AI Magazine Volume 14 Number 2 (1993) (© AAAI)

The Ninth International Conference on Machine Learning

Alberto Maria Segre

■ The Ninth International Conference on Machine Learning was held in Aberdeen, Scotland, from 1–3 July 1992, with 198 participants in attendance. The conference covered a broad range of topics drawn from the general area of machine learning, including concept-learning algorithms, clustering, speedup learning, formal analysis of learning systems, neural networks, genetic algorithms, and applications of machine learning. This article briefly touches on six selected talks that were of exceptional interest.

The Ninth International Conference on Machine Learning was held in Aberdeen, Scotland, from 1–3 July 1992. Conference organizers were Derek Sleeman (conference chair) and Peter Edwards (local arrangements chair), both of the University of Aberdeen.[1]

Since the first machine-learning workshop was held at Carnegie-Mellon University (CMU) in July 1980, meetings have been held regularly, alternating between a more formal conference format and a more informal workshop format. This summer's conference was the first to be held outside the United States. A total of 198 delegates from Europe, Australia, Japan, Taiwan, and North America attended the 3-day event.

All conference sessions were held on the stately campus of King's College (founded 1483) of the University of Aberdeen. Three 1-hour invited lectures were presented, one on each day of the conference. Ivan Bratko (Josef Stefan Institute, Slovenia) described his recent work on the interface between machine learning and qualitative reasoning, an important area for expert system applications; David Klahr (Carnegie Mellon University [CMU], Pittsburgh, Pennsylvania) spoke on the relationship between machine-discovery systems and discovery in children and adults; and Jude Shavlik (University of Wisconsin at Madison) delivered a presentation on his work in integrating symbolic and neurally inspired learning systems. In addition to the invited presentations, there were nine 1-1/2–hour plenary sessions (with three papers presented at each session), an evening poster session, and a final panel discussion. After the conference, some participants elected to stay an extra day to participate in one of several informal workshops.

Because it is impossible to review all the papers presented, this article briefly touches on six selected talks that were, in my opinion, of exceptional interest. These capsule summaries are intended to serve as a representative sample of the research presented at the conference; the interested reader is referred to the conference proceedings.[2]

On the first day, Philip Laird (NASA Ames Research Center, Mountain View, California) presented his work on dynamic optimization of pure Prolog programs. Laird's work is an example of *speedup learning*, where the objective is to make a system improve its performance with experience. His dynamic optimization system relies on collecting use statistics at run time that are subsequently analyzed and used to reconfigure a pure Prolog program through unfolding and clause reordering. This reconfiguration produces a logically equivalent but faster version of the original program. Laird's system is a greedy one, using hill climbing to suggest incrementally faster transformations but never guaranteeing that the optimal program is attained. The general idea is that one can automate the process of careful analysis of run-time statistics collected over a set of examples to produce appropriate program transformations. The claim is that these techniques are applicable to a wide variety of nondeterministic typed-term languages. Laird's results were illustrated with several examples and justified through an empirical evaluation.

A second memorable presentation, also on the first day, was given by Alan Christiansen (CMU), who described his work on learning to predict in uncertain, continuous environments. The problem addressed here is that of learning descriptions of primitive actions in domains with noise and uncertainty. His algorithm, FB, is an example of an inductive-learning algorithm that is tolerant of noise and operates on geometric descriptors. FB acquires descriptions of *funnels*, operators whose preconditions and postconditions are defined by geometric regions of the state-action space. Once the descriptions are obtained, planning becomes simply a matter of search to identify a sequence of funnels with high probability of success. Like Laird, Christiansen used an empirical test on the classic tray-tilting manipulation problem to evaluate the effectiveness of his approach.

In a similar vein, work presented by Sridhar Mahadevan (IBM T. J. Watson Research Center, Yorktown Heights, New York) applied statistical clustering techniques to the problem of learning action models for a robot-navigation problem. Mahadevan uses a reinforcement learning technique called Q learning to acquire an estimator that specifies which action is most appropriate under certain sets of input. The interesting result is that the action models acquired are not task specific but, rather, carry over from one task to the next (unlike the output of Christiansen's FB algorithm, which is task specific). Mahadevan's ideas are applied to a robot-navigation task, where the robot, guided by simple sonar and infrared sensors, learns to follow walls and avoid obstacles.

