

Preface

This proceedings contains a written record of the Twentieth International Conference on Machine Learning (ICML-2003) held in Washington, D.C. on August 21-24th, 2003. This year's conference marked several firsts for ICML:

For the first time, ICML was collocated with KDD (the International Conference on Knowledge Discovery and Data Mining). We are hopeful that the collocation will encourage further cross-fertilization between the two communities. ICML-2003 and KDD-2003 shared a joint session and an invited speaker. We thank Pedro Domingos, Christos Faloutsos and Ted Senator of KDD-2003 for helping make this possible.

For the first time, the conference offered a Student Paper Award program, sponsored by Kluwer and the *Machine Learning* Journal. We were fortunate to have Robert Holte chair and set the precedent for the Student Paper Award program.

In 2002 the International Machine Learning Society (IMLS) was formed as an oversight/advisory committee for the international conference. ICML-2003 was the first conference held under the Society's guidance.

This year we also continued with many of the past ICML traditions. We once again welcomed a collocation with our theoretical counterpart COLT '03 (the Sixteenth Annual Conference on Learning Theory). The conditional accept program was continued although changed in a minor but helpful way. To prevent overuse of the conditional accept option, the conditional accept decision was moved from the program committee members to the area chairs and the area chairs reviewed the revised submissions. As a result, far fewer papers were conditionally accepted—only about 7%.

The response to our call for papers was strong. We received 371 full-length research papers, of which 96 were accepted outright and 27 were conditionally accepted. Almost all conditionally accepted papers were finally accepted (23 of 27), yielding a final acceptance rate of about 32%.

We were delighted to host three distinguished invited speakers. Daphne Koller gave a joint ICML-KDD plenary talk on learning probabilistic relational models. Michael Kearns and Foster Provost gave invited presentations depicting the complementary roles of theory and applications in machine learning.

Assembling a conference of this magnitude requires a significant effort and we were very fortunate to have a talented and supportive team. Our twenty area chairs served as experts in various subfields of machine learning. The area chairs nominated program committee members, moderated discussions between reviewers whenever there was a disagreement, and made final accept/reject decisions. The program committee provided competent, high quality reviews. Dunja Mladenic, our workshops and tutorials chair, arranged an interesting and timely set of workshops and tutorials.

We are indebted to Richard van de Stadt of Borbala Conference Services for his skillful and responsive support of the CyberChair Pro conference software. Richard made countless changes to his software to accommodate our many specific needs. We owe a gigantic thank you to Richard. Stanford University graciously allowed us to use their systems to host the paper submission web site. Hewlett-Packard hosted the conference web site.

Terran Lane, our proceedings and publications chair, had the enormous task of assembling the proceedings. In addition, when our original publisher backed down, Terran quickly scrambled to find a new publisher. Tim Oates, our local arrangements chair, found a superb location for the conference and facilitated our intermingling with KDD. And Andrew McCallum, our fundraising chair, lined up corporate sponsors and also arranged student scholarships.

Many individuals offered us guidance throughout the process. The advisory committee (whose names are listed later in the proceedings) aided in the selection of invited speakers and many other conference-related matters. We are especially grateful to Tom Dietterich, Pat Langley, and Umesh Dayal for their frequent words of wisdom and encouragement.

We are grateful for the financial support provided by our corporate sponsors Kaidara Software, Microsoft Research, and Overture. In addition, Hewlett-Packard Labs generously allowed both of us to spend a significant fraction of our time putting the conference program together.

Finally, we wish to thank the authors and all other contributors for their time and effort in creating such a fine program.

– Tom Fawcett and Nina Mishra
ICML-2003 Program Cochairs