

Christopher Manning

Thomas M. Siebel Professor in Machine Learning,
Professor of Linguistics and of Computer Science



[Natural Language Processing Group](#), [Linguistics](#) and [Computer Science](#), [Stanford University](#)

Bio

Christopher Manning is the inaugural Thomas M. Siebel Professor in Machine Learning in the Departments of Computer Science and Linguistics at Stanford University. His research goal is computers that can intelligently process, understand, and generate human language material. Manning is a leader in applying Deep Learning to Natural Language Processing, with well-known research on Tree Recursive Neural Networks, sentiment analysis, neural network dependency parsing, the GloVe model of word vectors, neural machine translation, and deep language understanding. He also focuses on computational linguistic approaches to parsing, robust textual inference and multilingual language processing, including being a principal developer of Stanford Dependencies and [Universal Dependencies](#). Manning has coauthored leading textbooks on statistical approaches to Natural Language Processing (NLP) (Manning and Schütze 1999) and information retrieval (Manning, Raghavan, and Schütze, 2008), as well as linguistic monographs on ergativity and complex predicates. He is an ACM Fellow, a AAAI Fellow, and an ACL Fellow, and a Past President of the ACL. Research of his has won ACL, Coling, EMNLP, and CHI Best Paper Awards. He has a B.A. (Hons) from The Australian National University and a Ph.D. from Stanford in 1994, and he held faculty positions at Carnegie Mellon University and the University of Sydney before returning to Stanford. He is the founder of the [Stanford NLP group \(@stanfordnlp\)](#) and manages development of the [Stanford CoreNLP](#) software.

Contact

M Dept of Computer Science, Gates Building 2A, 353 Serra Mall, Stanford CA
94305-9020, USA

E manning@cs.stanford.edu

T [@chrmanning](#)

W +1 (650) 723-7683

F +1 (650) 725-1449

R Gates 248

O by appointment

A Grayce Ujihara, Gates 215, gujihara@stanford.edu

Brief CV

- I'm Australian ("I come from a land of wide open spaces ...")
- BA (Hons) Australian National University 1989 (majors in mathematics, computer science and linguistics)

- PhD Stanford Linguistics 1994
- Asst Professor Carnegie Mellon University Computational Linguistics Program 1994–96
- Lecturer University of Sydney Dept of Linguistics 1996–99
- Asst Professor Stanford University Depts of Computer Science and Linguistics 1999–2006
- Assoc Professor Stanford University Depts of Linguistics and Computer Science 2006–2012
- Professor Stanford University Depts of Linguistics and Computer Science 2012–
- President of the Association for Computational Linguistics 2015

Papers

Here is [my publications list](#). However, I've become lazy, so for recent stuff, you're often more likely to find it on the [NLP Group publications page](#). Or things might be more up-to-date at [Google Scholar](#), [Semantic Scholar](#), or [Microsoft Academic Search](#).

Books

[Introduction to Information Retrieval](#), with Hinrich Schütze and Prabhakar Raghavan (Cambridge University Press, 2008). Manning and Schütze, [Foundations of Statistical Natural Language Processing](#) (MIT Press, 1999). [Complex Predicates and Information Spreading in LFG](#) (1999). [Ergativity: Argument Structure and Grammatical Relations](#) (1996).

Conferences and Talks

A few of my [talks](#) are available online.

In 2013, I was one of the program co-chairs for the then new [International Conference on Learning Representations](#) (see: [ICLR 2013](#)). The 2013 edition was a really good workshop-scale event. Since then, it's been growing in size exponentially.

In 2013, I helped organize the first [CVSC workshop](#). It was a really lively workshop. I also helped organize [a second Workshop on Continuous Vector Space Models and their Compositionality](#) at EACL 2014.

I helped organize a [Workshop on Interactive Language Learning, Visualization, and Interfaces](#) to be held at [ACL 2014](#), trying to build an interdisciplinary community interested in the intersection of NLP, HCI, and data visualization.

Students

I made a page listing all my [Ph.D. graduates](#). You can find all my students on the [Stanford NLP Group People](#) page.

Research Projects

The general area of my research is robust but linguistically sophisticated natural language understanding, and opportunities to use it in real-world domains. Particular current topics include deep learning for NLP, Universal Dependencies and dependency parsing, language learning through interaction, and reading comprehension.

My research at Stanford is currently supported by the NSF, DARPA, Bloomberg, Tencent, and Ford.

