
Preface to the Proceedings of AISTATS 2011

Geoffrey J. Gordon

Machine Learning Dept., Carnegie Mellon

David Dunson

Dept. of Statistical Science, Duke

The Society for Artificial Intelligence and Statistics (SAIAS) is dedicated to facilitating interaction between researchers in Artificial Intelligence and Statistics. The Society's main responsibility is to organize the International Conference on Artificial Intelligence and Statistics, AISTATS (www.aistats.org). We acted as program chairs for the fourteenth conference, which took place April 11–13, 2011, in Fort Lauderdale, Florida, USA. As has been the tradition in recent odd-numbered years, AISTATS was colocated with the Learning Workshop (also known as “Snowbird”), which took place April 13–16; registrants of both conferences were invited to attend a day of joint sessions on Wednesday, April 13.

We made a number of changes to the reviewing process this year, intended to reduce reviewer workload and increase the quality of final accepted papers. By and large we believe these changes were successful (see details below): we attracted the largest number of submissions ever for an AISTATS meeting held in the United States, 272 papers, and had a great deal of trouble narrowing down to 99 acceptances (also a record for U.S. meetings) due to the high quality of submissions. The 99 acceptances (36.4% of submissions) were divided into 78 poster presentations and 21 oral presentations, each of which was allocated 8 pages in the proceedings, plus extra space for the bibliography. We kept reviewing loads low to ensure high-quality reviews: no more than about 15 papers per senior program committee member, and no more than about 5 papers per reviewer (and typically less than these limits). Each paper received at least two, and usually three or more reviews; overall, as the people who had to read all of the reviews, we are happy to report that the reviews were generally of extremely high quality.

In addition to the 21 contributed oral presentations, we had 3 invited speakers, selected by the senior program committee and the chairs for their expertise and gen-

eral interest: Adrian Dobra (University of Washington) spoke on “Multi-way Gaussian Graphical Models with Application to Multivariate Lattice Data,” Martin Wainwright (Berkeley) spoke on “Convex Relaxation and Estimation of High-Dimensional Matrices,” and Bin Yu (Berkeley) spoke on “Sparse modeling: some unifying theory and ‘topic-imaging.’”

We made four main changes to the reviewing process this year: Notable Paper awards, discussion papers, an expanded author feedback phase, and a two-phase reviewing process. The first three changes were unqualified successes, and we understand that future AISTATS meetings plan to incorporate them; the last change (two-phase reviewing) had both positive and negative effects.

We instituted the Notable Paper awards in place of the usual Best Paper awards. Our observation was that the tiny number of Best Paper awards, as well as the inherent variability in the review process, meant that many authors viewed a Best Paper award as a somewhat random event, and therefore were not strongly motivated by the possibility of receiving one. To combat this perception, we increased the number of awards to 7, or 2.6% of submissions. While still small enough to confer prestige, this number of awards meant that authors viewed them as achievable and therefore a significant motivating factor.

We instituted discussion papers in an attempt to avoid wasting the excellent work of our volunteer reviewers. In most conferences, reviewers spend long hours of work understanding papers and placing them in relation to the literature, but this work results in only a single bit of information that is available to the rest of the world (the final acceptance decision). Instead, for each notable paper, we invited an expert discussant to write a short summary of the reviews for the benefit of future readers, and to present this summary at the conference immediately after the authors' oral presentation. The discussants were sometimes Senior Program Committee members, sometimes reviewers, and sometimes outside experts. In each case, the discussant had anonymous access to the paper's reviews. (We chose to invite discussants only for notable papers

Appearing in Proceedings of the 14th International Conference on Artificial Intelligence and Statistics (AISTATS) 2011, Fort Lauderdale, FL, USA. Volume 15 of JMLR: W&CP 15. Copyright 2011 by the authors.

for two reasons: first, since it was an experiment, we didn't want to try it for all papers, and second, the additional attention was an extra benefit for award recipients.) We believe that the experiment was successful, and that the discussion papers will help future readers quickly evaluate their corresponding original papers, but you can check for yourself: each discussion paper appears next to its original paper in this proceedings volume.

We instituted the expanded author feedback phase in response to author frustration with the low bandwidth of the typical author feedback channel. Instead of allowing only a plain-text response of a few hundred words (as had been previous practice), we allowed the authors to submit a completely-revised version of their paper taking into account the reviewers' feedback, as well as a plain-text response summarizing the changes for the reviewers' benefit. As a result, we noticed a clear benefit to the readability of the final papers. In addition, the higher-bandwidth feedback channel helped authors convince reviewers: as we read over the final reviews, we saw again and again that reviewers commented on how well the authors had addressed their questions—and this comment was noticeably more frequent in papers that eventually were accepted.

We instituted two-phase reviewing for two reasons: first, we hoped to reduce reviewer effort by avoiding a third review in clear-cut cases. Second, we hoped to support the expanded author feedback phase by assigning a new reviewer to most revised papers, so that the new reviewer could approach the revised paper with a fresh mind. We did realize both of these benefits, although the number of cases where we could avoid a third review was smaller than we had hoped—there were very few clear-cut rejections, and at the other end of the scale, we needed to distinguish among poster, oral, and notable papers. (We decided to reject another common solution for reducing reviewer effort, “flash reviews”: these are short-time-frame reviews intended to quickly distinguish clear accepts and rejects. By talking with other conference chairs, we reached the conclusion that these flash reviews are too noisy, and therefore not a good way to allocate effort: it just takes time to understand and evaluate a paper well.)

Unfortunately, however, the two phases significantly complicated the administration of the reviewing process, both for us and for the reviewers. While some of this complication was probably a result of lack of familiarity, some of it is probably unavoidable. There was another, unforeseen problem as well: since the second-phase reviewers submitted their reviews after the author feedback phase, the authors were unable to respond to the second-phase reviews, resulting in frustra-

tion from a small number of authors who believed there were easily-correctable bugs in their second-phase reviews. As a result of these problems, we believe that it is not a clear-cut win to include a two-phase review process in future years.

To close, we are very grateful to our anonymous reviewers and our senior program committee for their excellent work. The SPC consisted of:

- Deepak Agarwal, Yahoo!
- Edo Airoldi, Harvard Statistics
- Drew Bagnell, CMU Robotics Institute
- Howard Bondell, North Carolina State Univ. Statistics
- Michael Bowling, Univ. of Alberta CS
- Brian Caffo, Johns Hopkins Biostatistics
- Lawrence Carin, Duke Univ. ECE
- François Caron, INRIA
- Vincent Conitzer, Duke Univ. CS
- Koby Crammer, Technion EE
- Nando de Freitas, Univ. of British Columbia CS
- Feng Liang, Univ. of Illinois Urbana-Champaign Statistics
- Brendan McMahan, Google
- Iain Murray, Univ. of Edinburgh Informatics
- Ron Parr, Duke Univ. CS
- Ben Taskar, Univ. of Pennsylvania CIS
- Ricardo Silva, University College London Statistics
- Alex Smola, Yahoo! and ANU
- Larry Wasserman, CMU Statistics
- Frank Wood, Columbia Statistics
- Hao Helen Zhang, North Carolina State Univ. Statistics

In addition, the conference would not have been possible without the efforts of our Sponsorship Chair, Charles Sutton (University of Edinburgh) and our Proceedings Chair, Miroslav Dudík (Yahoo! Research).

Finally, we would like to thank our sponsors: Google and Microsoft sponsored the Notable Paper awards and the student travel assistance; PASCAL2 sponsored recording the oral presentations and hosting them on videlectures.net, and also sponsored the invited speakers; and IBM and NEC sponsored the coffee breaks.