



Workshop on Fielded Applications of Machine Learning Final Report on ONR Grant No. N00014-93-1-0209

Background and objectives

One of the central insights of artificial intelligence is that expert performance requires domain-specific knowledge, and work on knowledge engineering has led to many AI systems that are now regularly used in industry and elsewhere. The ultimate test of machine learning, the subfield of AI that studies the automated acquisition of knowledge, is the application of its techniques to produce similar results. Recent successes in real-world applications of machine learning suggest the time was ripe for a meeting on this topic.

For this reason, *Pat Langley* (Siemens Corporate Research) and *Yves Kodratoff* (Universite de Paris, Sud) organized an invited workshop on applications of machine learning. The goal of the gathering was to familiarize participants with existing applications of computational learning methods and to explore the potential for additional ones in the private and public sector. To this end, it emphasized fielded applications that are in actual use, and it downplayed differences among the specific learning methods employed, focusing instead on the machinations necessary to obtain successful results in real-world domains.

The meeting took place at the University of Massachusetts, Amherst, on June 30 and July 1, 1993, immediately following the Tenth International Conference on Machine Learning. Approximately 30 participants listened to 12 invited presentations, most of which dealt with specific applications of machine learning. The attendees also took part in lively discussions about the issues that arise in developing fielded applications, the relation of such work to the rest of machine learning, and the potential for future applications.

In this report we summarize the talks presented at the workshop, in each case describing the application domain, the basic approach taken, and the status of the resulting system. After this, we make some general observations about the state of the field and its potential for the future. Appendix A presents a list of the invited speakers and their addresses.

Reducing delays in rotogravure printing

Robert Evans (R. R. Donnelley & Sons) reported on his work with Doug Fisher (Vanderbilt University) on process control for rotogravure printing. This task involves pressing a continuous supply of paper against a chrome-plated, engraved copper cylinder that has been bathed in ink. Sometimes grooves or bands develop on the cylinder during the printing process, appearing in turn on the printed pages; this requires the print run to be halted and, in some cases, the cylinder to be replaced, costing time and money for the printer. The reasons for banding are largely unknown, but Evans and Fisher collected positive and negative cases of banding, along with environmental factors present in each case, then used machine learning methods to induce a decision tree that predicts the probability of banding. One Donnelley plant now uses the decision tree to set ink viscosity and similar factors, which has almost entirely eliminated the banding effect. Evans' talk addressed the relative roles of data collection, representation engineering, weak domain expertise, and induction in their discovery of banding rules.

Autoclave layout for aircraft parts

David Hinkle (Lockheed AI Center) described his work with Chris Toomey on CLAVIER, a case-based reasoning system for layout design. Most modern aircraft are made from composite materials, which must be cured in a large convection oven called an autoclave. Tables that hold sets of such parts pass slowly through the oven, but some layouts cook swiftly, others slowly, and some actually damage the parts, forcing them to be replaced. A good layout includes many parts and cooks quickly without causing any damage, but the heating properties of autoclaves are only poorly understood. Hinkle and Toomey's CLAVIER uses

DISTRIBUTION STATEMENT A

Approved for public release
Distribution Unlimited

DTIC QUALITY INSPECTED 3

94 8 15 253

19950922 065

a case-based method to store successful layouts, retrieve them for use in novel situations, and adapt them where necessary. The system also uses a heuristic scheduler to generate a sequence of loads that best meets production goals while satisfying operational constraints. The resulting advisory system (which the domain expert can always overrule) has been in daily use on the shop floor at a Lockheed factory since 1990, where it suggests nearly identical layouts as does the expert.

Diagnosis of Mechanical Devices

Lorenza Saitta (University of Torino) described joint work between researchers at her university and ones at Sogesta, a large Italian chemical company, on the use of machine learning in developing an expert system. The task involved fault diagnosis for electric motor pumps, which play a major role in the company's production process. Starting from an initial diagnostic system that had been manually elicited through standard knowledge acquisition techniques, they generated improved versions of the system using induction methods that were capable of drawing on background knowledge, including a causal model of the domain. The resulting learned knowledge base has replaced the hand-crafted one in the operational expert system. Saitta's talk focused on the motivations for the development effort, the difficulties they encountered, evaluation of the fielded system, and on the reasons for its success.

