Reports

Machine Learning and Inference Laboratory

Distributed Decision Support
in Disruptive Environments

Janusz Wojtusiak and Otthein Herzog

MLI 10-1

December, 2010

College of Health and Human Services

George Mason University

DISTRIBUTED DECISION SUPPORT IN DISRUPTIVE ENVIRONMENTS

Janusz Wojtusiak[1,2] and Otthein Herzog[1,3]

(1) Machine Learning and Inference Laboratory
George Mason University, MSN 1J3
Fairfax, VA 22030-4444, USA

(2) Department of Health Administration and Policy
George Mason University

(3) TZI - Universitaet Bremen
Am Fallturm 1
28359 Bremen, Germany

Abstract

New class of methods is needed to support decision makers in complex and dynamically changing environments. Agent-based simulation combined with human-oriented machine learning provides robustness and efficiency, and at the same time ability to explain reasons for selected plans to decision makers. This makes the method ideal for decision support in healthcare, medicine, security, disaster planning and recovery, and other critical areas.

Keywords: agent-based modeling, decision support, disruptive environment, machine learning

Acknowledgments

Part of this research has been funded by the German Research Foundation (DFG) within the Collaborative Research Centre 637 “Autonomous Cooperating Logistic Processes – A Paradigm Shift and its Limitations” at Universitaet Bremen, Germany.

1 INTRODUCTION

Decision makers need to make decisions in complex situations and constantly changing environments. They are also usually under strong time pressure. This is particularly the case when disruptions to plans occur, for example during emergencies such as natural disasters, terrorist attacks, as well as mechanical interruptions in production or other events that affect scheduling. To aid decision makers computer-based decision support tools can be used. One specifically important approach to decision support that can be used in advanced planning is based on simulation and modeling.

Simulation and modeling methods are used in decision making in three main situations:

Training: Decision makers can be trained to quickly respond in critical situations by engagement in training that involves multiple realistic scenarios.

Preparation: To deal with unknown future events, preparation is needed. Decision makers may want to know how to respond and what will be the consequences of something happening (i.e. a disaster or other interruption in normal functioning of an organization).

Planning: In order to prepare for anticipated future events decision makers can use simulation methods to plan for the best causes of action.

Real time decision making: Real time decision making is needed to make fast informed decisions to update existing plans whenever needed.

2 DECISION SUPPORT THROUGH SIMULATION AND MODELING

To make informed decisions in complex situations, decision makers need to be provided with all relevant information about what is currently happening, predictions about the future, and available expertise (knowledge that can be applied to support decision making in a specific situation). All three of these elements are interrelated (Figure 1). The current situation needs to be known to decision makers, but also is needed to make predictions about the future. Traditional predictions about the future are made using *static* models that consider only the current situation. It is crucial, however, to provide decision makers with *dynamic* predictions that not only account for the current situation, but also a proposed cause of action. Furthermore, automated methods can be used to discover the most optimal plans/causes of action.

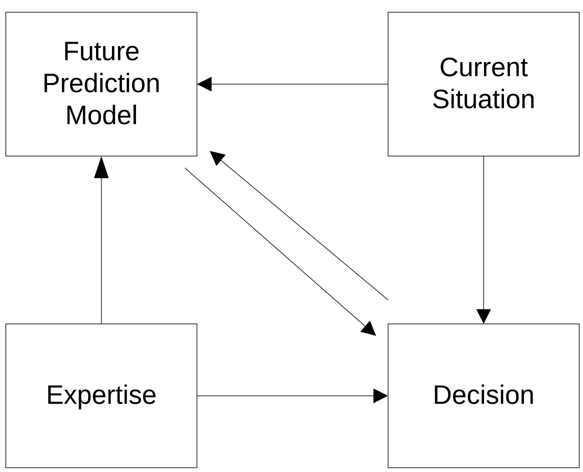

Figure 1: Decisions depend on the current situation, predictions of the future, and available expertise.

Simulation models provide decision makers with the ability to perform "*what-if*" analyses, to see how their decisions affect outcomes. Additionally some simulation methods, such as agent-based, allow decision makers and analysts to observe step-by-step how their actions affect the environment being simulated. Multiple simulations can be executed to test different scenarios, and to measure variation in outcomes due to nondeterministic nature of simulations.

