Ben Taskar

Address	Computer and Information Science University of Pennsylvania 3330 Walnut Street Philadelphia, PA 19104	215-898-7703 (voice) 215-573-8190 (fax) taskar@cis.upenn.edu www.cis.upenn.edu/~taskar

Education

Stanford University

Ph.D. in Computer Science, January 2005, Advisor: Daphne Koller
Thesis: *Learning Structured Prediction Models: A Large Margin Approach*

M.S. in Computer Science, Honors Cooperative Program, June 2000

B.S. in Computer Science with Distinction, June 1998

Academic Positions

University of Pennsylvania

Assistant Professor, Computer and Information Science, 2007–

University of California, Berkeley

Postdoctoral Fellow, January 2005–August 2006, Advisor: Michael I. Jordan

Research Interests

Machine learning, probabilistic graphical models, structured prediction, weakly supervised learning, natural language processing, computer vision.

Honors and Awards

Office of Naval Research Young Investigator Award (2010-13)
Sloan Research Fellowship (2010–12)
Microsoft New Faculty Fellowship Finalist (2010)
DARPA Computer Science Study Group (2009–)
Magerman Chair Assistant Professorship, University of Pennsylvania (2007–)
Arthur Samuel Best Thesis Award Runner-up, Stanford University (2005)
Best paper award, Empirical Methods in Natural Language Processing (2004)
Best student paper award, Neural Information Processing Systems (2003)
School of Engineering Fellowship, Stanford (2000)

Teaching

CIS 520 – Machine Learning – Fall10, Fall09, Fall08
CIS 620 – Probabilistic Graphical Models – Spr07, Spr09
CIS 700 – Monte Carlo Methods & Nonparametric Bayesian Models – Spr08
CIS 521 – Fundamentals of Artificial Intelligence – Fall07

Ph.D. Theses Supervised

Timothee Cour, *Weakly Supervised Learning from Multiple Modalities*, graduated June 2009, now a postdoc at INRIA, France

Kuzman Ganchev, *Posterior Regularization in Probabilistic Latent Variable Models*, expected graduation August 2010, co-advised with Fernando Pereira.

Benjamin Sapp, expected graduation 2012

Jennifer Gillenwater, expected graduation 2013

David Weiss, expected graduation 2013

Professional Service

Department Service: Organizer of Undergraduate Summer Research Program (`http://www.cis.upenn.edu/ugrad/research/`)

School Service: Member of MKSE Faculty Search Committee

University Service: Co-director of PRiML (Penn Research in Machine Learning, `http://priml.upenn.edu/`), Member of Weigle Information Commons Faculty Advisory Group.

Journal Editorial Board: Action Editor for Journal of Machine Learning Research, Editorial Board of Journal of Artificial Intelligence Research.

Conference Area Chair: European Conference on Computer Vision (ECCV2010), Neural Information Processing Systems (NIPS09,NIPS10), Empirical Methods in Natural Language Processing (EMNLP08).

Senior Program Committee: Uncertainty in Artificial Intelligence (UAI08 and UAI09), International Joint Conference on Artificial Intelligence (IJCAI09), Association for the Advancement of Artificial Intelligence (AAAI2008).

Conference Organization: Tutorial Chair for International Conference on Machine Learning (ICML), 2010.

Workshop Organization: *Kernels and Graphical Models* at Neural Information Processing Systems (NIPS04), *Learning in Structured Output Spaces*, at International Conference on Machine Learning (ICML06).

Journal Reviewing: Journal of Machine Learning Research and Journal of Artificial Intelligence Research, Journal of Computer and System Sciences, Machine Learning Journal, IEEE Transactions on Pattern Analysis and Machine Intelligence.

Conference Reviewing: Association for the Advancement of Artificial Intelligence (AAAI), Artificial Intelligence and Statistics (AISTATS), Empirical Methods in Natural Language Processing (EMNLP), Human Language Technology conference – North American Association for Computational Linguistics (HLT-NAACL), International Conference on Machine Learning (ICML), International Joint Conference on Artificial Intelligence (IJCAI), International Conference on Intelligent Robots and Systems (IROS), Neural Information Processing Systems (NIPS), Uncertainty in Artificial Intelligence (UAI).

Grant Proposal Reviewing: National Science Foundation (NSF), Israeli Science Foundation (ISF)

Funding Agency Service: Member of DARPA Computer Science Study Group, Participant of DARPA 2009 ISAT Study on Machine Learning on Multicores.

