

COLING/ACL 2006 Program Co-Chairs:

Claire Cardie
Pierre Isabelle

- Pierre and I began talking by phone last May of 2005 to determine various aspects of the joint conference.
- PLANNING MEETING AT ACL-2005. We met with Nicoletta Calzolari (general chair), and Robert Dale and Cecile Paris (local arrangements) at ACL-2005 in Ann Arbor to discuss the goals and general structure of the conference. In addition, we presented our ideas for the conference at a lunch meeting with representatives from the ACL executive committee and the ICCL, which runs COLING conferences. The main conference program will be fairly standard for an ACL conference:
 - parallel sessions
 - invited speakers or panels, one per day
 - poster sessions
 - demos (demo chair: James Curran)

At these meetings, it was decided that we would try to take advantage of the conference's location in Asia to highlight papers focusing on Asian language processing, possibly in a separate session of the main conference.

- PAPERS+POSTERS. In order to meet the combined desires of the ICCL and the ACL executive committee, we opted for a single submission and reviewing process for papers and posters. Papers were submitted to one of two categories: the REGULAR PAPER category or the POSTER category. Authors had to designate one of these categories at submission time. Although both categories of submission were the same maximum length, the presentation format would vary --- regular papers would be presented during one of the parallel paper sessions at the conference, and poster presentations would be repeated several times before small groups of people at one of the conference poster sessions. Regular papers are most appropriate for presenting substantial research results, while posters are more appropriate for presenting an ongoing research effort. Importantly for ACL, regular papers and poster papers would need to appear in separate volumes of the proceedings; and the acceptance rate for regular papers would need to stay below 25%. Importantly for ICCL, the combined acceptance rate among regular and poster papers should be around 40%.

In spite of the guidelines we placed in the call for papers and on the review form for each category of submission, the distinction caused some confusion among both authors and reviewers. In addition, the vast majority of authors opted for paper presentations (see below).

- CALL FOR PAPERS AND POSTERS. The call for papers and posters went out in early October. Handouts were included in conference materials at HLT-EMNLP in Vancouver. Richard Power offered to coordinate a mentoring service for authors from regions of the world where

English is not the language of scientific exchange.

- SCHEDULE. We created and basically stayed on the following schedule:

February 15 (weds)	reviewers recruited
February 28 (tues)	submissions arrive
March 1 (weds)	papers assigned to primary track/area
March 1-4 (weds-sun)	reviewer bidding phase
March 5 (mon)	initial assignment of reviewers to papers using START
March 7 (weds)	final assignment of reviewers to papers by area chairs
March 8 (thurs)	reviewing stage begins
April 17 (mon)	reviews due
April 17-24 (tues-mon)	e-mail/START discussion amongst reviewers and area chairs on "disagreement" papers
April 25-27 (tues-thurs)	area chair discussion/decisions on accept/reject; paper/poster
Fri, April 28	author notification

- AREA CHAIRS. Based on submission numbers from ACL 2005 (436 submissions) and ACL 2004 (351 submissions), we prepared for 500 submissions. We established 19 areas and recruited 20 area chairs (two for the machine translation area, which had 56 submissions in 2005).

Johan Bos (University of Edinburgh)
Jason Chang (National Tsing Hua University)
David Chiang (USC Information Sciences Institute)
Eva Hajicova (Charles University)
Chu-Ren Huang (Academia Sinica)
Martin Kay (Stanford University)
Emiel Krahmer (Tilburg University)
Roland Kuhn (National Research Council of Canada)
Lillian Lee (Cornell University)
Yuji Matsumoto (Nara Institute of Technology)
Dan Moldovan (University of Texas)
Mark-Jan Nederhof (University of Groningen)
Hwee Tou Ng (National University of Singapore)
John Prager (IBM Watson Research Center)
Anoop Sarkar (Simon Fraser University)
Donia Scott (Open University UK)
Simone Teufel (University of Cambridge)
Benjamin Tsou (City University, Hong Kong)
Ming Zhou (Microsoft Beijing)
ChengXiang Zhai (University of Illinois)

- SUBMISSIONS. We got more submissions than expected --- 628 vs. ~500.

paper submissions: 558 (88.9%)
poster submissions: 70 (11.1%)

Submission stats by area, showing the expected vs. the actual number of submissions are below.

AREA	expected #	actual #
	of subs	of subs

Phonology, Word Segmentation, Morphology;			
POS tagging, chunking	40	35	
Grammars/syntax	15	16	
Parsing	25	45	
Lexical Semantics	10	28	
WSD	30	20	
Inference, pragmatics	20	47	
Coreference, Discourse, Dialog, Prosody and Multi-modality	45	38	
Speech and Language Modeling	40	25	
Machine Learning methods	50	22	
Language resources, corpus annotation		10	22
Machine Translation and Multilinguality		60	81
Information retrieval and text classification, including sentiment analysis		10	56
Information Extraction	45	56	
Question-Answering	20	19	
Summarization	25	20	
Generation	25	12	
NLP applications and tools (e.g. tutoring)		20	39
Asian languages	15	47	
Totals	505	628	

As a result, we made IR its own area, and recruited ChengXiang Zhai (University of Illinois) as its area chair. In addition, many more reviewers needed to be recruited, 384 in all.

- SUBMISSION STATS BY COUNTRY/REGION. Note that these numbers were computed AFTER a number of papers were withdrawn or rejected without review.

We received submissions from 40+ countries: 39% from 13 countries in Asia, 29% from 17 countries in Europe, 25% from Canada and the United States, 4% from Australia and New Zealand, less than 1% from South America (Brazil) and from Africa (South Africa and Tunisia), 3% from 4 countries in the Middle East.