Claude Sammut (University of New South Wales, Australia), with Scott Hurst, Dana Kedzier, and Donald Michie, presented a case study of a system that uses inductive-learning methods to acquire a motor skill by observing human subjects. The skill in this case is flying an airplane; the input to the system are logs of human subjects guiding a flight simulator. The output of the learning system is then used in autopilot mode to fly the airplane. This task is obviously a complex control task with many input parameters, yet it illustrates how a well-understood technique (here, inductive inference as performed by Ross Quinlan's C4.5 system) can be applied in a novel fashion to problems of practical import with good results. In fact, the authors report that the synthetic autopilots that their system generated were often able to perform the same flight tasks in a smoother fashion than their original human subjects.

Cullen Schaffer (City University of New York) reported—in his usual energetic and entertaining manner—on the wisdom of applying overfitting avoidance strategies in inductive decision tree learning. *Overfitting*, or overspecialization, refers to a problem encountered by inductive-learning methods, whereby the output they construct are too closely linked to the training set. Several strategies for avoiding overfitting have been proposed, and the conventional wisdom is that applying such strategies helps improve the accuracy of induced decision trees. In this talk, however, Schaffer illustrated how all these techniques simply encode some implicit bias, and unless the bias they encode happens to match the particular application domain, they cannot be expected to reliably improve the performance accuracy of the output. The message of his talk was that because one cannot propose a domain-independent avoidance strategy, one should focus on identifying characteristics of the application domain and selecting an overfitting avoidance strategy on this basis.

Larry Hunter (National Library of Medicine, Washington, D.C.), with Nomi Harris and David States, presented work on the application of an unsupervised learning concept algorithm to large, unsegmented data streams in the presence of noise. This research area is interesting given the importance of biological applications of computer science as well as the need to deal effectively with large amounts of data in other areas of science (for example, astronomy, drug manufacture). The application explored here involves detecting *motifs*, or repeated patterns in protein sequences. Proteins are composed of amino acids in a linear sequence and then folded into three dimensions. The problem is to detect recurring subsequences of varying sizes, where each subsequence can differ slightly thanks to evolution (where one can think of evolution injecting noise into the matching process). Using the existing protein string matching algorithm BLAST, which identifies variable-length portions of a protein pair exhibiting greater than random similarity, the authors propose a means for clustering individual BLAST hits to identify longer sequences of reasonable similarity. The system is tested on both synthetic data as well as data from a real protein sequence database containing roughly 62,000 proteins. The results obtained with BLAST correlate well with the results obtained manually (and with great difficulty) by the biological community.

The remaining papers presented in plenary session covered a broad range of areas—formal analysis of learning systems, genetic algorithms, neural networks, and a variety of other areas of interest to the machine-learning community as well as the AI community as a whole. A final conference banquet was held at the Pittodrie House Hotel in pleasant and sophisticated surroundings: It was a fine ending, befitting an exceptionally educational and informative three days. The next conference will be held in Amherst, Massachusetts, from 27–29 June 1993.[3]

Notes

1. Financial support for the conference was provided by the U.S. Office of Naval Research, Computer Rentals (Scotland), the European Community (DGXIII), the Bank of Scotland, the City of Aberdeen, and the University of Aberdeen.

2. *Machine Learning: Proceedings of the Ninth International Workshop (ML92)*, eds. D. Sleeman and P. Edwards, San Mateo, Calif.: Morgan Kaufmann, 1992. Note that the proceedings volume is mistakenly labeled workshop instead of conference.

3. Inquiries should be directed to Paul Utgoff (utgoff@cs.umass.edu), program chair.

Alberto Maria Segre has been an assistant professor at Cornell University since 1987, where his research involves machine-learning systems for planning. After receiving undergraduate degrees in both music and electrical engineering from the University of Illinois, Segre spent a year as a Fulbright scholar, working in computer music at the University of Milan, Italy. He returned to the University of Illinois where he obtained his Ph.D. in electrical engineering for his work on machine learning and robotic assembly.