Media Publicity

Structured Prediction: A Framework for Machine Learning, Penn Engineering News, Spring 2010
`http://www.seas.upenn.edu/media/feature-structured-prediction.php`

Teaching computers how to learn, Penn Current, December 3, 2009,
`http://www.upenn.edu/pennnews/current/research/111209.html`.

Dans une vidéo, les dialogues sont un mode de référencement comme un autre, www.atelier.fr, October 12, 2009,
`http://www.atelier.fr/inventions/10/12102009/`

Machine Learning by Watching and Listening, PhysOrg.com, October 5, 2009, `http://www.physorg.com/news173980519.html`.

PUBLICATIONS *(Citation counts from: Google Scholar, retrieved on 03/22/2010)*

Books

2. *Predicting Structured Data*, Gökhan BakIr, Thomas Hofmann, Bernhard Schölkopf, Alex Smola, **Ben Taskar**, and S.V.N. Vishanathan, MIT Press, 2007. 360 pages.
1. *Introduction to Statistical Relational Learning*, Lise Getoor and **Ben Taskar**, MIT Press, 2007. 580 pages, **cited by 162**.

Journal Articles

5. *Posterior Regularization in Latent Variable Models*, Kuzman Ganchev, Joao Graca, Jennifer Gillenwater and **Ben Taskar**, Journal of Machine Learning Research, accepted March 2010, 51 pages.
4. *Learning Tractable Word Alignment Models with Complex Constraints*, Joao Graca, Kuzman Ganchev, and **Ben Taskar**, Journal of Computational Linguistics, accepted March 2010, 26 pages.
3. *Joint Covariate Selection and Joint Subspace Selection for Multiple Classification Problems*, Guillaume Obozinski, **Ben Taskar**, and Michael I. Jordan. Journal of Statistics and Computing, Published online January 2009, 22 pages.
2. *Structured Prediction, Dual Extragradient and Bregman Projections*, **Ben Taskar**, Simon Lacoste-Julien, and Michael I. Jordan, Journal of Machine Learning Research, Volume: 7, Pages: 1627-1653, 2006.
1. *Learning Probabilistic Models of Link Structure*, Lise Getoor, Nir Friedman, Daphne Koller and **Ben Taskar**, Journal of Machine Learning Research, Volume 3, Pages: 679-707, 2002, **cited by 128**.

Peer-reviewed Full Conference Papers

36. *Cascaded Models for Articulated Pose Estimation*, Ben Sapp, Alexander Toshev and **Ben Taskar**, 14 pages, European Conference on Computer Vision (ECCV), Crete, Greece, September 2010, **acceptance rate (oral) 3.4%**.
35. *Sparsity in Dependency Grammar Induction*, Jennifer Gillenwater, Kuzman Ganchev, Joao Graca, and **Ben Taskar**. 6 pages, Association for Computational Linguistics (ACL), Uppsala, Sweden, July 2010, **acceptance rate 21%**.
34. *Adaptive Pose Priors for Pictorial Structures*, Ben Sapp and Chris Jordan and **Ben Taskar**. 8 pages, Computer Vision and Pattern Recognition (CVPR), San Francisco, California, June 2010, **acceptance rate 22.3%**.
33. *Object Detection via Boundary Structure Segmentation*, Alexander Toshev and Kostas Daniilidis and **Ben Taskar**. 8 pages, Computer Vision and Pattern Recognition (CVPR), San Francisco, California, June 2010, **acceptance rate 22.3%**.

32. *Talking Pictures: Temporal Grouping and Dialog-Supervised Person Recognition,* Timothee Cour, Ben Sapp, Akash Nagle and **Ben Taskar**. 8 pages, Computer Vision and Pattern Recognition (CVPR), San Francisco, California, June 2010, **acceptance rate 22.3%**.

31. *Parsing Architecture at City-Scale from Range Data*, Alexander Toshev and Philippos Mordohai and **Ben Taskar**. 8 pages, Computer Vision and Pattern Recognition (CVPR), San Francisco, California, June 2010, **acceptance rate 22.3%**.

30. *Structured Prediction Cascades*, David Weiss and **Ben Taskar**, 8 pages, Artificial Intelligence and Statistics (AISTATS), Sardinia, Italy, May 2010, **acceptance rate 40.6%**.

