

Yisong Yue

Contact Information	H. John Heinz III College Carnegie Mellon University Hamburg Hall 5000 Forbes Avenue Pittsburgh, PA 15213	<i>website:</i> www.yisongyue.com <i>email:</i> yisongyue@cmu.edu
Research Interests	Machine Learning, Structured Prediction, Online Algorithms, Information Retrieval, Learning to Rank	
Education	Cornell University	2005 - 2010
	Ph.D. in Computer Science Graduate Minor in Statistics Dissertation: New Learning Frameworks for Information Retrieval Thesis Committee: Thorsten Joachims (advisor), Robert Kleinberg, Christopher Burges, Ping Li, John Hopcroft	
	University of Illinois at Urbana-Champaign	2001 - 2005
	Bachelor of Science in Computer Science Graduated with Highest Honors (Summa Cum Laude)	
	Illinois Math and Science Academy	1998 - 2001
Honors and Awards	Microsoft Research Graduate Fellowship, 2008-2010 Google Student Award Winner, NYAS Machine Learning Symposium, 2009 Yahoo! Key Scientific Challenges Award, 2008 Outstanding TA Award, Cornell Department of Computer Science, 2006, 2007 Dean's List, UIUC College of Engineering, 2002, 2003, 2004, 2005 Finalist, Scientific Blogging University Writing Competition, 2009	
Research Appointments	Carnegie Mellon University	<i>September 2010 - Present</i>
	Position: Postdoctoral Researcher. Working on applying and developing new machine learning techniques to understanding rich datasets and designing intelligent information systems. Under supervision of Carlos Guestrin and Ramayya Krishnan.	
	Cornell University	<i>May 2006 - August 2010</i>
	Position: Research Assistant / Research Fellow. Developed machine learning approaches for learning and prediction problems in information retrieval. Developed new structured prediction and on-line learning problem formulations and algorithmic solutions. Under supervision of Thorsten Joachims and Robert Kleinberg.	
	Google, Inc.	<i>June 2009 - September 2009</i>
	Position: Search Quality Analyst Intern. Developed techniques for automatically evaluating and optimizing information retrieval systems using implicit user feedback such as click data. Analyzed presentation bias effects in user click behavior. Under supervision of Rajan Patel.	
	Microsoft Research	<i>May 2007 - August 2007</i>

Position: Research Intern. Explored stochastic optimization techniques for optimizing rank-based performance measures. Analyzed empirical optimality conditions of learning to rank algorithms. Under supervision of Christopher Burges.

Service	SVM-map Software	http://projects.yisongyue.com/svmap/ Software package for structural support vector machine training to optimize for mean average precision.
	SVM-div Software	http://projects.yisongyue.com/svmdiv/ Software package for structural support vector machine training for diversified retrieval.
	SVM-sle Software	http://projects.yisongyue.com/svmsle/ Software package for structural support vector machine training to classify sentiment with latent explanations.
	Reflections Projections 2004	http://www.acm.uiuc.edu/conference Co-organized Reflections Projections 2004. Reflections Projections is a student-run computing conference at UIUC, typically drawing in hundreds of students from the Midwest region.
	Illini Book Exchange	http://www.illinibookexchange.com Worked on development, management and marketing of Illini Book Exchange (IBX). IBX is a free book exchange website for UIUC students, first launched in December 2002. Since then, over \$2 million worth of transactions has taken place via IBX.

Teaching Experience

Tutorials

- “Practical Online Retrieval Evaluation.” co-taught with Filip Radlinski, SIGIR '11 Tutorial, Beijing, China, July 2011.
- “Learning to Rank.” co-taught with Filip Radlinski, NESCAI '08 Tutorial, Ithaca, NY, May 2008
- “An Introduction to Structured Output Learning Using Support Vector Machines.” Microsoft Research Web Learning Group, Redmond, WA, August 2007.

Lectures & Sections Taught

- Advanced Topics in Machine Learning (Spring 2010). Taught two lectures on structured prediction and online learning techniques for information retrieval.
- Introduction to Programming (Fall 2005 - Spring 2006). Taught weekly section for introductory course in computer programming.

