reasoning, Qualitative Reasoning, and Truth Maintenance Systems.

Case-Based Reasoning is the process of “remember and adapt” or “remember and compare”. A reasoner remembers previous situations similar to the current one and uses them to help solve the new problem. It is a model of reasoning that incorporates problem solving, understanding and learning and integrates all with the memory processes. Learning occurs as a natural consequence of reasoning where novel procedures applied to new problems are indexed in the memory. Feedback and analysis of feedback through follow-up procedures and explanatory reasoning are necessary parts of the complete reasoning-learning circle. Without feedback, the reasoner might repeat its mistakes and never increase its capabilities.

There is much evidence that people do, in fact, use Case-Based Reasoning in their daily reasoning. People learning a new skill often refer to previous problems. Past experience can often provide guidelines on how to deal with current problems. The computer can be used as a retrieval tool to augment people's memories. References to old cases is advantageous in dealing with situations that recur. However, in applying Case-Based Reasoning to the current problem domain of change and negotiation in the construction process, it is important to keep in mind, that since the description of problems are often incomplete, the further step of understanding or interpreting the problem is a necessary prerequisite for reasoning. Furthermore, since no previous case is ever exactly the same as a new one, it is usually necessary to adapt an old solution.

The questions that then rise is that is it possible to interpret new data and adapt old solutions effectively in the construction industry with the available information? How frequently do situations recur in the problem domain? This research will evaluate the analysis of similar situations. Considering that few things facilitate a decision as much as a precedent, Case-Based Reasoning can provide objective standards in the negotiation process.

Humans reason not in terms of discrete symbols and numbers but in terms of fuzzy sets. Fuzzy sets are actually functions that map a value that might be a member of the set to a number between zero (value is not in the set) and one (value completely representative of the set) indicating its actual degree of membership. Fuzzy logic is a calculus of compatibility. Unlike probability, which is based on frequency distributions in a random population, fuzzy logic deals with describing the characteristic of properties. Fuzziness is a measure of how well an instance (value) conforms to a semantic ideal or concept. It describes the degree of membership in a fuzzy set. Knowledge acquisition is “the bottleneck problem” of Fuzzy systems applications. Fuzzy Logic is being applied to the problem domain of modelling the negotiation process: knowledge acquired from the negotiation process by applying learning algorithms like C4.5, Neural Networks, and Genetic Algorithms, added to the knowledge of the experts, is helping to model the negotiation of change by reducing its complexity and improving the handling of uncertainty and possibilities.

The Heuristic reasoning system also known as the Expert System is a computer program that represents and reasons with the knowledge of some specialist. Such a system may completely fulfill a function that normally requires human expertise, or it may play the role of an assistant to a human decision maker. An Expert System simulates human reasoning about a problem domain rather than simulating the domain itself. It performs reasoning over representations of human knowledge. In addition to doing numerical calculations or data retrieval, it solves problems by heuristic (rule of thumb which encodes a piece of knowledge) or approximate methods. Expert Systems do not require perfect data because solutions may be proposed with varying degrees of certainty. The expert must be able to perform the task, know how to perform the task, be able to explain how to perform the task, have the time to explain how to perform the task, and be motivated to cooperate in the enterprise. These systems fall short when inputs are not exactly as requested by the stored knowledge. Any deviation from the patterns they expect tends to result in a breakdown or impracticable behaviour. They fall short too in the learning process. It is also difficult to integrate information with the already existing information. Knowledge acquisition is "the bottleneck problem" of expert systems applications.

How then can we elicit knowledge from a human expert and codify it? It must be taken into consideration that specialists have their own jargon, that facts and principles in many domains cannot be characterised precisely in terms of a mathematical theory or a deterministic world, that experts need to know more than the mere facts or principles of a domain in order to solve problems, and that the human expertise is enhanced by a lot of common sense (knowledge about the everyday world). However, like Fuzzy Logic, Heuristic Reasoning can be applied to our problem domain to help model the negotiation process.

Qualitative reasoning automates the process of determining all of the possible outcomes of a deterministic system where some factors necessary to analyse the system are unknown. It is thus a simulation technique, as well as an inference mechanism. Qualitative analysis replaces the ordinary differential equations (ODE) of a deterministic system with qualitative differential equations (QDE), that inherently include uncertainty in their construction, development, and analytical process. The introduction of uncertainty means that the system is not deterministic and multiple outcomes will result from a simulation analysis. Qualitative reasoning may be useful in the current problem domain of change and negotiation in the construction process as a high-level gateway mechanism. That is, if other algorithms or systems are set up to discover the underlying mechanisms of the construction change and negotiation process, and those mechanisms structured as a qualitative reasoning framework, then modelling of a large set of indeterminate options could be accomplished by the combined system.