29. *Posterior vs. Parameter Sparsity in Latent Variable Models*, Kuzman Ganchev, Joao Graca, Jennifer Gillenwater and **Ben Taskar**. 8 pages, Neural Information Processing Systems (NIPS), Vancouver, Canada, December 2009, **acceptance rate (spotlight) 7.9%**.

28. *Dependency Grammar Induction via Bitext Projection Constraints*, Kuzman Ganchev, Jennifer Gillenwater and **Ben Taskar**. 8 pages, Association for Computational Linguistics (ACL), Singapore, August 2009, **acceptance rate 21%**.

27. *Learning from Ambiguously Labeled Images*, Timothee Cour, Ben Sapp, Chris Jordan, and **Ben Taskar**. 8 pages, Computer Vision and Pattern Recognition (CVPR), June 2009, **acceptance rate 22.0%**.

26. *Learning Sparse Markov Network Structure via Ensemble-of-Trees Models.* Yuanqing Lin, Shenghuo Zhu, Daniel Lee, **Ben Taskar**. 8 pages, Artificial Intelligence and Statistics (AISTATS), Clearwater Beach, Florida, April 2009, **acceptance rate 30.0%**.

25. *Movie/Script: Alignment and Parsing of Video and Text Transcription.* Timothee Cour, Chris Jordan, Eleni Miltsakaki and **Ben Taskar**. 14 pages, European Conference on Computer Vision (ECCV), Marseille, France, October 2008, **acceptance rate 27.9%**.

24. *Multi-View Learning over Structured and Non-Identical Outputs.* Kuzman Ganchev, Joao Graca , John Blitzer and **Ben Taskar**. 8 pages, Uncertainty in Artificial Intelligence (UAI), Helsinki, Finland, July 2008, **acceptance rate (oral) 13%**.

23. *Better Alignments = Better Translations?* Kuzman Ganchev, Joao Graca and **Ben Taskar**. 8 pages, Association for Computational Linguistics (ACL), Columbus, Ohio, June 2008, **acceptance rate 25%**.

22. *Online, Self-supervised Terrain Classification via Discriminatively Trained Submodular Markov Random Fields*, Paul Vernaza, **Ben Taskar** and Daniel D. Lee. 8 pages, International Conference on Robotics and Automation (ICRA), Pasadena, California, May 2008, **acceptance rate 44.7%**.

21. *Expectation Maximization and Posterior Constraints*, Joao Graca, Kuzman Ganchev and **Ben Taskar**. 8 pages, Neural Information Processing Systems (NIPS), Vancouver, Canada, December 2007, **acceptance rate (spotlight) 10.4%**.

20. *Mixture-of-Parents Maximum Entropy Markov Models*, David Rosenberg, Dan Klein, **Ben Taskar**. 8 pages, Uncertainty in Artificial Intelligence (UAI), Vancouver, Canada, July 2007, **acceptance rate 32%**.

19. *A Permutation-Augmented Sampler for Dirichlet Process Mixture Models*, Percy Liang, Michael I. Jordan, **Ben Taskar**. 8 pages, International Conference on Machine Learning (ICML), Corvallis, Oregon, June 2007, **acceptance rate 29%**.

18. *An End-to-End Discriminative Approach to Machine Translation*, Percy Liang, Alexandre Bouchard-Cote, Dan Klein and **Ben Taskar**. 8 pages, Association for Computational Linguistics (ACL), Sydney, Australia, July 2006, **cited by 72**, **acceptance rate 23%**.

17. *Word Alignment via Quadratic Assignment*, Simon Lacoste-Julien, **Ben Taskar**, Dan Klein and Michael I. Jordan, 8 pages, Human Language Technology conference - North American chapter of the Association for Computational Linguistics (HLT-NAACL), New York, June 2006, **acceptance rate 24%**.

16. *Alignment by Agreement*, Percy Liang, **Ben Taskar** and Dan Klein, 8 pages, Human Language Technology conference - North American chapter of the Association for Computational Linguistics (HLT-NAACL), New York, June 2006, **cited by 70**, **acceptance rate 24%**.

15. *Structured Prediction via the Extragradient Method*, **Ben Taskar**, Simon Lacoste-Julien, and Michael I. Jordan, 8 pages, Neural Information Processing Systems (NIPS), Vancouver, Canada, December 2005, **acceptance rate (poster) 25%**.