Students & Projects Advised

- Christina Brandt and Jacob Bank (Spring - Summer 2010), “Dynamic Ranked Retrieval.” Co-advised with Thorsten Joachims.
- Aaron Sidford, Undergraduate Senior Thesis (Spring 2008), “Online Local Search Optimization of Search Engines using Clickthrough Data.” Co-advised with Thorsten Joachims.
- Artificial Intelligence & Robotics Practicum (Fall 2007). Advised fifteen undergraduate student projects, including prediction of sporting event outcomes, learning to play board games, automatic music generation, object tracking and navigation of wheeled robots, and learning to walk using six-legged hexapod robots.

Teaching Assistant Appointments

- **AI/Robotics Practicum** (Fall 2007). Taught by Hod Lipson
- **Networks** (Spring 2007). Taught by Jon Kleinberg and David Easley.

- **Empirical Methods in Machine Learning** (Fall 2006). Taught by Rich Caruana.
- **Introduction to Programming** (Spring 2006). Taught by K.-Y. Daisy Fan.
- **Introduction to Programming** (Fall 2005). Taught by K.-Y. Daisy Fan.

Reviewing

Journals

- Data Mining and Knowledge Discovery
- Information Processing & Management
- Information Retrieval
- Journal of Artificial Intelligence Research
- Neural Networks
- Transactions on Knowledge and Data Engineering
- Transactions on the Web

Conferences

- ACL 2012
- ACML 2011
- COLING 2010
- ECML/PKDD 2008
- KDD 2011
- ICML 2007, 2008, 2009, 2010, 2011, 2012
- NAACL-HLT 2012
- NIPS 2008, 2009, 2010, 2011
- SIGIR 2008, 2009, 2010
- SoCG 2010
- WSDM 2011, 2012
- WWW 2011, 2012

Book Chapters

- Introduction to Information Retrieval, Chapter 18, Matrix decompositions & latent semantic indexing

Invited Articles

T. Joachims, T. Hofmann, Y. Yue, C.-N. Yu. (2009). “Predicting Structured Objects with Support Vector Machines.” *Communications of the ACM (CACM)*, Research Highlight, 52(11), 97–104, November 2009.

Journal Papers

O. Chapelle, T. Joachims, F. Radlinski, Y. Yue. “Large Scale Validation and Analysis of Interleaved Search Evaluation.” *ACM Transactions on Information Systems (TOIS)*, (to appear).

Y. Yue, J. Broder, R. Kleinberg, T. Joachims. “The K -armed Dueling Bandits Problem.” *Journal of Computer and System Sciences (JCSS)*, Special Issue on Learning Theory, (to appear).

Conference Papers

Y. Yue, C. Guestrin. (2011). “Linear Submodular Bandits and their Application to Diversified Retrieval.” In Proceedings of Neural Information Processing Systems (NIPS), Granada, Spain, December 2011.

Y. Yue, T. Joachims. (2011). “Beat the Mean Bandit.” In Proceedings of the International Conference on Machine Learning (ICML), Bellevue, WA, June 2011.

C. Brandt, T. Joachims, Y. Yue, J. Bank. (2011). “Dynamic Ranked Retrieval.” In Proceedings of the ACM Conference on Web Search and Data Mining (WSDM), Hong Kong, China, February 2011.

A. Yessenalina, Y. Yue, C. Cardie. (2010). “Multi-level Structured Models for Document-level Sentiment Classification.” In Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP), Cambridge, MA, October 2010.

Y. Yue, Y. Gao, O. Chapelle, Y. Zhang, T. Joachims. (2010). “Learning More Powerful Test Statistics for Click-Based Retrieval Evaluation.” In Proceedings of the ACM Conference on Information Retrieval (SIGIR), Geneva, Switzerland, July 2010.

Y. Yue, R. Patel, H. Roehrig. (2010). “Beyond Position Bias: Examining Result Attractiveness as a Source of Presentation Bias in Clickthrough Data.” In Proceedings of the World Wide Web Conference (WWW), Raleigh, NC, USA, April 2010.

