

Yu-Han Chang

USC Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292

ychang@isi.edu
(617) 678-2486

Research Interests

My research centers on learning in rich multi-agent environments. It draws on ideas from machine learning, game theory, and cognitive science. Learning algorithms based on mathematical models for agent behavior, knowledge representation, and appropriate metrics are important parts of this work. I am interested in fundamental theory as well as application domains such as interactive games and intelligent tutoring systems.

Education

Ph.D., Electrical Engineering and Computer Science, 2005
Massachusetts Institute of Technology
NSF Graduate Research Fellowship 1999-2003
Dissertation: Approaches to Multiagent Learning
Advisor: Professor Leslie Pack Kaelbling

S.M., Computer Science, Harvard University 1999

A.B., Mathematics and Economics, *magna cum laude*, 1999
Harvard University
Senior Thesis: Exogenous Uncertainty in Computable
Generable Equilibria Models
Advisor: Professor Dale Jorgenson

Experience

Department of Computer Science, 2008-present
University of Southern California
Research Assistant Professor

Information Sciences Institute, Intelligent Systems Division 2005-present
University of Southern California
Computer Scientist

Computer Science and Artificial Intelligence Laboratory, 1999-2005
Massachusetts Institute of Technology
Graduate Research Assistant

Program on Technology and Economic Policy, 1997-1999
Harvard University Kennedy School of Government
Undergraduate Research Assistant

Publications

Journal Articles

1. No Regrets about No Regret

Yu-Han Chang. *Artificial Intelligence*, Volume 171, Issue 7 (2007).

Refereed Conferences

2. Learning and Playing in Wubble World

Wesley Kerr, Paul Cohen, and Yu-Han Chang. *Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 2008.

3. Learning in Wubble World

Wesley Kerr, Shane Hoversten, Daniel Hewlett, Paul R. Cohen and Yu-Han Chang. *International Conference on Development and Learning (ICDL)*, 2007.

4. Language Learning in Children's Games.

Wesley Kerr, Daniel Hewlett, Yu-Han Chang, and Paul Cohen. *Artificial Intelligence and Interactive Digital Entertainment (AIIDE)*, 2007.

5. Learning and Transferring Image Schemas

Yu-Han Chang, Paul Cohen, and Carole Beal. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2007.

6. Piagetian Adaptation Meets Image Schemas

Yu-Han Chang, Paul R. Cohen, Clayton T. Morrison, Robert St. Amant, and Carole Beal. *International Conference on the Simulation of Adaptive Behavior (SAB)*, 2006.

7. On Constructive Network Coding for Multiple Unicasts

Tracey C. Ho, Yu-Han Chang, and Keesook J. Han. *Allerton Conference on Communication, Control, and Computing*, 2006.

8. The Jean System: Transfer Learning and Experimental State Splitting

Yu-Han Chang, Clayton T. Morrison, Wesley Kerr, Paul R. Cohen, and Robert St. Amant. *International Conference on Development & Learning (ICDL)*, 2006.

9. An Image Schema Language

Robert St. Amant, Clayton T. Morrison, Yu-Han Chang, Wei Mu, Paul R. Cohen, and Carole Beal. *International Conference on Cognitive Modeling (ICCM)* 2006.

10. Hedged learning: Regret minimization with learning experts

Yu-Han Chang and Leslie Pack Kaelbling. *International Conference on Machine Learning (ICML)*, 2005.

11. Network monitoring in multicast networks using network coding

Tracey Ho, Ben Leong, Yu-Han Chang, Yonggan Wen, and Ralf Koetter. *International Symposium on Information Theory (ISIT)*, 2005.

12. A reinforcement learning approach to mobilized ad-hoc networks

Yu-Han Chang, Tracey Ho, and Leslie Pack Kaelbling. *International Conference on Autonomic Computing (ICAC)*, 2004.

13. All learning is local: Multi-agent learning in global reward games

Yu-Han Chang, Tracey Ho, and Leslie Pack Kaelbling. *Advances in Neural Information Processing Systems (NIPS)*, 2003.

14. Text Bundling: Statistics-Based Data Reduction

Kai Shih, Jason Rennie, Yu-Han Chang, and David Karger. *International Conference on Machine Learning (ICML)*, 2003.

15. Playing is believing: The role of beliefs in multi-agent learning

Yu-Han Chang and Leslie Pack Kaelbling. *Advances in Neural Information Processing Systems (NIPS)*, 2001.

