

Curriculum Vitae

JOSHUA B. TENENBAUM

August 15, 2006

BIOGRAPHICAL INFORMATION

Address: Department of Brain and Cognitive Sciences
Massachusetts Institute of Technology
Building 46-4015
Cambridge, MA 02139

Phone: 617-452-2010
Internet: <http://web.mit.edu/cocosci/josh.html>
Email: jbt@mit.edu

Date of Birth: August 21, 1972

EDUCATION

B.S. *magna cum laude* in Physics (with distinction in the major), May 1993.
Yale University.

Ph.D. in Brain and Cognitive Sciences, February 1999. Massachusetts Institute of
Technology. Thesis: *A Bayesian Framework for Concept Learning*.

PROFESSIONAL POSITIONS

2002-present, Assistant Professor of Cognitive Science and Computation, MIT.

2003-present, Principal Investigator, Computer Science and Artificial Intelligence
Laboratory (CSAIL), MIT.

2004-2007, Paul E. Newton Career Development Professor, MIT.

2004-2008, Associate Editor, *Cognitive Science*.

1999-2002, Assistant Professor of Psychology, Stanford University.

2000-2002, Assistant Professor of Computer Science, Stanford University (by
courtesy).

1999, Postdoctoral Associate, Massachusetts Institute of Technology, Artificial
Intelligence Laboratory.

HONORS AND AWARDS

Phi Beta Kappa

Howard Hughes Medical Institute (HHMI) Predoctoral Fellowship, 1993-1998.

Selected for graduate fellowships from ONR and NSF (declined).

Outstanding Paper Award for “Learning bilinear models for two-factor problems in vision”, with William T. Freeman, IEEE Conference on Computer Vision and Pattern Recognition, June 1997.

Distinguished Visiting Junior Scholar, Peter Wall Institute for Advanced Studies, University of British Columbia, July 2000.

Best Student Paper Award (honorable mention) for “Bayesian models of inductive generalization”, with Neville Sanjana (student advisee), Neural Information Processing Systems conference, 2002.

Paul E. Newton Career Development Professorship, MIT, 2004-2007.

Best Student Paper Award for “From algorithmic to subjective randomness”, with Thomas L. Griffiths (student advisee), Neural Information Processing Systems conference [Natural Systems category], 2003.

Best Student Paper Award for “Hierarchical topic models and the nested Chinese restaurant process”, with David Blei, Thomas L. Griffiths (student advisee), and Michael I. Jordan, Neural Information Processing Systems conference [Synthetic Systems category], 2003.

Marr Prize for Best Student Paper (honorable mention) for “Using physical theories to infer hidden causal structure”, with Tom Griffiths (student advisee) and Liz Bonawitz, Twenty-Sixth Annual Conference of the Cognitive Science Society, 2004.

New Investigator Award, Society for Mathematical Psychology, 2005.

Distinguished Lecturer, University of Edinburgh, Division of Informatics, 2006.

Distinguished Lecturer, University of British Columbia, Institute for Computing, Information, and Cognitive Systems, 2007.

PROFESSIONAL ACTIVITIES

Editorial and reviewing

Associate Editor, *Cognitive Science*, 2004-2008. Editorial board member since 2001.

Program Committee member, Annual Conference of the Cognitive Science Society, 2002-2006.

Program Committee member, Neural Information Processing Systems Conference, 2000-2002. Chair of Cognitive Science section.

Tutorials Chair, Neural Information Processing Systems Conference, 2005.

Program Committee member, National Academy of Sciences Sackler Colloquium on Mapping Knowledge Domains, May 2003.

Program Committee member, Workshop on “Statistical Network Analysis: Models, Issues, and New Directions”, June 2006.

Journal reviewing: *Behavioral and Brain Sciences, BioMedCentral Bioinformatics, Cognition, Cognitive Science, Communications of the ACM, IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Transactions on Neural Networks, Journal of Classification, Journal of Experimental Psychology: General, Journal of Machine Learning Research, Journal of Mathematical Imaging and Vision, Machine Learning, Memory and Cognition, Neural Computation, Proceedings of the National Academy of Sciences, Psychological Review, Psychological Science, Psychonomic Bulletin and Review, SCIENCE, Trends in Cognitive Sciences.*

Conference reviewing: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), National Conference on Artificial Intelligence (AAAI), Neural Information Processing Systems (NIPS), Annual Conference of the Cognitive Science Society (COGSCI).

Grant reviewing: National Science Foundation (NSF), Natural Sciences and Engineering Research Council of Canada (NSERC), Human Frontiers Science Program (HFSP).