14. *A Discriminative Matching Approach to Word Alignment*, **Ben Taskar**, Simon Lacoste-Julien and Dan Klein, 8 pages, Empirical Methods in Natural Language Processing (EMNLP), Vancouver, Canada, October 2005, **cited by 76**, **acceptance rate 32%**.

13. *Learning Structured Prediction Models: A Large Margin Approach*, **Ben Taskar**, Vassil Chatalbashev, Daphne Koller and Carlos Guestrin. 8 pages, International Conference on Machine Learning (ICML), Bonn, Germany, August 2005, **cited by 133**, **acceptance rate 27.3%**.

12. *Discriminative Learning of Markov Random Fields for Segmentation of 3D Scan Data*, Dragomir Anguelov, **Ben Taskar**, Vassil Chatalbashev, Daphne Koller, Dinkar Gupta, Geremy Heitz, and Andrew Y. Ng. 8 pages, Computer Vision and Pattern Recognition (CVPR), San Diego, California, June 2005, **cited by 71**, **acceptance rate 21.6%**.

11. *Exponentiated Gradient Algorithms for Large-margin Structured Classification*, Peter Bartlett, Michael Collins, **Ben Taskar** and David McAllester. 8 pages, Neural Information Processing Systems (NIPS), Vancouver, Canada, December 2004, **acceptance rate (spotlight) 7.9%**

10. *Max-Margin Parsing*, **Ben Taskar**, Dan Klein, Michael Collins, Daphne Koller and Christopher Manning. 8 pages, Empirical Methods in Natural Language Processing (EMNLP), Barcelona, Spain, July 2004 (Best paper award), **cited by 134**, **acceptance rate 24%**.

9. *Learning Associative Markov Networks*, **Ben Taskar**, Vassil Chatalbashev and Daphne Koller. 8 pages, International Conference on Machine Learning (ICML), Banff, Canada, July 2004, **cited by 62**, **acceptance rate 32%**.

8. *Max-Margin Markov Networks*, **Ben Taskar**, Carlos Guestrin and Daphne Koller. 8 pages, Neural Information Processing Systems (NIPS), Vancouver, Canada, December 2003 (Best student paper award), **cited by 451**, **acceptance rate (oral) 3.6%**.

7. *Link Prediction in Relational Data*, **Ben Taskar**, Ming-Fai Wong, Pieter Abbeel and Daphne Koller. 8 pages, Neural Information Processing Systems (NIPS), Vancouver, Canada, December 2003, **cited by 92**, **acceptance rate (poster) 27.6%**.
6. *Learning on the Test Data: Leveraging Unseen Features*, **Ben Taskar**, Ming-Fai Wong and Daphne Koller. 8 pages, International Conference on Machine Learning (ICML), Washington, DC, August 2003, **acceptance rate 32.1%**.
5. *Discriminative Probabilistic Models for Relational Data*, **Ben Taskar**, Pieter Abbeel and Daphne Koller. 8 pages, Uncertainty in Artificial Intelligence (UAI), Edmonton, Canada, August 2002, **cited by 339**, **acceptance rate 34%**.
4. *Probabilistic Classification and Clustering in Relational Data*, **Ben Taskar**, Eran Segal and Daphne Koller, 8 pages, International Joint Conference on Artificial Intelligence (IJCAI), Seattle, Washington, August 2001, **cited by 184**, **acceptance rate 25.0%**.
3. *Rich Probabilistic Models for Gene Expression*, Eran Segal, **Ben Taskar**, Audrey Gasch, Nir Friedman, and Daphne Koller, 9 pages, Intelligent Systems for Molecular Biology (ISMB), Copenhagen, Denmark, July 2001, **cited by 186**, **acceptance rate 22%**.
2. *Learning Probabilistic Models of Relational Structure*, Lise Getoor, Nir Friedman, Daphne Koller and **Ben Taskar**, 8 pages, International Conference on Machine Learning (ICML), Williamstown, Massachusetts, June 2001, **cited by 139**, **acceptance rate 32.1%**.
1. *Selectivity Estimation using Probabilistic Models*, Lise Getoor, **Ben Taskar** and Daphne Koller, 12 pages, ACM SIGMOD International Conference on Management of Data, San Diego, California, May 2001, **cited by 109**, **acceptance rate 15%**.