Automatic classification of sky objects

Usama Fayyad (Jet Propulsion Laboratory) reviewed the results of the second Palomar Observatory Sky Survey, which has produced about three terabytes of image data, containing nearly a billion sky objects. Clearly, astronomers could not hope to classify these objects manually, and in response, Fayyad and his colleagues have developed SKICAT, a system that automatically catalogs sky objects in the survey's digitized photographic plates. First they used image processing techniques to describe a set of objects in the images, which astronomers then labeled for use in training. They then used machine learning methods to induce a decision tree that classified objects as members of one class or another. The classification accuracy on new images (94%) was above the level specified by astronomers as necessary for use in scientific data analysis, and the decision tree is currently being used to automatically classify all objects in the Sky Survey images, which would be impractical for humans. The objects classified in this manner are ten times fainter than any cataloged in large-scale surveys to date, producing a catalog at least three times the size possible had machine learning not been employed. Fayyad's talk dealt with both the techniques needed to apply the learning algorithms and the database work needed to make the tool useful to astronomers.

Predicting pilot bids

Pieter Adriaans (Syllogic) described his experiences in developing CAPTAINS, an AI system that enables a planner to maintain strategic, tactic, and operational models of pilot populations. Twice a year, pilots for the airline KLM can express their preference for "seats" on different airplanes, and the company is required by contract to give each seat to the most senior qualified pilot. Accurate prediction of these bids would let the airline decide how many new pilots to train for vacated positions, reducing their costs considerably. Adriaans used genetic algorithms and historical data on pilot bids to produce a set of predictive rules, which he then embedded in the CAPTAINS system and which KLM now uses in its planning process.

Automated Completion of Repetitive Forms

Jeffrey Schlimmer (Washington State University) presented his work with Leonard Hermens on a learning apprentice for the completion of forms. This activity occupies much of people's time in both business and government agencies, since most of them are filled in by hand. Yet much time and effort has been expended to automate form-filling by programming specific systems on computers. The high cost of programmers and

Dist	Special
A-1	

other resources prohibits many organizations from benefiting from efficient office automation. Schlimmer argued that a learning apprentice can be used to acquire the knowledge for such repetitious form-filling tasks in a cost-effective manner. He also described a framework for such a system, explained the difficulties of form filling, and presented empirical results of a form-filling system used in his department for eight months. The form-filling apprentice saved up to 87% in keystroke effort and correctly predicted nearly 90% of the values on the form.

Machine Learning Support for Help Desks

Brad Allen (Inference Corporation) reported on CBR EXPRESS, a software tool for constructing help desk advisory systems. When the users of a computer, copier, or other complex device encounter difficulties, they often call the maker's help desk for advice on how to correct the problem. However, typically few people in the company have the expertise needed to answer all such queries, and their time is valuable. This led Allen and his colleagues to develop CBR EXPRESS. The system stores specific cases of previously encountered problems, along with their solutions, in memory, and uses a simple nearest neighbor algorithm to retrieve cases that are similar to ones described by callers. The retrieval process is iterative, with the help desk consultant asking questions which lead to promising cases, which in turn suggest additional discriminating questions, and so forth, eventually leading to a few likely cases with recommended actions. The consultant adds solved problems to the case library for future use, so the knowledge base grows over time. CBR EXPRESS has been sold to over 80 companies, a number of which have used it to develop fieldied advisory systems.

Predicting Activity in the Automobile Market

Reza Nakhaeizadeh (Daimler-Benz) described his use of machine learning to produce predictive models of automobile activity. Each year, the marketing department of Mercedes-Benz predicts the the number of cars and trucks that will be registered in more than 80 countries. The management then uses these predictions to develop short-term and long-term plans for production. Thus, they are interested both short-term predictions (quarterly, annual) and long-term ones (five to ten year). The data used are the historical time series for cars and trucks and the historical values of the external economic attributes like GNP, prices, inflation rate, interest rate, most of which are available in quarterly and annual periods. In contrast to the classification tasks that predominate in machine learning research, this domain requires the prediction of continuous values, but some machine learning algorithms – like NEWID and CART – can handle such numeric data sets. After some initial experiments with the data for European countries, Nakhaeizadeh concluded that, despite the small number of training cases, machine learning approaches can predict the time series at least as well as the regression analyses that were currently in use within the company. Thus, his group developed an advisory tool that incorporated a PC version of NEWID, a version of regression analysis, and a data preprocessing algorithm. This tool is now in use by the Mercedes-Benz marketing department and supports users in predicting market activity, letting them compare the results achieved with the different approaches.

Machine Learning in Text Retrieval

David Waltz (Thinking Machines Corporation) discussed the task of information retrieval in large databases. For example, financial analysts would like rapid access to news stories that are relevant to their concerns without having to wade through irrelevant ones. In response to demand for such capabilities, Waltz and his colleagues developed a case-based learning system that stores specific stories, indexes them in memory by key words and phrases, and selectively retrieves them, using an appropriate distance metric, in response to queries formulated by users. The resulting system has been used extensively by hundreds of clients of a major financial company.