Truth maintenance systems (TMS) have five principal uses: (1) Identifying responsibilities for conclusions: TMS allows a problem solver to identify responsibility for its conclusions by providing rational explanations of how its conclusions follow from the premises. Generally, just providing the answer is not enough. By providing explanations, the problem solver enables the user (or itself) to figure out what to change when things go wrong. (2) Recovering from inconsistencies: In an ideal world all data would be valid and every constraint imposed would be perfectly satisfied. Neither we nor our programs live in such a world. For example, the data we feed to our program can be wrong. (3) Maintaining a cache of inferences: Most artificial intelligence solvers are a special kind of search engines. Since they search, they often go over parts of the search space again and again. If a problem solver cached its inferences, then it would not need to retrieve conclusions that it had already derived earlier in the search. (4) Guiding backtracking: When the search detects an inconsistency while exploring the solution TMS chronological backtracking, as the name implies, backtracks to the most recent choice the search has made and explores the next alternative. (5) Default reasoning: Many artificial intelligence applications require the problem solver to make conclusions based on insufficient information. TMS assumes that the solution is the generic one unless there is some evidence to the contrary. TMS will be very useful in providing justifications to the decisions adopted by the different reasoning mechanisms being applied in the multi-reasoning system being developed.

3.1.4 Knowledge Discovery in Databases

The overall process of finding and interpreting patterns from data is referred to as the Knowledge Discovery in Databases Process (KDD process)It is typically interactive and iterative, involving the repeated application of specific data mining methods or algorithms and the interpretation of the patterns generated by the algorithms. The KDD process is a multi-step process, that involves data preparation, search for patterns, knowledge evaluation, and refinement involving iteration after modification. KDD is, thus the non-trivial process of identifying valid, novel, potentially useful, and ultimately understandable patterns in data.

The information available from the KDD and Data Mining research is proving to be very helpful for the development of the multi-reasoning system. It is important to realize that the majority of the KDD systems are supported by one data mining tool without considering the possible advantages of the interaction between different tools. It is equally important to know that their main objective is data description that focuses on finding human interpretable patterns describing the data. This objective is different from the main objective of the multi-reasoning mechanism that is, data prediction that involves using some variables or fields in the database to predict unknown or future values of other variables of interest.

3.2 Modelling

The Modelling Phase seeks to develop the meta-representation of negotiations as well as the model of communications and co-operation between the multiple reasoning mechanisms. This phase is divided into two stages both using the CAT project as a case to obtain data. In the first stage, the problem domain is being defined. Research techniques like interviews, questionnaires, observation, and analysis of available change orders are being applied to help define the problem. In the second stage, a model is being developed so that Machine-learning algorithms like Neural Networks[22, 44], C4.5[43], Genetic Algorithms[19, 32], and Case-based Reasoning [30], are processing raw data available in change orders.

The most powerful way known for discovering how to solve a hard problem is to find a method that splits it into several simpler ones, each of which can be solved separately. The division of the problem into its smallest unit parts has the objective of imitating the way the human brain deals with complex problems. According to Minsky[38], each mind is made of many smaller processes. These, he calls agents. Each mental agent by itself can only do some simple thing that needs no mind or thought at all. Yet when those agents are joined in societies - in certain very special ways - this leads to true intelligence. The human brain when dealing with a complex problem uses different reasoning mechanisms to solve each part of the divided problem. No reasoning mechanism by itself is powerful enough to solve all different kinds of problems that one may encounter in a negotiation or conflict.

