

# **NAACL HLT 2010: Report of the General Chair**

Ronald M. Kaplan  
Microsoft Corporation

The 2010 NAACL HLT conference was held at the Millennium Biltmore Hotel in downtown Los Angeles. The hotel was an elegant venue that added a touch of grace and class to the intellectual activities and its ambience contributed to a positive social experience as well. The Local Arrangements Committee from ISI did an outstanding job of organizing all of the logistical components of the conference.

The main conference included a strong technical program of 146 long papers, short papers, and poster presentations. These were selected from 450 submissions, for an overall acceptance rate of 32%. The separate report of the Program Committee contains details of the submission and selection process. The program also included 12 demos (out of 20 submitted), presentations by two invited speakers, a Student Research Workshop, an Industry Panel lunch, and a few industry/publisher exhibits. The Program Committee chose “Noisy Genre” as a special theme for the conference and especially encouraged submissions on that topic. The main conference was preceded by a day of eight tutorials and followed by a weekend of 16 workshops. It was indeed a very full agenda.

One of the recommendations from the 2009 conference was to reduce the number of Program Committee co-Chairs from 4 to 3, recognizing that previous attempts to reach out to the Information Retrieval community had not born much fruit, presumably due to competition with conferences (like SIGIR) that focus directly on that topic. I thus selected two co-Chairs to represent different aspects of core NLP and one with a special focus on Speech. It was somewhat surprising, then, that we receive almost 40 submissions in that Information Retrieval and Extraction category, second only to the number that we received in Machine Translation. It may be worth revisiting our outreach to IR in future years.

This year there were several departures from past practices:

1. For the first time we did not create or distribute hardcopies of the proceedings. And for the first time the machine readable proceedings were created on USB sticks instead of CD-ROMs. This came at some additional cost, but recognizes the fact that CD drives are not available on new classes of machines (Macbook Air and Netbooks). I think this was generally well accepted (I have heard no complaints). We also put the full suite of papers up on the ACL Anthology to enable prereading in advance of the conference.
2. We included an “author response” period in the reviewing process, giving authors a few days to react to (and perhaps correct) the opinions expressed in the written reviews. The Program Committee report highlights a number of issues that arose in the way we were able to implement this form of

- interaction, and the NAACL board may want to formulate more specific policies if this is to be carried over into future conferences.
3. We did not have separate submission deadlines or reviewing phases for long and short papers.
  4. We included a “One-Minute Madness” plenary session immediately in front of the Poster/Demo session, as a way of promoting those contributions and perhaps to make those kinds of presentations more attractive.
  5. Unlike last year, we did not legislate a division of time in the Poster/Demo session, allowing individual presenters to decide how they wanted to allocate their time.
  6. We were the first conference to make use of the fact that the ACLPUB templates and machinery have been installed in the SoftConf/SMART conference management system, providing a more integrated set of tools for the Publications Committee. As pioneers, we confronted a number of glitches and non-features, but SoftConf addressed them fairly rapidly and this configuration should smooth out the publication process for future conferences.
  7. This year (and perhaps for the first time) there were no physical meetings of the various committees and subcommittees, all the interactions and discussions were handled by conference calls (subsidized by Microsoft) and electronic exchanges. This seemed to work well, and obviated the need for special travel arrangements in the budget (the only travel was my brief trip as General Chair to meet with the Local Arrangements Committee and to review the set up at the hotel.)
  8. The Demo Chair (Carolyn Penstein Rosé) made a concerted effort this year to solicit and review the demos. This resulted in a much more complete program than I think we saw in previous years, and the One-Minute-Madness made the demos seem like first class participants.

A large number of issues arose during the course of the planning and execution of the conference, mostly minor but some may bear consideration by the NAACL board (and ACL Executive) for future conferences:

1. Coordination with the ACL conference—Submission/notification deadlines  
The NAACL conference was the first of 3 NLP conferences this year, to be followed by the international ACL conference in Uppsala and Coling in Beijing. We were asked to arrange our submission and reviewing schedule so that authors would be notified of the NAACL decision in time for them to revise and resubmit (presumably rejected) papers by the time of the ACL submission deadline. This put undue pressure on our submission and reviewing timeline: our submission deadline was earlier than we would have liked, our review period overlapped with the Christmas holidays, we had to arrange our period for author responses in a less than optimal way. We also spent a fair amount of time negotiating with the ACL Program Committee to push them back from what would have been completely unworkable deadlines.