When the number of possible decisions is large, or the decisions consist of multi-step pans, it is not possible to manually test all possibilities. Further, for large planning problems that involve synchronization of thousands of objects and constraints, even large computers will require many hours to prepare plans. When the plan is disrupted, re-planning needs to be performed very quickly to avoid delays. This is presents din the next section that describes centralized approach to planning in which one global plan is created, and contrasts it with distributed approach in which planning is delegated to local entities responsible for optimally performing their tasks.

3 DISTRIBUTED AND CENTRALIZED APPROACHES TO SIMULATION AND PLANNING

Planning has a long standing scientific tradition in Artificial Intelligence (AI). On the other hand industry efforts are concentrated very successfully on planning for large tasks at hand, like, e.g., manufacturing processes, route planning, and process planning in general. In contrast, in AI there was a concentration on the principles of planning problems, i.e., given an initial state of the planning world, how to build a plan to achieve a given goal? It's not very surprising that the latter approach delivered a host of theoretical results and very fast planning algorithms, but it turned out that it was quite difficult to use these algorithms under real-world conditions which usually are much less deterministic than assumed. Thus some extensions to the "classical planning problem" have been proposed (cf. de Weerdt et al. 2005) such as the handling of time (Do and Kambhampati, 2001, Penberthy and Weld, 1992, Smith and Weld, 1999), utility maximization (Haddawy and Hanks, 1998), planning with limited resources (Köhler, 1998; Wolfman and Weld, 2001), and planning under uncertainty (Boutilier et al., 1999). These propositions turned out to exhibit the computational complexity of EXPTIME, rendering the planning of real-life processes to be impossible because they would require resources, which are far beyond the possibilities of everyday computing (de Weerdt et al. 2005). In addition, uncertainties effects like probabilities or partial observabilities and contribute to the inherent complexity of these algorithms caused by disruptive effects in real-world applications, For example in the logistics transportation domain, an important part of integrated supply networks, there is always a probability that a set of deliveries will not arrive in time because of traffic conditions.

Even with these (maybe even discouraging) theoretical results it is usually possible to use planning algorithms, which are useful to real-life. By providing additional knowledge about the application domain it is often possible to arrive at tractable algorithms, as well as through subdividing the planning problem. There are even domains, e.g., the logistic domain, where it is natural to assign simple tasks to agents as part of a multi-agent system, especially where the task at hand is by definition a subtask of a general framework. Although the agents in a multi-agent network will still need to communicate in order to prepare optimal or near-optimal plans, this plan communication can considerably reduce the effects of uncertainty in a multi-agent planning network (cf. Balakirsky and Herzog 2004). The task is then to derive local plans and coordinate them (maybe, e.g., because of locality reasons) even with only a limited number of other agents. This approach can reduce complexity, can deal with the dynamics of processes, can achieve an additional degree of robustness against disruptions, and can handle the natural distribution of tasks (Timm et al. 2001; Toenshoff et al. 2002; Stefan Kirn et al. 2006; Lorenzen et al. 2006; Bemeleit et al. 2007; Langer et al. 2006; 2007; Gehrke et al. 2008; Schuldt 2009; Gehrke et al. 2010).

The essence of planning for real-time processes is decision-making, which requires the anticipation of some strands of the future. Weerdt et al. (2005) argues that planning alone is not sufficient, but that planning must include coordination (in a multi-agent system), this idea is extended here: as it is possible to use multi-agent systems (basically with an additional synchronization algorithm among the agents) for simulation purposes (Becker et al. 2006). It is also possible to extend planning to

Planning = Planning + Coordination + Simulation

This will add an additional validation layer to the planning framework thus allowing the tailoring of communication needs and to identify any bottlenecks early on. This method shows (Gehrke et al., 2008) that multi-agent systems can be used in the real world as well as the amount of cost savings by modeling real-world processes (Jedermann et al., 2006; Schuldt, 2010). Furthermore, this approach with multi-agent systems turns out to be especially valuable as the multi-agent system used in the simulation can be transferred easily into a production system by removing the simulation synchronization.

In logistics it became clear very early on that the resulting dynamic and structural complexity of logistics networks renders the conventional central planning and control structures as outdated. This results in difficultly in providing all necessary information for central planning and control instance in a timely manner during the planning phase as well as reactions to incoming information during the execution phase. A possible solution to these challenges is the development of autonomous logistic processes, which have the ability and capabilities for decentralized coordination and decision making, i.e., autonomous logistic entities, which can be implemented by multi-agent systems [Bemeleit et al., 2007]. These autonomous entities are self-contained and follow their local goals. In the basic assumption they are individually rational decision makers in the sense of game theory, each aiming at maximizing their individual utility function.