Book proposals reviewed for: MIT Press, CRC Press, Oxford University Press.

Organizing symposia, workshops, and conferences

Co-organizer, Symposium on Bayesian Approaches to Cognitive Modeling, Twenty Second Annual Conference of the Cognitive Science Society, University of Pennsylvania, August 2000.

Co-organizer, Symposium on Bayesian Models in Cognitive Development, Biennial Meeting of the Society for Research on Child Development, Minneapolis, MN, April 2001.

Co-organizer, Workshop on Causal Learning and Inference in Humans and Machines, Neural Information Processing Systems conference, Whistler, B.C., Canada, December, 2001.

Co-organizer, Workshop on Spectral Methods in Dimensionality Reduction, Clustering, and Classification, Neural Information Processing Systems conference, Whistler, B.C., Canada, December, 2002.

Co-organizer, Workshop on Structured Data and Representations in Probabilistic Models for Categorization, Neural Information Processing Systems conference, Whistler, B.C., Canada, December, 2004.

Co-organizer, 5-Day Conference on “Probabilistic Models of Cognition: The Mathematics of Mind”, Institute for Pure and Applied Mathematics, UCLA, January 2005.

Organizer, McDonnell Causality Research Network meeting on “Prior Knowledge and Explanation in Causal Learning”, MIT, November 2005.

Co-organizer, Three-week summer school on “Probabilistic Models of Cognition”, Institute for Pure and Applied Mathematics, UCLA, July 2007.

Media features and educational outreach

Interview in The New York Times Magazine, “The Odds of That”, August 11, 2002.

Invited talk, Knight Fellows Science Journalism Seminar, November 2002.

Interview on CBC (Canadian Broadcasting Corporation) Radio One, “Coincidences”, August 2, 2003.

Invited talk, Knight Fellows Medical Evidence Bootcamp, December 1, 2003.

Interview for The Infinite Mind (internationally syndicated radio show on NPR and PRI), “Coincidences”, December 3, 2003.

Interview and research feature in *Psychology Today*, “Pattern & Circumstance: The Power of Coincidence”, August 2004.

Research feature, “Bayes Rules”, *The Economist*, January 2006.

Visiting lecturer, Canada/USA Mathcamp, 1999-2006.

RESEARCH INTERESTS

Mathematical theories, empirical studies, and computational models of human learning and inference, including:

- Constructing perceptual representations
- Learning concepts and words
- Reasoning: similarity, representativeness, coincidence detection, predicting the future
- Causal induction
- The structure and acquisition of intuitive theories

Modeling methodologies drawing on:

- Bayesian statistics and probability theory
- Geometry of manifolds
- Graph theory
- Multilinear models

GRANTS AND RESEARCH SUPPORT

The Schlumberger Foundation, September 1999, \$30,000, “Manifold learning for nonlinear inverse problems.”

Center for the Study of Language and Information, Stanford University, November 1999, \$12,500, “A Bayesian framework for example-driven database search.”

MERL (Mitsubishi Electric Research Labs), January 2000, \$20,000, “Bayesian models of perception and cognition.”

The Schlumberger Foundation, May 2000, \$30,000, “Manifold learning for nonlinear inverse problems” (continuation of previous award).

DARPA/ONR, October 2000 – July 2002, \$137,613, “Human ID and the structure of perceptual face space.”

NTT Communications Sciences Laboratory, April 2001-April 2002, \$115,000, “Theory and applications of Bayesian concept learning.”

MERL (Mitsubishi Electric Research Labs), April 2001, \$10,000, “Bayesian models of perception and cognition”.

NSF Focused Research Group, July 2001-July 2004, \$1,000,000, “Topological methods in data analysis”, with Gunnar Carlsson and Persi Diaconis.

NTT Communications Sciences Laboratory, July 2002-June 2006, \$340,000, “Learning semantic structure.”

DARPA/SRI CALO Project, July 2004-October 2006, \$320,000, “Topic modeling and identification.”

AFOSR MURI, June 2005-June 2008, \$4,600,000, “Computational modeling of adversary attitudes and behaviors: computational models for belief structures and intentions”, with Whitman Richards (PI) and eight other investigators across five institutions. MIT component, supporting Richards and Tenenbaum, is \$840,000.

James S. McDonnell Foundation Causal Learning Research Collaborative, September 2005-July 2007, \$87,773, “Prior knowledge in causal learning and categorization.”

DARPA BICA (Biologically Inspired Cognitive Architectures) Program, October 2005-November 2006. Support for Phase I study by CSAIL team, with Patrick

Winston (PI) and several other investigators. Total Phase I support for Tenenbaum is \$75,000.