I am interested in new students, at or accepted to Stanford, wanting to work in the area of Natural Language Processing. To find out more about what I do, it's best to look at my [papers](#), or my [group research page](#).

- Unadmitted students: The above statement is directed at students who have already been admitted to Stanford. I don't do admissions. You need to apply to a program in the usual manner; for Linguistics see <http://www-linguistics.stanford.edu/graduate/admissions.shtml>, and for Computer Science, see <http://cs.stanford.edu/Admissions/>.
- PhD students in CS/Linguistics or allied fields: please contact me directly.
- Masters students: I often employ a couple of masters students. Most appealing are people with a background in NLP, and time to devote to an RAship. It helps your case to have done well in [CS 224N: NLP](#).
- Undergraduate students in CS/Linguistics or allied fields: please contact me directly.

Courses

Online videos! You can find complete videos for two courses on NLP that I co-taught online on YouTube.

- [Natural Language Processing](#) (a.k.a. the 2012 Coursera NLP-class) by Dan Jurafsky and Christopher Manning.
- [Lecture Collection | Natural Language Processing with Deep Learning](#) (a.k.a. CS224N Winter 2017) by Christopher Manning and Richard Socher.

If you don't have much background in AI, ML, or NLP, you should start with the first course. If you're ready to dive into the latest in deep learning for NLP, you should do the second. The second course assumes more mathematics prerequisites (multivariate calc, linear algebra).

In Fall 2016, I will teach (for the first time ever!) [Linguistics 278: Programming for linguists](#) (and any other digital humanities or text-oriented social science students who think it might be a good match). I'm trying to learn a little more Python beforehand.

I co-taught tutorials on [Deep Learning for NLP](#) at ACL 2012 with Yoshua Bengio and Richard Socher, and at NAACL 2013 with Richard Socher. Slides, references, and videos are available.

In 2012, I co-taught [a free online course on Natural Language Processing](#) on [Coursera](#) with [Dan Jurafsky](#). We haven't found the time to revise it and teach a second version, but you can watch all the videos by selecting "Preview Lectures", and all the slides are available [here](#).

In June 2011, I taught a tutorial [Natural Language Processing Tools for the Digital Humanities](#) at [Digital Humanities 2011](#) at Stanford.

Nearly every year, I teach [CS 276: Information Retrieval and Web Search](#), with [Pandu Nayak](#). Earlier versions of this course include two years of two-quarter sequences CS276A/B on information retrieval and text information classification and extraction, broadly construed ("IR++"): [Fall quarter course website](#). [Winter quarter course website](#). This course started in 2001. Early versions were also co-taught by me, [Prabhakar Raghavan](#), and [Hinrich Schütze](#).

Nearly every year, I teach [CS 224N / Ling 237. Natural Language Processing](#) -- Develops an in-depth understanding of both the algorithms available for the processing of linguistic information and the underlying computational properties of natural languages. Morphological, syntactic, and semantic processing from both a linguistic and an algorithmic perspective. Focus on modern quantitative techniques in NLP: using large corpora, statistical models for acquisition, disambiguation, and parsing. Examination and construction of representative systems. Prerequisites: 121/221 or Ling 138/238, and programming experience. Recommended: basic familiarity with logic and probability. 3 units. I've taught this course yearly since Spr 2000. Many [previous student projects](#) are available online.

In fall 2007 I taught [Ling 289: Quantitative and Probabilistic Explanation in Linguistics](#) MW 2:15–3:45 in 160–318. I previously taught it in winter 2002 (née [Ling 236](#)) and Winter 2005 (as [Ling 235](#)).

In the summer of 2007, I taught at the LSA Linguistic Institute: [Statistical Parsing](#) and [Computational Linguistics in Industry](#).

In fall 1999 and winter 2001, I taught [CS 121 Artificial Intelligence](#). The text book was S. Russell and P. Norvig, [Artificial Intelligence: A Modern Approach](#).

I ran the NLP Reading Group from 1999–2002. [The NLP Reading Group](#) is now student organized.

Other stuff

LaTeX: When I used to have more time (i.e., when I was a grad student), I used to spend some of it writing [\(La\)TeX macros](#). [Actually, that's a lie; I still spend some time doing it....]

We've got two sons: Joel and Casey. Here are my opinions on [books for kids](#).

<http://www.stanford.edu/~manning/>

Christopher Manning -- manning@cs.stanford.edu -- Last modified: Nov 19, 2016