Y. Yue, J. Broder, R. Kleinberg, T. Joachims. (2009). “The K -armed Dueling Bandits Problem.” In Proceedings of the Conference on Learning Theory (COLT), Montreal, Quebec, Canada, June 2009.

Y. Yue, T. Joachims. (2009). “Interactively Optimizing Information Retrieval Systems as a Dueling Bandits Problem.” In Proceedings of the International Conference on Machine Learning (ICML), Montreal, Quebec, Canada, June 2009.

Y. Yue, T. Joachims. (2008). “Predicting Diverse Subsets Using Structural SVMs.” In Proceedings of the International Conference on Machine Learning (ICML), Helsinki, Finland, July 2008.

Y. Yue, T. Finley, F. Radlinski, T. Joachims. (2007). “A Support Vector Method for Optimizing Average Precision.” In Proceedings of the ACM Conference on Information Retrieval (SIGIR), Amsterdam, The Netherlands, July 2007.

Workshop Papers

Y. Yue. (2009). “Online Gradient Descent using Interactive User Feedback.” NIPS Workshop on Analysis and Design of Algorithms for Interactive Machine Learning, Whistler, Canada, December 2009.

Y. Yue, T. Joachims. (2008). “Interactively Optimizing Information Systems as a Dueling Bandits Problem.” NIPS Workshop on Beyond Search: Computational Intelligence for the Web, Whistler, Canada, December 2008.

Y. Yue, C. Burges. (2007). “On Using Stochastic Perturbation Stochastic Approximation for Learning to Rank; and, the Empirical Optimality of LambdaRank.” NIPS Workshop on Machine Learning for the Web, Whistler, Canada, December 2007.

Additional Presentations

“An Interactive Learning Approach to Optimizing Information Retrieval Systems.”

- Carnegie Mellon University, Pittsburgh, PA, September 2010
- Yahoo! Research, Santa Clara, CA, August 2010
- Google Zürich, Zürich, Switzerland, July 2010
- Microsoft Research Asia, Beijing, China, June 2010

“New Learning Frameworks for Information Retrieval.”

- Microsoft Research, Redmond, WA, March 2010
- Google, Mountain View, CA, March 2010
- Johns Hopkins University, Baltimore, MD, March 2010
- Yahoo! Research, Sunnyvale, CA, February 2010
- Carnegie Mellon University, Pittsburgh, PA, February 2010
- Cornell University, Ithaca, NY, February 2010
- IBM TJ Watson, Hawthorne, NY, December 2009

“Interactively Optimizing Information Retrieval Systems as a Dueling Bandits Problem.”

- New York Academy of Sciences Machine Learning Symposium, New York, NY, November 2009
- Cornell University, Ithaca, NY, February 2009

“Diversified Retrieval as Structured Prediction.”

- Google, Mountain View, CA, August 2009
- SIGIR 2009 Workshop on Redundancy, Diversity, and Interdocument Relevance, Boston, MA, July 2009
- Cornell University, Ithaca, NY, April 2008

“Towards Interactive Approaches to Learning to Rank.”

- SIGIR 2009 Workshop on Learning to Rank, Boston, MA, July 2009

“Information Retrieval as Structured Prediction.”

- University of Massachusetts Amherst, Amherst, MA, April 2009
- Microsoft Research Asia, Beijing, China, August 2008

“A Support Vector Method for Optimizing Average Precision.”

- Microsoft Research Machine Learning Discussion Group, Redmond, WA, July 2007
- Cornell University, Ithaca, NY, April 2007

References**Carlos Guestrin**

Finmeccanica Associate Professor
Machine Learning Department
Carnegie Mellon University
Pittsburgh, PA 15213 USA
Email: guestrin@cs.cmu.edu

Thorsten Joachims

Associate Professor
Department of Computer Science
Cornell University
Ithaca, NY 14853 USA
Email: tj@cs.cornell.edu

Robert D. Kleinberg

Assistant Professor
Department of Computer Science
Cornell University
Ithaca, NY 14853 USA
Email: rdk@cs.cornell.edu

Christopher J. C. Burges

Principal Researcher
Microsoft Research
One Microsoft Way
Redmond, WA 98052 USA
Email: cburges@microsoft.com