4 ADAPTABILITY IN DISTRIBUTED ENVIRONMENTS

Continuous and efficient operation in real complex environments requires the adaptation of agents performing tasks and supporting decision making. What is mainly understood by adaptation is the agents' ability to operate in a changed environment, while also including increases of performance in their tasks. In a new environment, where it is no longer possible to repeat tasks as they were done before, agents need to perceive the current situation and modify their usual actions to reflect necessary changes. The previous section indicated that it is more efficient to deal with disruptions to normal operation in distributed system in which autonomous agent make local decisions, rather than in a centralized system in which each disruption requires global re-planning which is often not feasible.

The second meaning of adaptability concerns agents automatically improving their performance in completing their tasks. This is traditionally done as a part of reinforcement learning in which agents are learning how to perform tasks in a way that maximizes a long term reward. Agents representing complex logistic objects perform several tasks leading to the final planning.

Intelligent agents often need to reason with multiple models, and upon their experience update these models. In the proposed research we will consider five immediate reasons for this, driven by the application area in the autonomous logistics:

1: Exploratory and long-term learning. Exploratory learning is based on a single agent's limited short term experience. Models created in the process of exploratory learning may apply only to specific situations, or be simply incorrect because of limited observations. On the other hand, exploratory learning allows for the capturing of unusual situations, it allows individual agents to use different learning methods, and therefore agents are able to explore possibilities not available when agents used the same knowledge.

Long-term learning is used to create knowledge of validated, typical, cases that are not updated each time new agents' experience is gathered (in contrast to exploratory learning based on every experience of an agent). In the proposed study, long-term learning is the process of converting validated exploratory knowledge/models into a more permanent form. It is the process of learning from already existing exploratory knowledge, not from data as typically considered in machine learning. In particular, we consider the process of long-term learning as the integrating of models created through exploratory learning.

Our claim, to be tested in the proposed research in the domain of logistics, is that by properly combining exploratory and long-term learning, the overall multiagent system is more robust, and able to adapt faster to changing environments.

2: Distribution of models. While exploratory models learned are created by individual agents only, long-term models can be either distributed or centralized. At the local level, models are created and available only to individual agents. This corresponds to exploratory models, to models that apply only to a specific agent (because of its properties that distinguishes it from other agents), or models that an agent does not want to share with other agents. On the other end of spectrum, there are global models which are available to all agents. These types of models come from publically available sources, and include, for example, calendars of national holidays, traffic prediction models (some of which are freely available in the internet), road network models, etc. Because of the public character of global knowledge, it may not be easily updated, as it is collected by different agencies, companies and groups. In the proposed research, we will use global models only for the purpose of making predictions, and will not consider the problems of learning or updating them.

At the most important level of distribution, are models that are between those that are local and global. These are long-term models shared by groups of agents (usually not all), and are created from local, exploratory models of contributing agents. The models may be available to some, but not all agents, e.g. agents representing trucks or containers within one company. The process of learning these shared models is based on integration of models created by individual agents (section 5).

3: Representation. Each agent may use a different type of model to make its predictions (e.g. rules, decision trees, Bayesian networks, regression). Even if models used by two agents are of the same type, they may represent data differently (different attributes, different discretizations of numeric attributes, etc.). Also, highly accurate, statistical models can be used to represent long-term models. Such models will be encountered for most typical, validated cases. On the other hand, to handle very few cases gathered by individual agents, in learning their local, exploratory, models a symbolic representation such as a rule-based representation, is more desirable. Additionally, knowledge provided by experts needs to be considered. This includes ontologies of concepts used in logistics, and business rules that enforce agents to comply with company strategies.

4: Different types of predictions. An autonomous logistics agent needs to make complex decisions. Such decisions may require making multiple predictions, and combining them together. For example, a transportation agent associated with a truck, when making its decision if to accept cargo, may use predictions of time needed to go to destination, the cargo available at the destination (to minimize travel without load), and the number of other trucks that will compete for that cargo. Similarly, a cargo agent associated with a container may make predictions about the approximate travel time, trucks available at intermediate destinations, cost of hiring trucks from the intermediate destinations, and others. Because of clearly different purposes, these predictions need to be made by different models.