ONR MURI, June 2007-June 2012 (white paper submitted, application pending), “Robust learning and skill transfer with video games”, with Daphne Bavelier (PI) and four other investigators across four institutions.

PATENTS

1. J. B. Tenenbaum & W. T. Freeman. System for analyzing and synthesis of multi-factor data. U.S. Patent #6,549,899, April 2003.

PUBLICATIONS

REFEREED JOURNALS

1. Shult, D. L. and Tenenbaum, J. B. (1988). A three-dimensional multiplication table inspires mathematical discovery. *The Mathematics Teacher* 81(2), 92-25.
2. Tenenbaum, J. B. and Freeman, W. T. (2000). Separating style and content with bilinear models. *Neural Computation*, 12 (6), 1247-1283.
3. Tenenbaum, J. B., de Silva, V., and Langford, J. C. (2000). A global geometric framework for nonlinear dimensionality reduction. *Science*, 290 (5500): 2319-2323.
4. Tenenbaum, J. B. and Griffiths, T. L. (2001). Generalization, similarity, and Bayesian inference. *Behavioral and Brain Sciences*, 24(4), 629-641.
5. Tenenbaum, J. B. and Griffiths, T. L. (2001). Some specifics about generalization. *Behavioral and Brain Sciences*, 24(4), 772-778.
6. Balasubramanian, M., Shwartz, E. L., Tenenbaum, J. B., de Silva, V., and Langford, J. C. (2002). The Isomap Algorithm and Topological Stability. *Science* Jan 4 2002: 7.
7. Freeman, W. T., Tenenbaum, J. B., and Pasztor, E. (2002). Learning style translation for line drawings. *ACM Transactions on Graphics*, 22(1), 33-46.
8. Steyvers, M., Tenenbaum, J. B., Wagenmakers, E., and Blum, B. (2003). Inferring causal networks through observations and interventions. *Cognitive Science*, 27, 453-489.
9. Sharma, J., Dragoi, V., Tenenbaum, J. B., Miller, E. K., and Sur, M. (2003). V1 neurons signal acquisition of an internal representation of stimulus location. *Science*, 300, 1758-1763.
10. Sobel, D., Tenenbaum, J. B., and Gopnik A. (2004). Children’s causal inferences from indirect evidence: Backwards blocking and Bayesian reasoning in preschoolers. *Cognitive Science* 28(3), 303-333.
11. Steyvers, M. and Tenenbaum, J. B. (2005). The large-scale structure of semantic networks: statistical analyses and a model of semantic growth. *Cognitive Science* 29(1), 41-78.
12. Griffiths, T. L. and Tenenbaum, J. B. (2005). Structure and strength in causal induction. *Cognitive Psychology* 51(4), 285-386.
13. Saxe, R., Tenenbaum, J. B., and Carey, S. (2005). Secret agents: Inferences about hidden

- causes by 10- and 12-month-old infants. *Psychological Science* 16(12), 995-1001.
14. Tenenbaum, J. B., Griffiths, T. L., and Kemp, C. (2006). Theory-based Bayesian models of inductive learning and reasoning. *Trends in Cognitive Sciences* 10(7), 309-318.
 15. Griffiths, T. L. and Tenenbaum, J. B. (in press). Optimal predictions in everyday cognition. *Psychological Science*.
 16. Griffiths, T. L. and Tenenbaum, J. B. (in press). From mere coincidences to meaningful discoveries. *Cognition*.
 17. Xu, F. and Tenenbaum, J. B. (in press). Sensitivity to sampling in Bayesian word learning. *Developmental Science*.
 18. Kemp, C., Perfors, A. and Tenenbaum, J. B. (in press). Learning overhypotheses with hierarchical Bayesian models. *Developmental Science*.
 19. T. Iwata, K. Saito, N. Ueda, S. Stromsten, T. L. Griffiths, J. B. Tenenbaum (revision submitted). Parametric Embedding for Class Visualization. *Neural Computation*.
 20. Tenenbaum, J. B., Sobel, D. M., Griffiths, T. L., and Gopnik, A. (revision submitted). Bayesian reasoning in adults' and children's causal inferences. *Cognition*.
 21. Krynski, T. R. and Tenenbaum, J. B. (revision submitted). The role of causality in judgment under uncertainty. *Journal of Experimental Psychology: General*.
 22. Griffiths, T. L., Steyvers, M. and Tenenbaum, J. B. (revision submitted). Topics in semantic representation. *Psychological Review*.
 23. Shafto, P., Kemp, C., Baraff, E. R., Coley, J., and Tenenbaum, J. B. (revision submitted). Inductive reasoning about causally transmitted properties. *Cognition*.
 24. Xu, F. and Tenenbaum, J. B. (revision submitted). Word learning as Bayesian inference. *Psychological Review*.
 25. Kemp, C. and Tenenbaum, J. B. (submitted). The discovery of structural form. *Science*.
 26. Kording, K., Tenenbaum, J. B., and Shadmehr, R. (submitted). Motor learning and memory as Bayesian inference over multiple timescales. *Nature Neuroscience*.
 27. Kording, K., Beierholm, U., Ma, W. J., Quartz, S., Tenenbaum, J. B., and Shams, L. (under revision). Causal inference in multisensory perception.
 28. Kemp, C. and Tenenbaum, J. B. (under revision). Structured statistical models of inductive reasoning.

