Terry Koo

maestro@mit.edu
630-336-7922

Education

Doctoral Candidate, MIT CSAIL 6/2004–Now

Graduating in 6/2010 (planned). Focusing in natural language processing and machine learning, specifically dependency syntax and discriminative modeling.

BS and M.Eng Degrees, MIT 9/1999–6/2004

5-year course with simultaneous award of BS and M.Eng. M.Eng thesis on hidden-variable models for discriminative parse reranking.

Experience

JHU CLSP Workshop 6/2002–8/2002

The Johns Hopkins Center for Language and Speech Processing holds a summer NLP workshop, where professors, grad students, and undergrads collaborate on selected projects. My project concerned generation of English text from Czech tectogrammatic structure, as part of a translation system. Supervisor: Jan Hajič.

MIT Media Lab UROP 10/1999–6/2001

Worked with the Media Lab's Explanation Architecture group on the Image Maps project, which aimed to collect and codify "local history": the history of day to day life in local areas. A key phase in the project was the collection and organization of large amounts of participants' personal photographs. I designed and implemented a tool that allowed users to visually specify the location and angle of personal photographs on an interactive map. Supervisor: Erik Blankenship and his advisor Brian K. Smith.

Internship, Vanu Inc. 6/2001–8/2001

Created and implemented backend software for an adaptive music selection system. Participated in regular meetings with the customer to discuss project progress. Implemented file parsing/generation module for a state chart manipulator. Supervisor: Rick Poyner.

Internship, Vanu Inc. 1/2001

Improved and updated analysis tools developed in the previous summer. Worked in pairs as part of an extreme programming trial. Supervisor: Rick Poyner.

Internship, Vanu Inc. 6/2000–8/2000

Designed and implemented tools for debugging and benchmarking software radio applications. Supervisor: John Chapin.

Publications

X. Carreras, M. Collins, and T. Koo. TAG, Dynamic Programming, and the Perceptron for Efficient, Feature-rich Parsing. *Proceedings of CoNLL*, 2008. (Best paper award).

T. Koo, X. Carreras, and M. Collins. Simple Semi-supervised Dependency Parsing. *Proceedings of ACL*, 2008.

M. Collins, A. Globerson, T. Koo, X. Carreras, and P. Bartlett. Exponentiated Gradient Algorithms for Conditional Random Fields and Max-Margin Markov Networks. *Journal of Machine Learning Research*, 9(Aug):1775–1822, 2008.

T. Koo, A. Globerson, X. Carreras, and M. Collins. Structured Prediction Models via the Matrix-Tree Theorem. *Proceedings of EMNLP*, 2007.

A. Globerson, T. Koo, X. Carreras, and M. Collins. Exponentiated Gradient Algorithms for Log-Linear Structured Prediction. *Proceedings of ICML*, 2007.

T. Koo and M. Collins. Hidden-Variable Models for Discriminative Reranking. *Proceedings of EMNLP*, 2005.

M. Collins and T. Koo. Discriminative Reranking for Natural Language Parsing. *Computational Linguistics*, 31(1):25–69, 2005.

Awards

Best Paper Award, CoNLL 2008

3rd place, US Wushu Union Nationals, Straightsword 2004

3rd place, US Wushu Union Nationals, Spear 2004

NSF Graduate Research Fellowship 2003

Best Overall Final Project, 6.170 2001

Our "Gizmoball" implementation was retained and shown to students in later semesters as an example of a well-designed project.

Languages

In order from most to least expertise: C, C++, Java, Perl, Korean, Spanish, Mandarin.