

Bill MacCartney

San Francisco, CA · 415-746-0904 · wcmac@cs.stanford.edu

Overview

- Ph.D. candidate in Computer Science at Stanford University (graduation expected June 2009).
 - Research focus: statistical natural language processing (NLP) and computational semantics.
 - Strong background in artificial intelligence (AI), machine learning, and linguistics.
 - Extensive business experience: entrepreneurship, finance, technology development.
-

Consulting

While my Ph.D. research is my primary focus, when time permits I serve as a consultant and advisor to companies seeking to apply state-of-the-art NLP and machine learning techniques to everyday business problems. I see my role as bridging the gap between theoretical research and practical implementation. Recent projects have involved:

- Extracting consumer opinions from online product reviews.
 - Developing a natural-language interface to a travel reservation system.
 - Applying machine learning methods in an online math tutoring system.
-

Professional Experience

2006: *Intern, GOOGLE RESEARCH*, Mountain View, California

- Pursued research on automated question answering.

2002-2003: *Researcher, KNOWLEDGE SYSTEMS LAB*, Stanford University, Stanford, California

- Pursued research in distributed AI: strategies for query-answering over multiple knowledge bases.
- Developed algorithms for efficient collaborative reasoning among disparate KBs and automatic decomposition of large KBs into minimally-connected partitions.
- Developed test suite for empirical validation and analysis of our algorithms using real-world KBs.

1999-2000: *Cofounder, VP Production, SAYIT, INC.*, San Francisco, California

- Secured \$5.3M in venture funding, with Softbank as lead investor.
- Designed, built, and tested initial website in only four months, including highly-scaleable content-ranking mechanisms, browse and search caches, and an interactive animation builder.
- Website ranked in top 250 within 3 months of launch, with over 2 million unique visitors per month.
- Managed web production team. Developed efficient production process, including project queuing and prioritizing, technical standards, and project specs.
- Directly responsible for information architecture, UI design, and usability analysis. Developed multidimensional content ontology.

1996-1998: *Quant/Trader, D. E. SHAW & CO.*, New York, New York

- Held primary responsibility for trading a \$5 billion portfolio of Japanese bonds and derivatives: JGB switches, swap spreads, box trades, basis trades, listed vs. unlisted JGBs, dollar-yen basis swaps, TIBOR-LIBOR swaps, euroyen futures, Japanese municipal bonds, FX forwards and options.
- Generated and analyzed trade ideas within a relative-value (expectations arbitrage) framework.
- Utilized quantitative valuation techniques and financial modeling, including multi-factor yield curve models, no-arbitrage derivatives pricing models, and mean-reverting stochastic processes.
- Developed and maintained fixed-income analytics code in C, C++, Perl, awk, SQL, bash, Tcl.

1994-1995: *Researcher, Knowledge Engineer, CYCORP, INC.*, Austin, Texas

- Led effort to deploy Cyc AI technology in a multi-agent architecture to achieve distributed inferencing.
- Developed a C library implementing most Lisp functions to facilitate the porting of Cyc from Lisp to C.
- Ported C version of Cyc from Unix to Mac OS. Designed and implemented Mac OS GUI for Cyc in C++.
- Developed HTML interface to Cyc, using cgi scripts to construct dynamic views of Cyc knowledge base.
- Contributed knowledge to the Cyc knowledge base, using the CycL knowledge representation language.

1991-1994: *Head of Research & Development, TESTTAKERS, INC.*, Roslyn Heights, New York

- Held primary responsibility for research and curriculum development.
- Initiated program to develop and market a new LSAT course. Achieved profitability within one year.
- Developed and maintained relational database application that ran the business.
- Managed multiple course locations, with responsibility for dozens of instructors and hundreds of students.

Education

2003-2009: **STANFORD UNIVERSITY**, Stanford, California

- Ph.D. in Computer Science (expected June 2009).
- Member of Stanford Artificial Intelligence Lab (SAIL) and Natural Language Processing (NLP) Group.
- Research on machine learning for computational semantics and natural language inference.

1990-1991: **J. W. GOETHE UNIVERSITÄT**, Frankfurt, Germany

- Awarded Fulbright scholarship for graduate work in philosophy.
- Coursework and independent work with concentration in philosophy of language.
- Mastered German reading, writing, and speaking skills.