5: Perspective. Different types of agents may create models differently for making their predictions. Two types of agents mentioned earlier are transportation agent, corresponding to transportation media such as trucks, ships, airplanes or trains, and cargo agent, corresponding to the shipped goods, usually represented by a container, but that can also represent smaller entities such as individual packages. There are additional types of agents that correspond to other entities in logistics systems, i.e. mediation agents and information agents. These agents build their models based on different data available to them and for different purpose.

To illustrate some of the above points, consider an agent depicted in Figure 2, which represents a truck making a decision about its possible route. In this very simplified model, the agent uses prediction about traffic, prediction about the status of different destinations (e.g. how much cargo will be available there at the time of arrival), and a road network information. Global, company, and local traffic prediction models are integrated (down arrows) to form an ad hoc model in context of a specific situation faced by the agent. The model is then used to make the prediction, and feedback (dashed arrows) is provided to the local and company model, which constitutes part of the models' learning. The novelty of the proposed approach is that the feedback not only includes the prediction and its result, but more importantly, it includes the entire ad hoc model (or its relevant parts) that can be integrated within the original models. The feedback is not provided to the global model which cannot be updated. Similarly, two company models and one local model for predicting the number of available cargo at destination are integrated into an ad hoc model to make prediction and provide feedback to the original models.

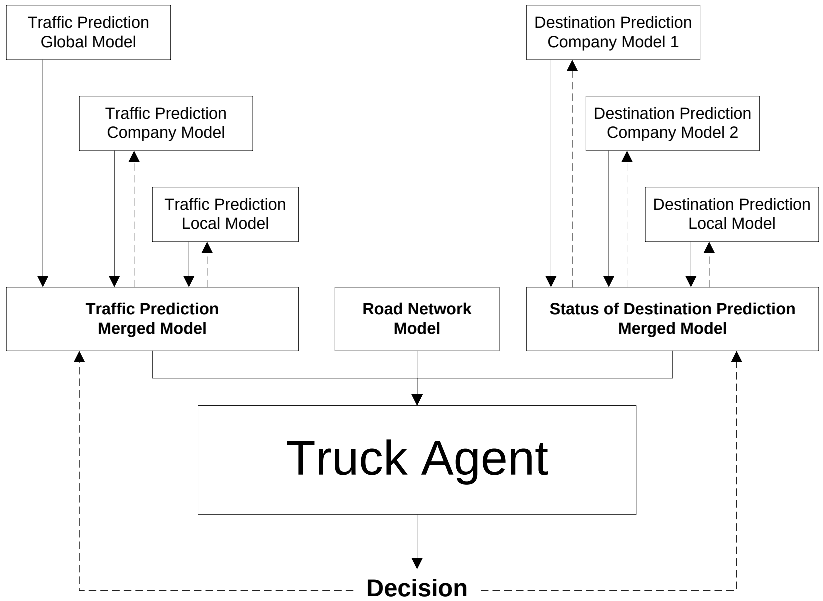

Figure 2: A simplified model of a transportation agent making decisions.

While the goal of the presented research is not investigating the actual decision making in logistics (it is extensively studied by our colleagues at the University of Bremen), it is important to understand the process. Our focus is on the way models of a specific type are integrated to create an intermediate model that is used to give the final prediction. In cases when models contributing to the intermediate model do not agree on their prediction, the model merging process enters into a negotiation phase in which models “explain” their decisions to other model, and try to arrive at a common solution. The challenge of this approach is to make models constructed in different representations, e.g. Bayesian network, a set of rules and a logistic regression model, understand each other. The feedback and solutions from the “merged” models are then propagated to the original models as a part of their learning process.

5 APPLICATION AREAS

High robustness and adaptability of the presented agent-based methodology makes it particularly applicable to important domains. These include disaster management in which agents may represent resources to be moved to the disaster site, people to be evacuated and transportation means; logistics in which agents may represent transported goods and means of their transportation; and healthcare in which agents may represent patients, supplies, hospital beds, operating rooms, and personnel.

The advantage of the described method over more traditional centralized approach is the most visible in application areas that are decentralized, dynamic, require simulated, and critical. A specific application area does not need, however, to have all of these features to deploy the methodology.