Helping Domain Experts Formulate Applications

Derek Sleeman (University of Aberdeen) reported on CONSULTANT, an advisory module for the Machine Learning Toolbox (MLT), a collection of software tools designed for intelligent data analysis and knowledge acquisition. These tools are complex programs and one must consider a variety of factors when selecting a tool for a particular machine learning application. Thus, Sleeman and his colleagues developed the CONSULTANT system to assist users in selecting a suitable tool. However, he noted that as the Toolbox was implemented, it became clear that they had underestimated the amount of assistance needed by nonexperts in machine learning, and that insights gained into the application of machine learning during the project identified more sophisticated forms of help that could be given to the user. These factors led to significant enhancements of the CONSULTANT system. Sleeman's talk described how the CONSULTANT evolved from its original specification and the motivation behind these changes. He also examined in detail the role played by the system in one of the applications efforts.

Conclusions from the workshop

In addition to the above talks, most of which focused on specific applications, two additional presentations attempted to draw some generalizations about the field as a whole. *Pat Langley* (Siemens Corporate Research) opened the meeting by reviewing some additional fielded applications and raising some challenges for future work. In particular, he noted the importance of problem formulation and representation engineering in many applications efforts, and that closer study of such activities might suggest ways to automate these processes. He encouraged speakers to emphasize their application domain and the obstacles encountered on the development path, and to downplay the particular induction algorithms used in their work.

Patricia Riddle (Boeing Aircraft Company) gave a commentary on the applications presented at the workshop. She proposed a number of distinctions among the approaches taken, laying the groundwork for a useful taxonomy of learning applications. For instance, she distinguished between *learned* systems, which are produced using machine learning techniques but do not use them during performance of their task, and *learning* systems, which also learn during their use. She also noted that the methods used in application efforts ran the gamut of learning techniques, including methods for inducing rules and decision trees to case-based and instance-based schemes, though connectionist algorithms were not represented at the meeting. The tasks addressed also covered a broad spectrum, including mechanical diagnosis, configuration and layout, planning, and process control, but in most cases developers had found ways to transform their problems into simple tasks that involved classification or prediction, for which robust induction algorithms exist.

A lively, extended discussion took place after the formal presentations, with many attendees contributing. Topics ranged from the types of problems encountered during the development process to ways to encourage increased applications work and increasing the academic respectability of such efforts. A proposal to establish a regular conference on applications of machine learning, separate from the annual research conference, was generally felt to have more disadvantages than benefits. Participants left the meeting with high hopes about the potential of learning algorithms for use on real-world problems, and agreed that the field of machine learning should concentrate a substantial portion of its energies toward developing additional fielded applications.

19950922 065

Appendix A: Invited Speakers for the Workshop

Pieter Adriaans
Syllogic B. V.
Postbus 26, NL-3990 DA Houten
THE NETHERLANDS
(31) 3403-51110
PIETER@SYLLOGIC.NL

Brad Allen
Inference Corporation
550 North Continental Blvd., Third Floor
El Segundo, CA 90245 USA
(310) 322-0200
ALLEN@INFERENCE.COM

Robert Evans
R. R. Donnelley and Sons
801 Steam Plant Rd
Gallatin, TN 37066 USA
(615) 230-1374

Usama Fayyad
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109 USA
(818) 306-6197
FAYYAD@AIG.JPL.NASA.GOV

David Hinkle
Lockheed AI Center
3251 Hanover Street
Palo Alto, CA 94304-1191 USA
(415) 354-5237
HINKLE@AIC.LOCKHEED.COM

Pat Langley
Siemens Corporate Research
755 College Road East
Princeton, NJ 08540 USA
(609) 734-6574
LANGLEY@LEARNING.SIEMENS.COM

Gholamreza Nakhaeizadeh
Daimler-Benz AG
Forschung und Technik
Wilhelm-Runge-Str. 11
D-7900, Ulm, GERMANY
REZA%FUZI.UUCP@GERMANY.EU.NET

Patricia J. Riddle
Boeing Computer Services
P. O. Box 24346, MS 7L-66
Seattle, WA 98124-0346 USA
(206) 865-3415
RIDDLE@GRACE.RT.CS.BOEING.COM

Lorenza Saitta
Department of Informatics
Università di Torino
Corso Svizzera 185, 10149-Torino, ITALY
(11) 771-2002
SAITTA@DI.UNITO.IT

Jeffrey C. Schlimmer
School of Electrical Engineering & Computer Science
Washington State University
Pullman, WA 99164-2752 USA
(509) 335-2399
SCHLIMME@EECS.WSU.EDU

Derek H. Sleeman
Department of Computing Science
University of Aberdeen
Aberdeen, AB9 2UB, SCOTLAND
(224) 272-288
SLEEMAN@COMPUTING-SCIENCE.ABERDEEN.AC.UK

David Waltz
Thinking Machines Corporation
245 First Street
Cambridge, MA 02142 USA
(617) 876-1111
WALTZ@THINK.COM