1986-1990: **PRINCETON UNIVERSITY**, Princeton, New Jersey

- B. A. in Philosophy, *summa cum laude*. Focus: philosophy of language, philosophy of mathematics.
- Coursework in computer science and mathematics, including logic, computability, abstract algebra.
- GPA: 4.0 in major, 3.9 overall. GRE: 800/800/800. LSAT: 180.

Awards

- Springer Best Paper Award at the 22nd International Conference on Computational Linguistics, 2008.
- Fulbright scholarship for advanced study in Germany, 1990-91.
- Dickinson Prize for best senior thesis in logic and epistemology, Princeton University, 1990.
- Phi Beta Kappa, Princeton University, 1990.

Professional Service

I've served on the program committee or as a reviewer for several NLP conferences:

- Reviewer, NAACL-HLT 2007
- Reviewer, EMNLP 2008
- Reviewer, NAACL-HLT 2009
- Program Committee, TextInfer 2009

Skills

- Java, C, C++, Python, Scala, Lisp, Perl, bash, SQL, MatLab, HTML, XML, ...
- Fluent in German, conversant in French, unintelligible in Spanish.
- Talented writer. Widely traveled. Massively parallel. Know where Elvis is.

Papers

An extended model of natural logic

Bill MacCartney and Christopher D. Manning

The Eight International Workshop on Computational Semantics (IWCS-8), Tilburg, Netherlands, January 2009

A phrase-based alignment model for natural language inference

Bill MacCartney, Michel Galley, and Christopher D. Manning

The Conference on Empirical Methods in Natural Language Processing (EMNLP-08), Honolulu, HI, October 2008

Modeling semantic containment and exclusion in natural language inference

Bill MacCartney and Christopher D. Manning

The 22nd International Conference on Computational Linguistics (Coling-08), Manchester, UK, August 2008
—received Springer Best Paper Award—

Natural logic for textual inference

Bill MacCartney and Christopher D. Manning

ACL Workshop on Textual Entailment and Paraphrasing, Prague, June 2007

Learning Alignments and Leveraging Natural Logic

Nathanael Chambers, Daniel Cer, Trond Grenager, David Hall, Chloe Kiddon, Bill MacCartney, Marie-Catherine de Marneffe, Daniel Ramage, Eric Yeh and Christopher D. Manning

ACL Workshop on Textual Entailment and Paraphrasing, Prague, June 2007

Aligning semantic graphs for textual inference and machine reading

Marie-Catherine de Marneffe, Trond Grenager, Bill MacCartney, Daniel Cer, Daniel Ramage, Chloé Kiddon, Christopher D. Manning

AAAI Spring Symposium, Stanford, 2007

Learning to recognize features of valid textual entailments

Bill MacCartney, Trond Grenager, Marie-Catherine de Marneffe, Daniel Cer, Christopher D. Manning

Proceedings of the Human Language Technology Conference of the North American Chapter of the Association for Computational Linguistics (HLT-NAACL 2006), New York, NY, 2006

Generating Typed Dependency Parses from Phrase Structure Parses

Marie-Catherine de Marneffe, Bill MacCartney, Christopher D. Manning

5th International Conference on Language Resources and Evaluation (LREC-06), Genoa, Italy, 2006

Learning to distinguish valid textual entailments

Marie-Catherine de Marneffe, Bill MacCartney, Trond Grenager, Daniel Cer, Anna Rafferty, and Christopher D. Manning

Second Pascal RTE Challenge Workshop, 2006

Robust Textual Inference using Diverse Knowledge Sources

Rajat Raina, Aria Haghighi, Christopher Cox, Jenny Finkel, Jeff Michels, Kristina Toutanova, Bill MacCartney, Marie-Catherine de Marneffe, Christopher D. Manning, Andrew Y. Ng

Proceedings of the First PASCAL Challenges Workshop, 2005

Solving Logic Puzzles: From Robust Processing to Precise Semantics

Iddo Lev, Bill MacCartney, Christopher D. Manning, Roger Levy

Proceedings of the ACL-04 Workshop on Text Meaning and Interpretation, July 2004

Practical Partition-Based Theorem Proving for Large Knowledge Bases

Bill MacCartney, Sheila A. McIlraith, Eyal Amir, Tomas Uribe

Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03), August 2003

The Cycic Friends Network: Getting Cyc agents to reason together

James Mayfield, Tim Finin, Rajkumar Narayanaswamy, Chetan Shah, Bill MacCartney, Keith Goolsbey

Proceedings of the ACM CIKM-95 Intelligent Information Agents Workshop, December 1995