This raises the policy issue of how important it is to provide for the resubmission opportunity. It would be useful to gather data on this question: how many papers were rejected or withdrawn from NAACL and were then later submitted *and accepted* to ACL (with or without revision). It may be that we are created pressure on the NAACL submission and reviewing process for something that benefits the community in a very marginal way. This is an issue that perhaps the NAACL Board should discuss with the ACL Executive

2. Coordination with ACL and Coling—workshop selection

As noted in the separate report from the NAACL Workshop co-Chairs (Richard Sproat and David Traum), workshops proposals were submitted and to a joint committee representing the three conferences. The attempt to match venues preferred by the submitters with slots available for the different conferences led to extra work and confusion in the review process, and it is not clear that this joint submission and review process is a net positive. The Workshop report points to some alternative arrangements that might improve the process, and the NAACL Board and the ACL Executive may also want to reexamine this aspect of coordination.

(A side note on Workshops: the finance and policy documents are 15 years old, and probably should be updated. It should also be made clear which parts of the policy applies to regional as opposed to main-ACL conferences.)

3. Although an author response period may lead to better acceptance decisions (by allowing correction of true misunderstandings), there were deficiencies in our first experiment with this additional phase of quality control. As noted in the Program Committee report, the time constraints forced us to overlap the author responses with discussions that were still going on among the reviewers and the Area Chairs. In some cases the author response information was not distributed and the loops were not closed, and we did not get full benefit from this extra part of the procedure. Some authors also were surprised when the initial scores and comments (which they responded to) were followed by a lowering of scores (based on further reviewer discussion that took place contemporaneous to the response period). If author responses are to become a normal part of our review process, then the schedule should allow for a more natural progression of the phases, to avoid confusion and consternation. This is related to the deadline pressure coming from ACL coordination (point #1 above).

4. Short and long paper reviewing

This year (unlike last) we did not have separate submission and reviewing dates for long and short papers (partly because of the time pressure, and the desire not to hit up reviewers with a task that is spread out in time). It is not clear that lumping them together for review is a good thing. Reviewers may

have had trouble keeping track of which review criteria to apply, and that may account for the fact that short papers tended to be scored lower than long papers. It may be better to revert to the policy of previous years, wherein long and short papers were reviewed at different points in time.

5. Student Research Workshop

The SRW committee this year raised some of the same concerns that appeared in the report of last year's conference. There is confusion about the purpose and attraction of the student workshop compared to the main conference, and the SRW competes in time with the main conference sessions. The submission rate is fairly low, suggesting that advisors are not encouraging students to submit, and there has been difficulty in recruiting senior researchers to participate as discussants and consultants. The SRW committee this year repeats the request from last year's committee, that the NAACL Board reexamine the purpose and organization of this workshop, and perhaps consider other mechanisms for achieving its goals.

An NSF policy change also affected the allocation of funding for SRW participants this year, and that may also call for some reconsideration. We received essentially the same size grant this year as last, but NSF now restricts funding only to U.S. students or students from developing countries. International students from developed countries, say from Europe, Canada, or Japan, cannot be supported from this grant. If NAACL cannot provide a work-around, this fact should be clearly stated in any call for future versions of the SRW—and that may lead to even fewer submissions.

6. Financial support for student participation

Somewhat late in the game the NAACL Board was asked to define a policy for travel grants to students who were not part of the SRW, and in particular, to determine the criteria for assigning priority to applications. The policy gave priority to North American students, to first-author students, and to student volunteers. This policy should be reviewed for the next conference and should be publicized on the conference web page, together with the application procedure. If there are multiple ways in which student support can be provided, these should be listed in a single description with a single application for all types of support.

7. Publication issues

As noted, for the first time we provided no hardcopy proceedings. The USB sticks were distributed at the beginning of the conference, and the papers were put up on the ACL Anthology a week or so before the conference. Given the absence of hardcopy, at least some of the traditional reasons for limiting the lengths of written papers do not seem to apply. In the call this year we allowed for 8 pages of content and one extra page for references, and we stuck to that policy against some number of requests for more space. The workshop calls were in some instances less clear about page length, and

workshop paper-lengths were not necessarily consistent with the overall specification.

The NAACL Board may want to revisit the paper-length policy, in light of new technologies. And the committees for the next conference should be clear in specifying whatever policy is adopted. The conflict, as I see it, is that we cannot burden the reviewers with papers of any greater length, but but allowing substantially more material to be added in the final copy means that authors can included large amounts of unreviewed content.

A second publication issue concerns the time at which papers are uploaded to the ACL Anthology. As noted in our earlier email exchanges, publication in advance of dates that authors expected can have serious consequences for patentability (and hence can have real world consequences for such things as financial rewards and promotions). Since the default expectation is that papers submitted to a conference with confidential reviewing will not be publicly disclosed until the opening date of the conference, it is extremely important to announce very clearly any exception to that convention. Anthology uploading should be listed as one of the critical dates in each of the Call for Papers (main conference, workshops, demos—whatever will go up in advance of the conference start).