Asia=====

15 countries

266/616: 43% of submissions

CHINA: 55 (8.93%)
HONGKONG: 8 (1.30%)
INDIA: 11 (1.79%)
JAPAN: 107 (17.37%)
MALAYSIA: 2 (0.32%)
PAKISTAN: 1 (0.16%)
PHILIPPINES: 1 (0.16%)
REPUBLIC-OF-KOREA: 13 (2.11%)
SINGAPORE: 20 (3.25%)
SRILANKA: 3 (0.49%)
TAIWAN: 18 (2.92%)
THAILAND: 4 (0.65%)
TURKEY: 1 (0.16%)

AUSTRALIA: 22 (3.57%)
NEWZEALAND: 1 (0.16%)

Europe=====

17 countries
177/616= 29%

AUSTRIA: 3 (0.49%)
CZECHREPUBLIC: 7 (1.14%)
FINLAND: 2 (0.32%)
FRANCE: 34 (5.52%)
GERMANY: 36 (5.84%)
GREECE: 3 (0.49%)
HUNGARY: 1 (0.16%)
IRELAND: 8 (1.30%)
ITALY: 9 (1.46%)
NETHERLANDS: 2 (0.32%)
PORTUGAL: 4 (0.65%)
ROMANIA: 1 (0.16%)
RUSSIANFEDERATION: 1 (0.16%)
SPAIN: 17 (2.76%)
SWEDEN: 6 (0.97%)
SWITZERLAND: 4 (0.65%)
UNITED-KINGDOM: 6 (0.97%)
UNITEDKINGDOM: 32 (5.19%) ***dup

South America =====

2/616= <1%

BRAZIL: 2 (0.32%)

North America =====

153/616= 25%

CANADA: 13 (2.11%)
UNITED-STATES: 140 (22.73%)

Middle East =====

4 countries
16/616= 3%

IRAN: 3 (0.49%)
ISRAEL: 11 (1.79%)
SAUDIARABIA: 1 (0.16%)
UNITED-ARAB-EMIRATES: 1 (0.16%)

Africa =====

2 countries
2/616= <1%

SOUTHAFRICA: 1 (0.16%)
TUNISIA: 1 (0.16%)

- CONFERENCE MANAGEMENT SYSTEM. As in previous ACL conferences, we used the START system to manage submissions and reviews. Things

generally went well and we got good support from Rich Gerber.

- BIDDING. One new component of the reviewing process for Coling/ACL was "bidding" for papers by reviewers. This had worked well for other conferences (AAAI, ICML, KDD, etc.). In general, it also worked well for us. In particular, reviewers generally received the papers that they preferred via bidding; they only needed to bid on papers for the areas that they were reviewers for; and START makes automatic assignments (which can be modified) based on the bids, so that, in general, it is easy for area chairs to assign papers to reviewers. We had hoped that bidding would give us greater flexibility in assigning reviewers to papers, e.g. for papers that span more than one area or if we wanted to share reviewers across areas. But START did not have a reasonable mechanism for doing this. Some reviewers (e.g. Peter Turney) were disappointed by that and expressed a strong preference for the way these matters were handled by CyberChairPRO at ICML and other conferences. Nevertheless, we were happy with the bidding process overall.
- REVIEWING. The reviewing process went very smoothly. As indicated in the schedule presented above, it included a discussion period for papers on which reviewers strongly disagreed. 100% of reviews were turned in on or nearly on the due date!
- ACCEPTANCE RATES AND STATS. Paper presentation: 23%; Poster presentation: 20%; Overall: 43%

N.B. (1) Since only 11% of the submissions were in the poster category, the majority of submissions that were accepted as posters had in fact been submitted in the paper category.

N.B. (2) The numbers above are in fact the ratios between the number of papers/posters accepted and the total number of submissions. The proportion of positive decisions from the PC was slightly higher, but some papers were pulled out, especially for paper submissions that were being redirected to the poster category.

For more details, see the .htm file.

- DOUBLE SUBMISSIONS. Double submissions were a bit of a problem. In spite of the wording in the call for proposals, some (30+?) authors failed to note parallel submissions. The current wording is the following:

"Papers that are being submitted in parallel to other conferences or workshops must indicate this on the title page, as must papers that contain significant overlap with previously published work."

We suggest that the wording be stronger next year. The new wording should probably make it clear that if the authors decide to submit to another conference only after they have completed their submission to the ACL conference, then they must update their submission info with the ACL.

There was a particular problem with EMNLP submissions; we needed confirmation on accepted Coling/ACL papers and posters BEFORE the EMNLP decision notification date. Authors of 8-10 EMNLP submissions

that we had accepted as posters, wanted to wait until the EMNLP decisions were made before committing to attend Coling/ACL for a poster presentation.

- INVITED SPEAKERS. With input from the area chairs, we decided on two invited speakers, one from within the Coling/ACL community and one from a field traditionally outside the purview of NLP/CL.
- ASIAN LANGUAGE EVENTS. In honor of the joint conference's location in Asia, we planned a special Asian language event that consists of the presentation of the top four Asian language papers in a parallel session and a plenary talk/panel focusing on issues in Asian language processing, followed by the presentation of the Best Asian Language Paper Award. The panel was organized by Aravind Joshi. The Asian language best paper was selected by the area chairs and program chairs.
- BEST REVIEWER AWARD. In order to acknowledge the crucial role from our reviewers, we were planning to grant a Best Reviewer Awards at the Closing Session. We asked our area chairs to designate some contenders for the title. Our plan was to randomly draw 5 names from the resulting list. However, we realized that some people were taking this matter too seriously ("this is stuff for beefing up one's CV") for the kind of lightweight informal process we had set up. This is why we decided to deemphasize this by dropping the name "Best Reviewer Award". Instead as part of the general "thanking" episode we will draw the names of 5 great reviewers.