Thus, in the Modelling Phase, this research will be concerned with finding methods that computers can use for splitting the change negotiation problem into smaller sub-problems and then solve them by mapping to the appropriate reasoning mechanisms

3.3 Mapping

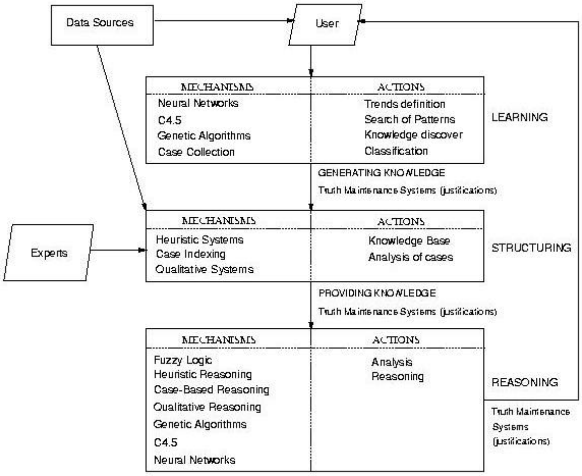

Figure 3 - Change Negotiation system

During the Mapping Phase each small part of the defined problem is being mapped to the most suitable technical solution. Figure 3 presents an architecture for using machine-learning algorithms and reasoning mechanisms to help humans perform negotiations.

After the modelling phase, the knowledge generated by the implementation of Neural Networks, C4.5, Genetic Algorithms, and Case-based reasoning together with the knowledge provided by the justifications from the Truth Maintenance Systems [15] and the knowledge provided by experts will allow the system to create the required rules to apply Heuristic Systems[24], Fuzzy Systems[7], Qualitative Systems[31], and to index the available cases.

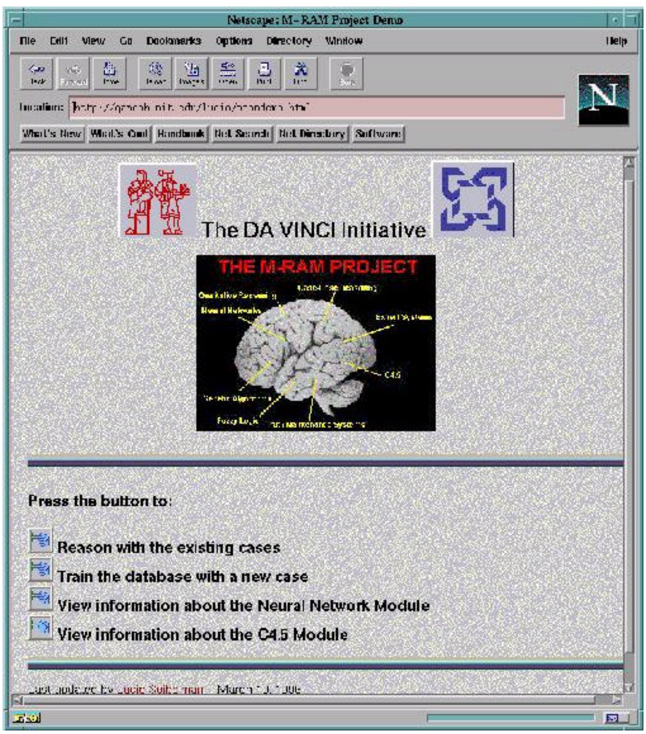

Figure 4 - System Demo

3.4 Simulation

In the Simulation Phase, several change negotiation problems will be simulated for the case project. The WWW interface being developed will allow with the proposed distributed and parallel architecture; - distributed because each different reasoning will be located in a different server and parallel because all mechanisms will be able to reason at the same time.

4. THE PROTOTYPE

A simplified prototype was developed to investigate the feasibility of the proposed system architecture. The prototype is composed of two reasoning modules: a C4.5 module and a Neural Network Module. The

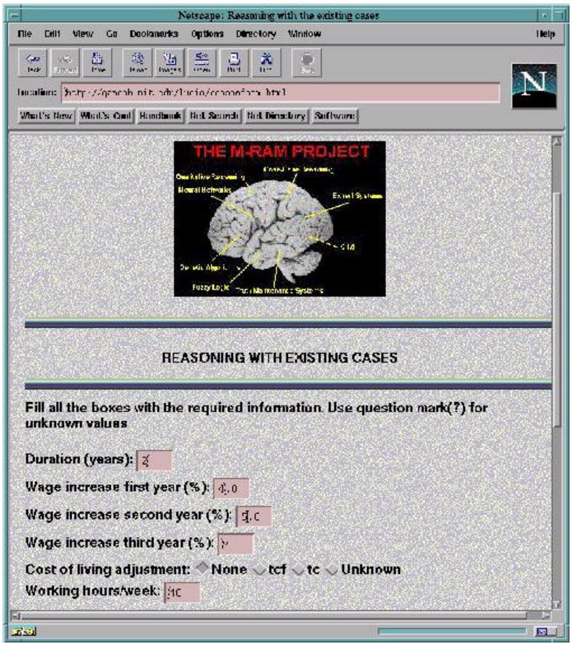

Figure 5- Form to Reason with Existing Database

algorithms for C4.5 and neural network were taken from[43], and from [44]. Figure 4 presents the first Web page of the prototype demo. In this page the user has four options: (1) to reason with the existing database, (2) to add new cases to the database, (3) to view information from the Neural Network Module (e.g. the network weights and a file with the neural network training data), and (4) to view information from the C4.5 module (e.g. the decision tree, the set of production rules and a file with the C4.5 training data).

