

Program committee

Dragomir Radev (co-chair)	University of Michigan
Chris Brew (co-chair)	Ohio State University
Robert Dale	Macquarie University
Graeme Hirst	University of Toronto
Eduard Hovy	University of Southern California and ISI
Jason Eisner	Johns Hopkins University
Andy Kehler	UC San Diego
Lillian Lee	Cornell University
Gina Levow	University of Chicago
Diane Litman	University of Pittsburgh
Chris Manning	Stanford University
James Martin	University of Colorado
Detmar Meurers	Ohio State University
Massimo Poesio	University of Essex
James Pustejovsky	Brandeis University
Ehud Reiter	University of Aberdeen
Philip Resnik	University of Maryland
Ellen Riloff	University of Utah
Matt Stone	Rutgers University
Rich Thomason	University of Michigan
Hans Uszkoreit	University of the Saarland and DFKI
Bonnie Webber	University of Edinburgh
Dekai Wu	Hong Kong UST

Introduction

How should we teach Natural Language Processing and/or Computational Linguistics? Throughout the forty year history of the organization, many ACL members have responded to this question, finding solutions appropriate to a huge range of different local contexts. The ensuing diversity in approach is in some ways extremely welcome, but could be isolating for the individual instructor. This workshop is a forum for sharing of useful tools and experience.

We work in an “interface discipline”. Most of our work is collaborative and multidisciplinary, and this has immediate consequences for training. If we are doing our job, we should be able to answer – or at least address – most of the following questions.

- Is there a common core of knowledge and capability that should be part of the training of all NLP/CL students, or are the differences in emphasis between the contributing disciplines so large that this is unrealistic?
- Does it work well to shoehorn NLP/CL into niches in linguistics, CS or EE, or is it preferable to set up new departments and cross-cutting research centres? If the latter, what is the strategy for making sure that our students have a good onward trajectory into more conventional academic contexts.
- Should we be advising linguists to learn algorithmic complexity, speech engineers to bone up on Montague semantics or computer scientists interested in language to learn phonetics? If so, how can we make palatable the inevitable extra work, if not, what other advice should we offer?
- Which subfields of NLP/CL do we judge to have the most important long term research and employment potential? If we know that we may be better placed to understand and respond to the educational needs of the students and researchers who will create the future of our subject .

We can expect the detailed answers to these questions to be site-specific, but we are sure that there are also general insights that can help us all to make our teaching relevant to and supportive of the long term goals of the CL/NLP community. These issues are noticed by many, but little discussed, and often solved individually on an ad-hoc basis. The point of this workshop is to facilitate sharing and discussion, and to act as a seed for a community of research and scholarship in the teaching of NLP.

This workshop was made possible, in part, by grant #0226408 from the National Science Foundation. We are also grateful for contributions of time and effort from many people, especially Mary Taffet, Lillian Lee, Jason Eisner, Robert Dale and the whole of the program committee. We are grateful to them (and you) for your willingness to spend time on developing an effective and relevant scholarship of teaching for NLP and CL.

Dragomir and Chris