REFEREED CONFERENCE PROCEEDINGS

29. Tenenbaum, J. B. (1994). Functional parts. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Erlbaum, 864-869. [poster]
30. Tenenbaum, J. B. and Todorov, E. V. (1995). Factorial learning by clustering features. In G. Tesauro, D. Touretzky, and T. Leen (eds.), *Advances in Neural Information Processing Systems 7*. Cambridge, MA: MIT Press, 561-568. [poster]
31. Tenenbaum, J. B. (1996). Learning the structure of similarity. In D. Touretzky, M. Mozer, and M. Hasselmo (eds.), *Advances in Neural Information Processing Systems 8*. Cambridge, MA: MIT Press, 3-9. [oral]
32. Tenenbaum, J. B. and Freeman, W. T. (1997). Separating style and content. In M. Mozer, M. Jordan, and T. Petsche (eds.), *Advances in Neural Information Processing Systems 9*. Cambridge, MA: MIT Press, 662-668. [oral]
33. Freeman, W. T. and Tenenbaum, J. B. (1997). Learning bilinear models for two-factor

- problems in vision. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. IEEE Press, 554-560. [oral]
34. Tenenbaum, J. B. (1997). A Bayesian framework for concept learning. *Proceedings of the Interdisciplinary Workshop on Similarity and Categorization 1997*, University of Edinburgh, 249-255. [oral]
 35. Tenenbaum, J. B. (1998). Mapping a manifold of perceptual observations. In M. Jordan, M. Kearns, and S. Solla (eds.), *Advances in Neural Information Processing Systems 10*. Cambridge, MA: MIT Press, 682-688. [poster]
 36. Tenenbaum, J. B. (1999). Bayesian modeling of human concept learning. In M. Kearns, S. Solla, and D. Cohn (eds.), *Advances in Neural Information Processing Systems 11*. Cambridge, MA: MIT Press, 59-65. [spotlight]
 37. Tenenbaum, J. B. (2000). Rules and similarity in concept learning. In S. Solla, T. Leen, and K. R. Mueller (eds.), *Advances in Neural Information Processing Systems 12*. Cambridge, MA: MIT Press, 59-65. [oral]
 38. Griffiths, T. L. and Tenenbaum, J. B. (2000). Teacakes, trains, toxins, and taxicabs: A Bayesian account of predicting the future. *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, 202-207. [oral]
 39. Tenenbaum, J. B. and Xu, F. (2000). Word learning as Bayesian inference. *Proceedings of the Twenty-Second Annual Conference of the Cognitive Science Society*, 517-522. [oral]
 40. Tenenbaum, J. B. and Griffiths, T. L. (2001). Structure learning in human causal induction. In T. Leen, T. Dietterich, and V. Tresp (eds.), *Advances in Neural Information Processing Systems 13*. Cambridge, MA: MIT Press, 59-65. [poster]
 41. Tenenbaum, J. B. and Griffiths, T. L. (2001). The rational basis of representativeness. *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, 1036-1041. [oral]
 42. Griffiths, T. L. and Tenenbaum, J. B. (2001). Randomness and coincidences: Reconciling intuition and probability theory. *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society*, 370-375. [oral]
 43. Griffiths, T. L. and Tenenbaum, J. B. (2002). Using vocabulary knowledge in Bayesian multinomial estimation. In T. G. Dietterich, S. Becker, and Z. Ghahramani (eds.), *Advances in Neural Information Processing Systems 14*. Cambridge, MA: MIT Press, 1385-1392. [poster]
 44. Tenenbaum, J. B. and Griffiths, T. L. (2003). Theory-based causal inference. In S. Becker, S. Thrun, and K. Obermayer (eds.), *Advances in Neural Information Processing Systems 15*. Cambridge, MA: MIT Press, 35-42. [poster]
 45. Sanjana, N. and Tenenbaum, J. B. (2003). Bayesian modeling of inductive generalization. In S. Becker, S. Thrun, and K. Obermayer (eds.), *Advances in Neural Information Processing Systems 15*. Cambridge, MA: MIT Press, 51-58. [spotlight]
 46. Danks, D., Griffiths, T. L., and Tenenbaum, J. B. (2003). Dynamical causal learning. In S. Becker, S. Thrun, and K. Obermayer (eds.), *Advances in Neural Information Processing Systems 15*. Cambridge, MA: MIT Press, 67-74. [poster]
 47. de Silva, V., and Tenenbaum, J. B. (2003). Local versus global approaches to nonlinear dimensionality reduction. In S. Becker, S. Thrun, and K. Obermayer (eds.), *Advances in Neural Information Processing Systems 15*. Cambridge, MA: MIT Press, 705-712. [poster]