Decentralized: Geographical distribution of entities makes communication and planning harder. This is particularly case when disruptions and emergencies happen and normal communication channels may not be functioning.

Dynamic: Environments that change over time require information systems to adapt to the changes. It is often not possible to manually detect and follow all the changes in the environment. Instead the changes can be automatically discovered and incorporated in the systems' operations.

Simulated: Application areas that require performing "*what-if*" analyses can benefit from simulation capabilities of the proposed methodology. In fact, the same platform can be used for real-time management of resources, and for simulated scenarios that provide decision makers with potential outcomes of strategic decisions.

Critical: The presented methodology provides robustness while preserving transparency of systems' operation thus leaving an audit trail and explanations for all internal decisions.

6 CONCLUSION

Distributed decision making is an effective solution, particularly in dynamically changing environments in which normal processes can be disrupted by unexpected events. An attractive solution to distributed systems is thorough multi-agent systems, in which each individual agent acts in environment to solve its tasks. The agents can then collaborate with each other or act independently.

One of the most important characteristics of agents is their ability to learn and adapt to changes in the environment. Individual agents can more quickly adjust their behavior in local context in which they operate than it takes to change entire global system, particularly when disruptions are local. The learning process can be done in centralized or distributed manner, be used to explore possibilities or create long-term knowledge, affect different types of agents' behavior, use different knowledge representations, and be done from sometimes conflicting perspective of different agents.

Among the most important application areas are these in which critical decision making needs to be done in timely, real time or near real time, fashion. These application areas include security, healthcare, and logistics.

REFERENCES

Balakirsky, S. and Herzog, O., “Parallel Planning In Partitioned Problem Spaces*”. In: Proc. of the 5th IFAC/EURON Symposium on Intelligent Autonomous Vehicles*, Instituto Superior Técnico, Lisbon, Portugal, p. 328-333. July 5 - 7, 2004.

Becker, M., Gehrke, J.D., Wenning, B.L., Lorenz, M., Görg, C., and Herzog, O., “Agent-based and Discrete Event Simulation of Autonomous Logistic Processes”. *In: W. Borutzky and A. Orsoni and R. Zobel (Eds.): Proceedings of the 20th European Conference on Modeling and Simulation,* Bonn, Sankt Augustin, Germany, pp. 566-571. 2006.

Bemeleit, B., Lorenz, M., Schumacher, J., and Herzog, O., “Autonomous Logistic Processes: An Application Oriented Approach for Risk Management”. *In Proc. The 12th International Symposium on Logistics (12th ISL),* Budapest. July 8-10, 2007

Bemeleit, B., Lorenz, M., Schumacher, J., and Herzog, O., “Proactive Knowledge-based Risk Management”. *In: Hülsmann, M. and Windt, K. [Eds.]: Understanding Autonomous Cooperation & Control in Logistics*.Springer, Berlin, pp. 233–254. 2007.

Boutilier, C., Dean, T.L., and Hanks, S., “Decision-theoretic planning: Structural assumptions and computational leverage”. *Journal of AI Research*, 11, pp.1–94. (1999).

Do, M. and Kambhampati, S.S., “A domain-independent heuristic metric temporal planner”. *In: Proceedings of the Sixth European Conference on Planning (ECP-01)*, pp. 109–120. (2001).

Gehrke, J.D., Langer, H., and Herzog, O., “Distributed Control for Robust Autonomous Logistic Processes”. *In: Hans-Christian Pfohl, Thomas Wimmer (Eds.): Robuste und sichere Logistiksysteme*. Verkehrs-Verlag, Hamburg. DVV Media Group-Deutscher, pp. 25-37. 2008.

Gehrke, J.D., Herzog, O., Langer, H., Malaka, R., Porzel, R, and Warden, T., “An Agent-Based Approach to Autonomous Logistic Processes”. *In: Künstliche Intelligenz*. February 2010.

Haddawy, P. and Hanks, S., “Utility models for goal-directed, decision-theoretic planners”. *Computational Intelligence*, 14:3, pp. 392–429. 1998.