The database used in this example came from [43]. This database represents labour salary negotiations and was used since it was readily available and allows the testing of the research hypothesis on a rapid prototype. The next step is to use a database with negotiations from the CAT Project.

If the user decides to reason with the existing database, he or she just needs to fill a form like the one presented in figure 5. When the form is submitted, a CGI script that communicates with the Neural Network meta-agent and the C4.5 meta-agent is automatically launched. These meta-agents translate the user inputs from the form format to a format understandable by the neural network and C4.5 programs and request a solution to the problem. Both modules send back to the user interface their recommended solution to the problem (figure 6).

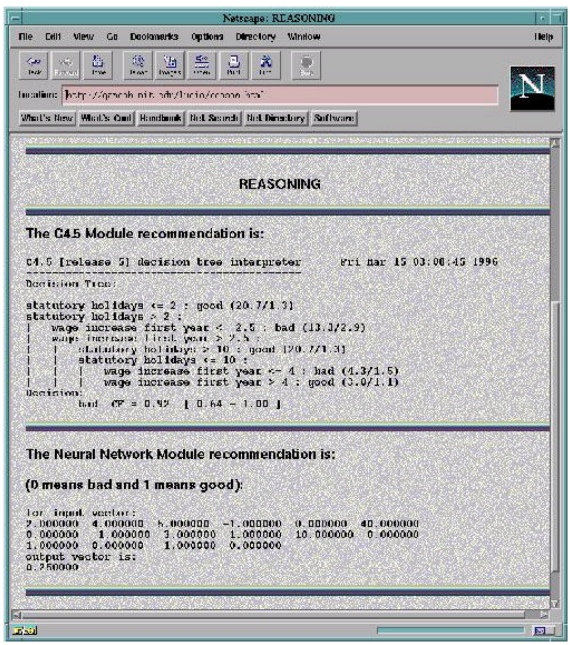

Figure 6 System Recommendations

If the user decides to add a new case to the database, he or she just needs to fill another form. When the form is submitted a CGI script is launched and it communicates with the modules meta-agents. These meta-agents translate the information available in the form to a format understandable by the modules, save that information in each module's training file, recalculate the decision tree, and the set of production rules in the C4.5 module or the new network weights in the Neural Network module(figure7). To try

