

Editorial

Machine learning continues to be an exciting and growing area of artificial intelligence, with new developments in topics as diverse as theoretical foundations, analytical (explanation-based) and inductive (similarity-based) techniques, integrated planning and learning architectures, connectionist learning methods, genetic algorithms, learning in analogical and case-based reasoning systems, and the gradual infusion of learning into commercial applications—primarily in the guise of decision tree induction and the semi-automatic acquisition of rules from data. The journal has played an active role as a dissemination forum for scientific papers combining theoretical analysis with empirical validation. Its purpose has been three-fold: the archival vehicle for established scientific results, a means of educating researchers and students in the field, and as a rapid communication device for new results and ideas. *Machine Learning* has succeeded admirably in the first objective (due to the quality of the research and the diligence of the authors), has done very well on the second (largely through the editorial insistence and hard work of Pat Langley), but has not as yet truly faced up to the third.

With the publication of this volume the journal is transitioning from the very able management of Pat Langley in years past to my more distributed (but I hope equally able) management style. Although some important changes will be instituted, in most ways *Machine Learning* will continue steady on course. We still seek the same high-quality, well-written scientific papers, combining where possible theory, system implementation, and thorough empirical evaluation. Pat Langley's past editorial on how to write the ideal paper is still very much valid. We intend to continue with our high standards in the review process, both for scientific content and coherent clear writing. The coming changes fall into two categories: broadening of perspectives for the journal and purely administrative issues. Most of the latter should remain invisible to our authors and readers; they are just different means to achieve the same ends.

I believe that *Machine Learning* should publish articles encompassing a broader range of topics, including: connectionist learning methods, knowledge acquisition for expert systems, and full integration of learning methods into complex systems (such as planners, problem solvers, schedulers, robotic and process controllers, etc.). We actively solicit papers in these areas as well as in the more traditional foci of machine learning. To promote this desired expansion of our scientific horizon, we are encouraging special-topic issues of the journal in each of these areas. In fact, a special issue on knowledge acquisition is already in the works and should appear shortly. Special issues on the other topics are currently being discussed and planned in a one- to two-year time frame.

In addition to broadening the scientific scope of the journal, we also wish to broaden its geographic scope by encouraging our European, Japanese, and other

international colleagues to submit their latest research results to *Machine Learning*. We hope that the additional burden of producing quality English prose will not deter too many potential international contributors. We will also make an effort to broaden somewhat the composition of the editorial board both in scientific and geographical directions. However, we remain a journal dedicated to quality research articles on machine learning, and we will not expand beyond that general mandate.

In order to make the journal more reactive to new research results, and to cut down the time submitted papers remain in the reviewing pipeline, we will institute the following two changes, effective immediately:

- **Multiple Editors:** We have a larger number of active editors, each expert in one or more subfields of machine learning. I will scan all submitted papers to ensure relevance to machine learning and to assign them to an editor, who will then assume the responsibility of the entire review process, including the selection of reviewers and all direct communications with the author. I hope that this distributed approach will expedite reviewing by removing the obvious central bottleneck once the paper is properly dispatched.
- **Scientific Notes:** In order to start addressing the rapid-communication function of the journal, we are starting a scientific notes section. Each note should contain one key technical contribution such as an important new idea not yet validated with a complete system implementation, an improvement of a published method, a tabular comparison of several methods on a comprehensive data set, a validation or refutation of an earlier result by repetition of the experiment under possibly different conditions, focused commentary on a recently published article, and so forth. All scientific notes will be reviewed, often by a single person (editor or reviewer), and for those accepted we guarantee publication in the next issue that enters the publication cycle after receipt of the final manuscript, regardless of the regular paper backlog. However, notes must be short, to the point, and clearly written. The contribution by Kurt VanLehn in this issue is an example of such a note, albeit a particularly detailed one—shorter ones are also welcomed.

I hope that *Machine Learning* continues to flourish in the coming years, and I am open to suggestions and remarks, best transmitted to me by electronic mail (Jaime.Carbonell@cs.cmu.edu on internet) with the keyword “journal” on the subject field.

JAIME G. CARBONELL
CARNEGIE MELLON UNIVERSITY