Refereed Workshops

16. Transfer learning in real-time strategy games

Clayton T. Morrison, Yu-Han Chang, Paul R. Cohen. *International Conference on Machine Learning (ICML) Transfer Learning Workshop*, 2006.

17. A reinforcement learning approach to mobilized ad-hoc networks

Yu-Han Chang, Tracey Ho, and Leslie Pack Kaelbling. *AAAI Fall Symposium on Artificial Multi-agent Learning*, October 2004.

18. On the utility of network coding in dynamic environments

Tracey Ho, Ben Leong, Yu-Han Chang, Muriel Medard, Ralf Koetter, and Michelle Effros. *International Workshop on Wireless Ad-hoc Networks*, June 2004.

19. Not Too Hot, Not Too Cold: The Bundled-SVM is Just Right

Kai Shih, Yu-Han Chang, Jason Rennie, and David Karger. *International Conference on Machine Learning (ICML) Text Workshop*, 2003.

Invited Presentations and Other Papers

20. Probabilistic Fluent Graphs: Planning in Continuous Time and Space

Yu-Han Chang, Paul Cohen. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2009, to be submitted.

21. Finding Episodic Structure in Multivariate Sensor Data

Wesley Kerr, Yu-Han Chang, and Paul Cohen. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2009, to be submitted.

22. Learning Policies for Intelligent Tutoring Systems

Yu-Han Chang, Shane Hoversten, and Paul Cohen. *International Joint Conference on Artificial Intelligence (IJCAI)*, 2009, to be submitted.

23. Toddler Machine Meets Preteen Children: Concepts and language from combining lots of computation with lots of free time

Invited Talk, National University of Singapore Computer Science Seminar, August 2007.

24. Evaluating Transfer: Methods and Metrics For Analyzing the Performance of Learning Agents Across Domains

Invited Talk, AAAI 2007 Workshop Evaluation Methods for Machine Learning II, July 2007.

25. From Games to Gambling: Multi-Agent Learning with No Regrets.

Invited Talk, Pomona and Harvey-Mudd Computer Science Colloquium, November 2006.

26. Beyond Equilibria: Reinforcement and Hedging in Multi-agent Learning

Invited Talk, Workshop on Learning and Information in Games and Control, Caltech Center for the Mathematics of Information, March 2006.

27. Reinforcement Learning and Game Theory in Mobilized Ad-hoc Networks

Invited Talk, NIPS Workshop on Multi-agent Learning, 2002.

28. Mobilized ad-hoc networks: A reinforcement learning approach

Yu-Han Chang, Tracey Ho, and Leslie Pack Kaelbling. *AI Lab Memo*, AIM-2003-025, 2003.

29. Approaches to multi-agent learning

Yu-Han Chang. MIT Ph.D. Thesis, May 2005.

Grants and Contracts

Co-PI, DARPA Deep Green: Crystal Ball component, \$1.3M. Interactive battlefield planning. 2008-2011.

Co-PI, DARPA Battlesmarts seedling, \$300K. Life-long intelligent tutoring systems. 2008-2009.

Teaching

Teaching Assistant, MIT 6.042 Discrete Mathematics, Fall 2001

Teaching Assistant, Harvard Stats 110, Introduction to Probability and Statistics, Spring 1998

Honors and Awards

National Science Foundation Graduate Research Fellowship	1999-2004
John Harvard Scholarship	1997
United States Mathematical Olympiad (USAMO) participant	1995
5 th Place, Massachusetts Mathematical Olympiad	1994

Professional Activities Program Committee for *NIPS* 2003, *NIPS* 2004, *NIPS* 2005
Program Committee for *IJCAI* 2003
Program Committee for *AAAI* 2007
Program Committee for *ICML* 2008, *ICML* 2009
Referee for *Artificial Intelligence*
Referee for *Neural Computation*
Referee for *Machine Learning Journal*
Referee for *Journal of Machine Learning Research*
Referee for *IEEE Transactions on Automatic Control*
Referee for *Computational Intelligence*
Referee for *ACM Transaction on Autonomous and Adaptive Systems*, 2008.
Peer Reviewer for Air Force Office of Scientific Research, 2008.

Professional Memberships AAAI

Other Activities President, MIT Edgerton House Executive Committee 2001-2003
President, Harvard University Taiwanese Cultural Society 1998-1999

Personal Information Citizenship: United States
Hobbies: Skiing, hiking, architecture.