#### 8. Final Paper date and Best Paper judging

In hindsight, I think we allowed too much time for the authors to work between notification and final submission, and that put unnecessary pressure on the downstream processes, the work of the Publication Committee and the constitution and work of the Best Paper Committee. Those committees can only start to operate on the final submissions, and they come in at a critical phase of the end game where there isn't much leeway for error. For future conferences, I suggest moving the Final Paper submission date back a week, to remove some of the end-game pressure.

The Best Paper Committee was created this year in advance of actually having the final nominations from the Area Chairs, and it was recognized late in the process that some of the BP judges had conflicts of interest with some of the papers they were asked to judge. In future conferences, the nominations should be made known to potential judges before they are appointed, so that we do not end up with unexpected recusals.

#### 9. Local arrangement issues

Although the arrangements this year were very smooth and well done, there were some surprises towards the end that made it difficult to set the registration fees. These issues may not arise for conferences that are held on university campuses or in less glamorous hotels. The initial contract with the hotel provided very good room rates and other amenities, but it did not specify some important additional features, and these turned out to be

surprisingly expensive. In particular, internet access and audio visual arrangements were outrageously priced. Priscilla did a fantastic job of negotiating some work-arounds (also with the incredibly expensive (but very good) food), but we had to settle for not having wireless access in any of the meeting rooms. And all the back-and-forth did result in delays in opening up the conference registration site.

These extras should not be left out of future initial negotiations (and I'm sure Priscilla will keep track).

This list includes most of the issues that emerged from my experience this year and the experiences of the various committees. Each of the major organizational tasks takes a substantial amount of knowledge. Past organizers have tried to capture their experience and make recommendations in documents that can be found by rummaging around on the ACL web site, but those documents are not maintained and not necessarily consistent with current practices. Priscilla maintains the corporate memory for many of the logistic issues, but much of the knowledge is gained and then lost from year to year. It would be helpful, I think, if there could be overlapping terms for some of the major conference positions (PC, Publications, Workshops). I recognize that that would be very unlikely—who would ever want to volunteer for a two-year responsibility. A possible compromise, for consideration of the NAACL Board, is to ask one person to precommit to take an operational role on each of the key committees for the next year's conference but to act only as an observer of that role for the current year (participate in phone calls, read email, etc. but not be required to act or respond). That would be one way of passing knowledge from year to year.

But to summarize for NAACL 2010, I think (and I hope you agree) that it was quite a successful conference. About 700 fully-engaged participants, a full assemblage of high-quality papers, workshops, and presentations, and the program and local arrangements all came together to make for a great experience. This is because of the hard work and careful attention from all the committee chairs and committee members—I am indebted to all of them, as I have indicated elsewhere, for making this the memorable event that I think it was.

Ron Kaplan, Microsoft Corporation  
General Chair

### **Brief report from the Local Arrangements Committee**

NAACL HLT 2010 was held at the Millennium Biltmore Hotel in downtown Los Angeles, June 1-6, 2010, with the banquet at the nearby Maguire Gardens in front of the central library. The local arrangements committee was based at USC Information Sciences Institute. Nearly 700 people were in attendance.

The main program consisted of 85 long papers, 25 short papers, 34 posters, and 12 demos. This year, there was a special focus on research with noisy data, including data from informal communications (such as Twitter, blogs, e-mail, SMS) and processed data (such as speech, OCR, historical data, and machine translation). It incorporated for the first time a "one-minute madness" session to raise awareness and excitement about the poster and demo presentations. Also for the first time, the proceedings were distributed on a USB memory stick, unaccompanied by a hardcopy volume.

The conference website ([naaclhlt2010.isi.edu](http://naaclhlt2010.isi.edu)) has been viewed over 60,000 times from more than 100 countries. In keeping with the conference's theme, we also note that #naacl2010 was the subject of hundreds of tweets on Twitter.

David Chiang, Jason Riesa, Jonathan May, Ed Hovy (co-Chairs)  
University of Southern California Information Sciences Institute

## **Report of the Tutorial Committee**

The tutorial chairs actively recruited submissions, received 9, and accepted 8. All 8 tutorials received sufficient enrollment to proceed. The complete list of tutorials is given below (attendance figures are final on-line at the time of this report plus expected on-site registration):

T1: Computational psycholinguistics  
Roger Levy, Klinton Bicknell and Nathaniel Smith  
28 participants

T2: Integer Linear Programming in NLP - Constrained Conditional Models  
Ming-Wei Wang, Nicholas Rizzolo and Dan Roth  
40 participants

T3: Data-Intensive Text Processing with MapReduce  
Jimmy Lin and Chris Dyer  
30 participants

