

Preface from the Program Co-chairs

One morning each of us received a phone call from Ed Hovy. "Are you sitting down?" he asked. He told us that as a way to combat conference overload, and to promote interaction among communities, a joint conference had been proposed to combine HLT and NAACL. A diverse oversight committee had been formed, and according to Ed, this committee had been able to agree on two people – and only two people – as program co-chairs, because together we represented all of the vested interests. Marti was meant to represent the standards and tastes of the NAACL and the SIGIR crowds, and Mari the speech community, and both have been working on research contracts with HLT funders. Ed told us that if either of us said no, the entire enterprise would come crashing down. There are few better ways to convince busy people to become program co-chairs.

Throughout the process, Ed provided the vision for and the drive behind this conference. We salute him for making this idea a reality, and for his enthusiastic and energetic phone calls that kept everything going.

This is an exciting time for research in human language technologies. After years of relative calm, the field seems suddenly to be moving by leaps and bounds. Evidence of this can be found in our conference panel on "Preparing for a Surprise Language" (and as embodied in the short paper "Desperately Seeking Cebuano"). This panel will discuss the experiences of several groups of researchers, who at the behest of DARPA, acquired and developed language resources for an entirely new language within a span of only 10 days. This experiment took place in March of 2003, and the language in question was Cebuano, a language spoken in the Philippines. Participants successfully collected a large body of lexical and textual resources and developed a range of tools, including stemmers and POS taggers. (In June, DARPA will announce a new surprise language.)

The existence of a variety of language resources, combined with advances in statistical analysis and modeling techniques, is resulting in fast-paced improvements in the field. Statistical parsers can now produce syntax trees for long sentences with high accuracy and great speed. Advances are starting to be made in automated semantic analysis. Great strides are being made in the sophistication and coverage of question answering systems. Speech recognition systems have achieved sufficiently high accuracy that it is now possible to do retrieval, information extraction and topic tracking on spoken documents. Large and growing collections of text and speech corpora – and the promise of much more from the web – have enabled many of these advances. New developments in weakly supervised and unsupervised learning algorithms are critical for taking advantage of many new data sources, and hence this was chosen as a special theme of the conference. Lexical resources such as FrameNet, WordNet, PropBank, MeSH, and the Penn TreeBank also play prominent roles in HLT advances.

As a field, human language technologies research should use, as motivation and guide, an understanding of the linguistic and cognitive bases of language. The invited talk by Dr. Elissa Newport, entitled "Statistical language learning: Mechanisms for language acquisition in human learners," should help enlighten the community by informing us about the latest in psycholinguistic research.

We received 162 submissions for full papers, of which 37 were accepted, resulting in a highly competitive acceptance rate of 22%. For the short (late-breaking) papers track, we received 80 submissions, of which 41 were accepted (2 later withdrawn). Some of these will be presented as short talks, and others as posters. Seventeen demonstrations will be shown.

We thank everyone who submitted full and short papers, whether or not your submission was accepted. Your participation is critical to the moving forward of the field, and we hope the reviewer feedback was helpful.

We were fortunate to be able to accept 15 papers that addressed the conference theme of unsupervised and weakly supervised methods. We also encouraged papers that described techniques that cross over or combine NLP, speech and/or IR, and several of the papers demonstrate this kind of crossover.

The full paper reviewing was done using a two-tier system. First, two first-tier reviewers read every paper. Then a third reviewer, known as the meta-reviewer, wrote their own review. Finally, the meta-reviewer summarized these reviews and introduced additional comments. In some cases, the meta-reviewer instigated discussion among the first-tier reviewers to work out controversial issues. The meta-reviewers also attended the program committee meeting in which all the papers were discussed and acceptances were decided. For the short papers, each short paper received at least two reviews. Those papers whose reviewers disagreed, or which received middling scores, were subsequently reviewed by a member of the program committee and the program co-chairs. Paper submission and reviewing was done online using Marti's conference reviewing software (Conga), which she updated for this conference. Marti also maintained the conference website.

Many people have contributed to the exciting program at this meeting. We thank our large and able reviewer pools; 59 people reviewed full papers, and 39 reviewed short. Their careful and timely reviews made it possible to include both high quality and late breaking results. We owe a great debt to the members of the program committee, who we believe are responsible for the high quality of the submissions (their presence on the PC no doubt encouraging top-notch research to be submitted). Their excellent judgment helped lead to the compilation of a stimulating and forward-looking set of papers. The PC was also responsible for suggesting and unanimously approving, our exciting invited speaker, Dr. Newport. Finally, all members of the PC cheerfully agreed to a second round of reviewing for the short papers, even though we did not tell them about this additional duty when we recruited them. We are also grateful to Doug Oard for the brilliant suggestion for the panel, as it combines the goals of the HLT funders with the talents of NLP, speech, and IR researchers.

Because of our goal of integrating researchers from different communities, tutorials play a particularly important role at this conference. We are pleased to have outstanding lecturers covering a broad range of HLT topics, including both statistical and non-statistical work in natural language processing, topics in speech and information retrieval, and general methods for optimization and estimation. We thank James Allen, Jason Eisner, and Wayne Ward for their efforts in bringing this program together. In putting together the schedule, we had much help from Dekang Lin, who very ably dealt with last minute changes in scheduling and arrangements for all the different session needs. We also thank Dragomir Radev and Steve Abney for their work on the conference proceedings.

Most of our decisions had to be approved by a large and diverse committee representing NAACL, HLT, ISCA, and SIGIR, as well as different government sponsors of HLT research. We are pleased to report that despite this seemingly cumbersome situation, our questions were always resolved within one or two days, if not hours, and with great alacrity. We send our thanks to all the board members!

We did our best to merge the different goals of the communities, and we think we succeeded for the most part. We also worked hard to solicit involvement from all of the associated communities,

and we are especially pleased about the interdisciplinary nature of many of the papers. As the conference becomes better-known, we hope to see this trend continue. We feel that, although there is room for improvement, the HLT-NAACL merge was a success and hope to see it continue next year.

Mar(t)i

Marti Hearst and Mari Ostendorf

April 17, 2003