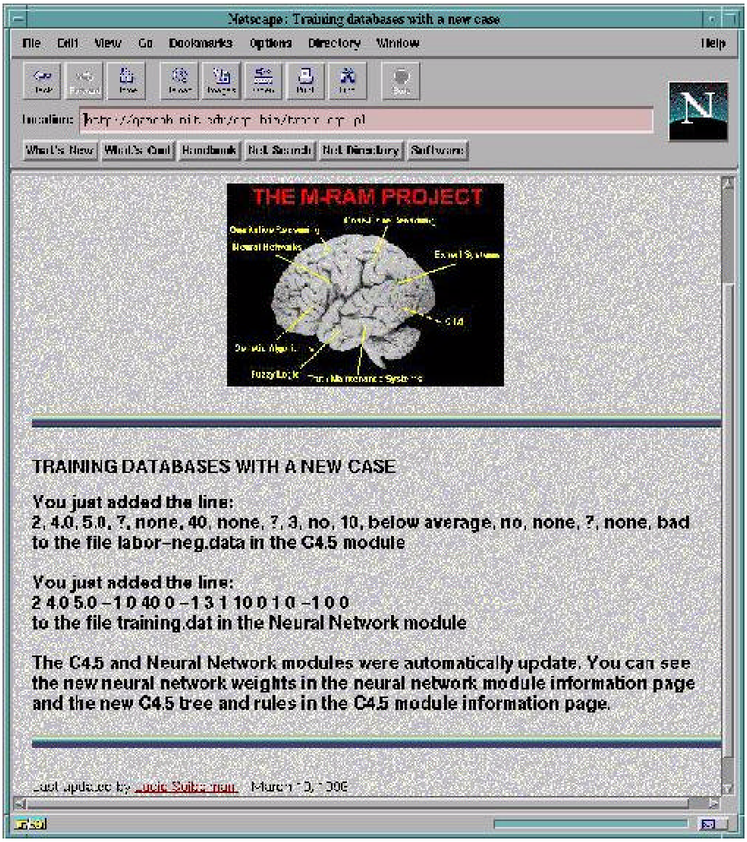

Figure 7 - Adding a New Case

the example, please see the URL: http://ganesh.mit.edu/lucio/mramdemo.html

5. CONCLUSIONS

The attempt of this research in interconnecting and distributing existing machine-learning algorithms, reasoning mechanisms, and negotiation models at the technical level, and on the exploration of the different expertise for negotiations will enable the future development of better methodologies for machine collaboration and a better understanding of the negotiation process. The results presented in this paper are preliminary but has shown that the research is feasible. The next steps of the research are: (1) implementation of other machine learning and reasoning mechanisms, (2) generation of the database for negotiations from the CAT Project change orders, and (3) development of better communication protocols between different reasoning mechanisms allowing them to communicate and to integrate their partial solutions.

6. REFERENCES

[1] Anandalingam, G. and Apprey, V., *"Multi-Level Programming and Conflict Resolution in International River Management"*, Center for Research in Conflict and Negotiation, Penn State University, 1992.

[2] Anson, R. and Jelassi, M., "A Development Framework for Computer-Supported Conflict Resolution", *European Journal of Operational Research*, pp. 181--199, 1990.

[3] Axelrod, R., *"The Evolution of Co-operation"*, Basic Books, 1984.

[4] Brown, D., "Failure Handling in a Design Expert System", *Computer-Aided Design*, pp.436-442, North-Holland, 1985.

[5] Byrnes, J., "Ten Guidelines for Effective Negotiating", *Business Horizons*, May-June, pp. 7-12, 1987.

[6] Cook, J., "The Base Selection Task in Analogical Planning", *Proc. of AAAI-86*, 1986.

[7] Cox, E., *"The Fuzzy Systems Handbook - A Practitioners Guide to Building, Using, and Maintaining Fuzzy Systems"*, AP Professional, 1994.

[8] Dixon, M. and de Kleer, J., "Massively Parallel Assumption-based Truth Maintenance", *Proceedings of the National Conference on Artificial Intelligence (AAAI-88)*, 1988.

[9] Doyle, M., Straus, D., *"How to Make Meetings Work"*, Jove Books, 1982.

[10] Drumheller, M., "Connection Machine Stereomatching", Proceedings of the National Conference on Artificial Intelligence (AAAI-86), 1986.

[11] Evett, M., Hendler, J., Mahanti, A, and Nau, D., "PRA A memory-limited heuristic search procedure for the connection machine", *Proceedings - Frontiers of Massively Parallel Computing*, 1990.

[12] Fisher, R., Davis, W.H., "Six Basic Interpersonal Skills for a Negotiator's Repertoire", *Negotiation Journal*, April, pp. 117-122, 1987.

[13] Fisher, R., Ury, W., Patton, B., "Getting to Yes - Negotiating Agreement Without Giving In", Penguin Books, 1991.

[14] Folberg, J., Taylor, A., "The Nature of Conflict, Lewicki, R.J., Litterer, J.A., Minton, J.W.,(Ed.)", *Negotiation Readings and Cases*, IRWIN, 1993.

[15] Forbus, K.D., de Kleer, J., *"Building Problem Solvers"*, The Mit Press, 1993.

[16] Fraser, N. and Hipel, K., "Using the Decision Maker Computer Program for Analysing Environmental Conflicts", *Journal of Environmental Management*, pp. 213-228, 1988.

[17] Geller, J., *"Advanced Update Operations in Massively Parallel Knowledge Representation"*, New Jersey Institute of Technology, CIS-91-28, 1991.

[18] Gero, J.S. and Qian, L., "A Design System Using Analogy", *Artificial Intelligence in Design '92*, Kluwer, 1992.

[19] Goldberg, D.E., "Genetic Algorithms in Search, Optimisation, and Machine Learning", Addison Wesley, 1989.

[20] Gray, B., "Collaborating: Finding Common Ground for Multiparty Problems", Jossey-Bass, 1989.