Jedermann, R., Gehrke, J. D., Lorenz, M., Herzog, O., and Lang, W., “Realisierung lokaler Selbststeuerung in Echtzeit: Der Übergang zum intelligenten Container.” *In: Pfohl, H. C.; Wimmer, T. (Hrsg.): Wissenschaft und Praxis im Dialog. Steuerung von Logistiksystemen - auf dem Weg zur Selbststeuerung, Deutscher Verkehrs-Verlag*, Hamburg, pp. 145-166. 2006

Kirn, S., Herzog, O., Lockemann, P., and Spaniol, O. Multiagent “Engineering - Theory and Applications in Enterprises”. Springer Verlag, 2006.

Koehler, J., “Planning under resource constraints”. *In: Proceedings of the Thirteenth European Conference on Artificial Intelligence (ECAI-98)*, John Wiley & Sons, pp. 489–493. 1998.

Langer, H., Gehrke, J.D., Hammer, J., Lorenz, M., Timm, I.J., and Herzog, O., “A Framework for Distributed Knowledge Management in Autonomous Logistic Processes”. *In: International Journal of Knowledge-Based & Intelligent Engineering Systems*, 10:4, pp. 277-290. 2006.

Langer, H., Gehrke, J.D., and Herzog, O., “Distributed Knowledge Management in Dynamic Environments”. *In: Hülsmann, M.; Windt, K. (eds.): Understanding Autonomous Cooperation in Logistics - The Impact on Management, Information and Communication and Material Flow*, Springer, pp. 215-231. 2007.

Lorenzen, L., Scholz, T., Timm, I.J., Woelk, P.-O., Denkena, B. and Herzog, O., “Process Planning and Production Control I”.*Chapter II.2 in Kirn, S. et al. (Eds.): Multiagent Engineering – Theory and Applications in Enterprises,* Springer, Berlin, pp. 91-113. 2006

Penberthy, J. and Weld, D.S., “UCPOP: A sound, complete, partial order planner for ADL”. In*: Proceedings of the Third International Conference on Knowledge Representation and Reasoning (KR&R-92),* Morgan Kaufmann Publishers, pp. 103–114. 1992.

Schuldt, A., “Decentralisation and Interaction Efficiency in Cooperating Autonomous Logistics Processes”. *In: Kreowski, H.-J., Scholz-Reiter, B., and Thoben, K.-D. (eds.), “2nd International Conference on Dynamics in Logistics” (LDIC 2009)*, Bremen, Germany, Springer-Verlag, August 17-21, 2009.

Schuldt, A., “Multiagent Coordination Enabling Autonomous Logistics”. *Doctoral Dissertation*, Universität Bremen, March 2010.

Smith, D. E. and Weld, D. S., “Temporal planning with mutual exclusion reasoning”. *In: Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence (IJCAI-99)*, San Mateo, CA. Morgan Kaufmann Publishers, 1999.

Timm, I., Herzog, O., Tönshoff, H.K., and Woelk, P.-O., “Flexible Process Planning and Production Control Using Co-operative Agent Systems”. *In: Proceedings of International Conference on Competitive Manufacturing (COMA ’01)*, Stellenbosch, S., GCC, University of Stellenbosch, pp. 442-449. January 1-February 2, 2001.

Tönshoff, H.K., Herzog, O., Timm, I.J., and Woelk, P.-O., “Integrated Process Planning: and Production Control Based on the Application of Intelligent Agents”. *In: Ueda, K. (Eds.): “Proceedings of the 4th International Workshop on Emergent Synthesis (IWES-02)*, Kobe University, Japan, pp. 217-224. May 9-10, 2002.

de Weerdt, M. , ter Mors, A., and Witteveen, C., “Multi-agent Planning: An introduction to planning and coordination”. *In Handouts of the European Agent Summer School*, pp. 1-32. 2005.

Wolfman, S.A. and Weld, D.S., “Combining linear programming and satisfiability solving for resource planning”. *Knowledge Engineering Review*, 16(1):85–99. 2001.

A publication of the *Machine Learning and Inference Laboratory*
Department of Health Administration and Policy, College of Health and Human Services
George Mason University
Fairfax, VA 22030-4444 U.S.A.
http://www.mli.gmu.edu

Editor: Janusz Wojtusiak

The *Machine Learning and Inference (MLI) Laboratory Reports* are an official publication of the Machine Learning and Inference Laboratory, which has been published continuously since 1971 by R.S. Michalski's research group (until 1987, while the group was at the University of Illinois, they were called ISG (Intelligent Systems Group) Reports, or were part of the Department of Computer Science Reports).