T4: Distributional Semantic Models  
Stefan Evert  
20 participants

T5: Markov Logic in Natural Language Processing: Theory, Algorithms, and Applications  
Hoifung Poon  
48 participants

T6: Noisy Text Analytics  
L. Venkata Subramaniam  
45 participants

T7: Recent Advances in Dependency Parsing  
Qin Iris Wang and Yue Zhang  
25 participants

T8: Textual Entailment  
Mark Sammons, Idan Szpektor and V.G.Vinod Vydiswaran  
20 participants

Jason Baldrige (co-Chair)  
The University of Texas at Austin

Peter Clark (co-Chair)  
The Boeing Company

Gokhan Tur (co-Chair)  
SRI International

## **Report of the Demonstration Committee**

20 demos were submitted. 12 were accepted. The demo chair specifically solicited demos from people who had papers at ACL, NAACL, or affiliated workshops last year that seemed like they included something that could be an interesting demo. Personal invitations were set to the first author of each such paper, with reference to the paper from last year that was the "trigger" for the invitation. We had 13 reviewers altogether, and every demo received at least 2 reviews. Reviews were checked before being sent out to make sure each demo received constructive feedback.

Carolyn Penstein Rosé (Chair)  
Carnegie Mellon University

## **Report of the Student Research Workshop/Doctoral Consortium**

The Student Research Workshop provided a venue for student researchers investigating topics in the broad field of Human Language Technologies to present their work and receive feedback from the community. This year we received a total of 20 submissions. Submissions were down from the 29 that were received last year, although last year the deadline was extended twice.

Of the 20 submissions, we accepted 8 of them as oral presentations (two of which were withdrawn from the workshop) and another 6 as poster presentations (one of which was withdrawn, and another of which will be presented by someone else as the author was unable to attend due to VISA issues). Final oral presentations and posters come from the following countries: United States (4), India (3), Germany (2), China – Hong Kong (1), Brazil (1). A total of 33 people (14 students and 19 postdocs/faculty/senior researchers) agreed to serve on the program committee, which allowed us to assign 3 reviewers per paper. 12 researchers who were largely different than the reviewers agreed to serve as oral presentation panelists (2 per presentation) during the workshop.

Both oral presentation and poster presentation sessions were scheduled during the main conference. The oral presentations were held in two sessions on the first day, separated by a lunch. The oral presentation slots were 30 minutes, with 20 minutes for presentation and 10 minutes for feedback from two panelists. Other conference participants attended the sessions as well. The posters were presented at the main conference's evening Poster and Demo session. Students with oral presentations were invited to participate in the poster session, and 5 out of the 6 accepted this invitation.

All eligible presenters received financial support from the U.S. National Science Foundation to assist them in their travel to Los Angeles for the conference. Altogether we received \$18,000 from the National Science Foundation to fund the workshop, which included support for student participants, student co-chairs, and the cost of the student lunch. \$16,400 was budgeted for travel funds and distributed as follows. Students studying in N. America were offered up to \$1000 to defray the cost of registration and hotel as well as travel from within North America. Students traveling from overseas were offered up to \$2500. This year NSF changed its policy, so we were only allowed to fund two categories of students (those studying at American institutions, or in developing countries). We were thus unable to fund our students from Germany and Hong Kong, and were restricted in what we could offer our Canadian student co-chair. Finally, \$1,400 was budgeted for the student lunch and any administrative costs.

There was a student lunch on the day of the SRW (consisting of the organizers, the student participants, and the panelists), providing social and networking opportunities. We also solicited feedback regarding how the SRW could be improved. An enormous amount of time was spent on this workshop by the organizers, and we would like to make sure that the payoff for the students justifies such an effort. Most of the issues noted in last year's report are still problematic this year. For example, a lack of submissions seems to be an ongoing problem, so perhaps students are not getting encouragement from their advisors to submit to the SRW, and/or there is confusion about the kind of papers appropriate for the workshop, which now seems to be somewhere between a workshop and a doctoral consortium. This year it also seemed difficult to get senior researchers to agree to

serve as both reviewers and panelists. We strongly recommend that the NAACL Board revisit these issues, which were discussed in more depth in last year's report.

Julia Hockenmaier (Faculty Advisor)  
University of Illinois, Urbana, Illinois, USA

Diane Litman (Faculty Advisor)  
University of Pittsburgh, Pittsburgh, Pennsylvania, USA

Adriane Boyd (NLP Co-Chair)  
Ohio State University, Columbus, Ohio, USA

Mahesh Joshi (NLP Co-Chair)  
Carnegie Mellon University, Pittsburgh, Pennsylvania, USA

Frank Rudzicz (Speech Co-Chair)  
University of Toronto, Toronto, Ontario, Canada

[Reports of the Program and Workshop Committees are in separate files]