[21] Greenhalgh, L., "Managing Conflict", *Sloan Management Review*, Summer, pp.45-51, 1986.

[22] Haykin, S., *"Neural Networks - A Comprehensive Foundation"*, Macmillan College Publishing Company, 1994.

[23] Huhns, M.N. and Bridgeland, D.M., "Multiagent Truth Maintenance", *IEEE Transactions on Systems. Man. and Cybernetics*, Vol.21,No.6. November-December, 1991.

[24] Jackson, P., *"Introduction to Expert Systems"*, Addison Wesley, 1990.

[25] Kettler, B., Andersen, W., Hnedler, J., and Evett, M., "Massively Parallel Support for Case based Planning System", *Proceedings of the Ninth IEEE Conference on AI Applications*, Orlando, Florida, March 1993.

[26] Kitano, A. and Hendler, J., *"Massive Parallel Artificial Intelligence"*, AAAI Press/The MIT Press, 1994.
[27] Klein, M., "DRCS: an Integrated System for Capture of Designs and their Rationale", *Artificial Intelligence in Design '92*, pp. 393-412, 1992.
[28] Kraus, S. and Lehmann, D., "Designing and Building a Negotiating Automated Agent", *Computational Intelligence*, 11(1):132-171, 1995.
[29] Kolodner, J., "Retrieving Events from a Case Memory: A Parallel Implementation", *Proceedings of Case-Based Reasoning Workshop*, 1988.
[30] Kolodner, J., *"Case-Based Reasoning"*, Morgan Kaufmann, 1993.
[31] Kuipers, B., "Qualitative Reasoning - Modelling and Simulation with Incomplete Knowledge", The MIT Press, 1994.
[32] Ladd, S.R., *"Genetic Algorithms in C++"*, MIS Press, 1996.
[33] Lander, S. and Lesser, V., "A Framework for the Integration of Co-operative Knowledge-Based Systems.", *Proceedings of IEEE International Symposium on Intelligent Control*, pp. 472--477, 1989.
[34] Maes, P., "Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back", The MIT Press, 1991.
[35] Maes, P., "Behavior-Based Artificial Intelligence", *Proceedings of the Second Conference on Adaptive Behavior*, MIT Press, 1993.
[36] Maher, M.L. and Zhao, F., "Using Network-Based Prototypes to Support Creative Design by Analogy and Mutation", *Artificial Intelligence in Design'92*, Kluwer, 1992.
[37] Michael, B. and Wilkenfeld, J., *"Crisis, Conflict, and Instability"*, Oxford, 1989.
[38] Minsky, M., *"The Society of Mind"*, Touchstone, 1988.
[39] Mintzberg, H., *"The Nature of Managerial Work"*, Harper and Row, 1973.
[40] Morgan, P.I., *"Resolving Conflict trough "Win-Win" Negotiating"*, Management Solutions, August, 1987.
[41] Peña-Mora, F., Sriram, R.D., Logcher, R., "Conflict Mitigation System for Collaborative Engineering" *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, Cambridge University Press, 1995.
[42] Pruitt, D., Rubin, J., *"Social Conflict"*, McGraw-Hill, 1986.
[43] Quinlan, J.R., *"C.4.5 Programs for Machine Learning"*, Morgan Kaufmann, 1993.
[44] Rao, V. and Rao, H., *"C++ Neural Networks and Fuzzy Logic"*, MIS Press, 1995.
[45] Robertson, G., "Parallel Implementation of Genetic Algorithms in a Classifier System", *Davis, L.,(Ed.) Genetic Algorithms and Simulated Annealing*, Morgan Kaufmann Publishers, 1987.
[46] Stanfill, C., Thau, R., and Waltz, D., "A Parallel Indexed Algorithm for Information Retrieval". *Proceedings of the SIG-IR-89*, 1989.
[47] Sussking, L., Cruikshank, J., *"Breaking the Impasse"*, Basic Books, 1987.
[48] Sycara, K., "Co-operative Negotiation in Concurrent Engineering Design", *Proceedings of MIT-JSME Workshop in Computer-Aided Co-operative Product Development*, pp. 269-297, 1989.
[49] Wilkenfeld, J. and Kraus, S., "A Strategic Negotiation Model with Applications to an International Crisis", *IEEE Transactions on Systems, Man. and Cybernetics*, Vol.23, No.01, January-February, 1993.