

The First Conference on E-mail and Anti-Spam

David Heckerman, Tom Berson,
Joshua Goodman, and Andrew Ng

■ The First Conference on E-mail and Anti-Spam was held from July 30 to July 31, 2004 in Mountain View, California. The conference, attended by 180 researchers, featured 29 papers that covered a number of topics, including e-mail in general, nonstatistical techniques for stopping spam, machine learning techniques, issues of identity in e-mail, as well as law and policy. The 2005 conference will be held at Stanford University from July 21 to 22.

The First Conference on E-mail and Anti-Spam was held from July 30 to July 31, 2004 in Mountain View, California. The conference organizers requested papers on all aspects of spam and e-mail; received 81 submissions, and accepted 29 papers.¹ There were 180 attendees.

E-mail has grown from a tool used by a few academics on the Arpanet to a ubiquitous communications tool. It has evolved from a piece of simple, plain text in an inbox into a rich graphical medium that can be viewed, sorted, signed, encrypted, shared, archived, searched, prioritized, and so on. Spam, following the growth of e-mail, has changed from a minor curiosity, to a nuisance, to a multibillion-dollar problem.

There have been numerous industrial conferences on spam and e-mail—mostly concentrated on shipping products and other practical concerns. The organizers of this conference, however, saw the need for an academic-style conference, peer-reviewed, with published proceedings, and, most importantly, focused on the kind of carefully done, thorough, long-term research needed by the academic and industrial research communities.

Of course, many academic papers,

in fields as diverse as machine learning, cryptography, natural language processing, systems, security, and human computer-interaction, have been published about spam and e-mail. The Conference on E-mail and Anti-Spam was the first to bring together researchers from these varied fields to a single academic conference where they could exchange ideas and compare techniques.

Many topics were covered at the conference. Five papers considered e-mail in general, unrelated to spam. Of course, spam is the biggest problem afflicting e-mail today, and seven papers considered nonstatistical techniques for stopping it. Additionally, four papers looked at statistical or machine learning techniques. Two papers considered network techniques for stopping spam, and two others investigated the techniques used by spammers and phishers (e-mail scam artists). Three papers considered issues of identity in e-mail. We were pleased to have six papers on law and policy, and an invited talk by Lawrence Lessig, professor of law and the John A. Wilson Distinguished Faculty Scholar at Stanford University. Hal Varian (Class of 1944 Professor at the School of Information Management and Systems, Haas School of Business, and the Department of Economics at the University of California, Berkeley) also presented an invited talk—entitled “Who Signed up for the Do Not Call List.” Finally, five panelists discussed the topic “Payment for Anti-Spam: Time, Money, and Puzzles.”

This conference was the first in an annual series. The Second Conference on E-mail and Anti-Spam will be held from July 21 to July 22, 2005 at Stanford University.²

Notes

1. The papers are available at www.ceas.cc.
2. See www.ceas.cc for more information.



David Heckerman is founder and manager of the Machine Learning and Applied Statistics Group at Microsoft Research. Since 1992, he has been a senior researcher at Microsoft, where he cocreated applications such as the first machine-learning spam filter, data-mining tools in SQL Server and Commerce Server, handwriting recognition in the Tablet PC, text-mining software in Sharepoint Portal Server, troubleshooters in Windows, and the Answer Wizard in Office. Heckerman received his Ph.D. (1990) and M.D. (1992) from Stanford University and is an AAAI Fellow.



Tom Berson consults on information conflict and on cryptography, which he describes as “the mathematics of trust and betrayal.” He is a Fellow of the International Association for Cryptologic Research and a former editor of the *Journal of Cryptology*.



Joshua Goodman is a researcher in Microsoft Research’s Machine Learning and Applied Statistics group. Goodman helped start and spent two years working with Microsoft’s Anti-Spam Technology and Strategy Team. He

has a Ph.D. from Harvard University in natural language processing. In 1998, Goodman joined Microsoft Research, where he has worked on language modeling, text input, and spam filtering.



Andrew Y. Ng is an assistant professor in the Computer Science Department at Stanford University. His research focuses on machine learning, including its application to text processing and to control.

He built “ML Papers,” the first automatically indexed, publicly available search engine for research papers on the web. His autonomous helicopter was also the first to be capable of sustained inverted (upside-down) flight under computer control. His Ph.D. is from the University of California, Berkeley