48. Griffiths, T. L. and Tenenbaum, J. B. (2003). Probability, algorithmic complexity, and subjective randomness. *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*, 480-485. [poster]
49. Kemp, C. S. and Tenenbaum, J. B. (2003). Theory-based induction. *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*, 658-663. [oral]
50. Krynski, T. R. and Tenenbaum, J. B. (2003). The role of causal models in reasoning under uncertainty. *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*, 693-698. [oral]
51. Tenenbaum, J. B. and Niyogi, S. (2003). Learning causal laws. *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*, 1152-1157. [oral]
52. Griffiths, T. L. and Tenenbaum, J. B. (2004). From algorithmic to subjective randomness. In S. Thrun, L. K. Saul, and B. Scholkopf (eds.), *Advances in Neural Information Processing Systems 16*, 953-960. [oral]
53. Blei, D., Griffiths, T. L., Jordan, M. I., and Tenenbaum, J. B. (2004). Hierarchical topic models and the nested Chinese restaurant process. In S. Thrun, L. K. Saul, and B. Scholkopf (eds.), *Advances in Neural Information Processing Systems 16*, 17-24. [oral]
54. Kemp, C., Griffiths, T. L., Stromsten, S., and Tenenbaum, J. B. (2004). Semi-supervised learning with trees. In S. Thrun, L. K. Saul, and B. Scholkopf (eds.), *Advances in Neural Information Processing Systems 16*, 257-264. [poster]
55. Balas, B. J. and Tenenbaum, J. B. (2004). Domain-specificity in the perception and representation of shape. *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society*. [oral]
56. Griffiths, T. L., Baraff, E., R., and Tenenbaum, J. B. (2004). Using physical theories to infer hidden causal structure. *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society*. [oral]
57. Kemp, C., Perfors, A. F., and Tenenbaum, J. B. (2004). Learning domain structure. *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society*. [oral]
58. T. L. Griffiths, M. Steyvers, D. Blei, and J. B. Tenenbaum (2005). Integrating topics and syntax. In L. K. Saul, Y. Weiss, and L. Bottou (eds.), *Advances in Neural Information Processing Systems 17*, 537-544. [spotlight]
59. T. Iwata, K. Saito, N. Ueda, S. Stromsten, T. L. Griffiths, J. B. Tenenbaum (2005). Parametric Embedding for Class Visualization. *Advances in Neural Information Processing Systems 17*, 617-624. [oral]
60. Kemp, C., Bernstein, A., and Tenenbaum, J. B. (2005). A generative theory of similarity. *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society*. [poster]
61. Shafto, P., Kemp, C., Baraff, L., Coley, J., and Tenenbaum, J. B. (2005). Context-sensitive induction. *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society*. [oral]
62. Perfors, A., Kemp, C., and Tenenbaum, J. B. (2005). Modeling the acquisition of domain structure and feature understanding. *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society*. [poster]
63. Xu, F., and Tenenbaum, J. B. (2005). Word learning as Bayesian inference: Evidence from Preschoolers. *Proceedings of the Twenty-Seventh Annual Conference of the Cognitive Science Society*. [oral]

64. Baker, C., Saxe, R., and Tenenbaum J. B. (2006). A Bayesian framework for human action understanding. In Y. Weiss, B. Scholkopf, and J. Platt (eds.), *Advances in Neural Information Processing Systems 18*, 99-106. [poster]
65. Kemp, C., Perfors, A., and Tenenbaum J. B. (in press). Learning overhypotheses. *Proceedings of the Twenty-Eighth Annual Conference of the Cognitive Science Society*. [oral]
66. Perfors, A., Tenenbaum J. B., and T. Regier (in press). Poverty of the stimulus? A rational approach. *Proceedings of the Twenty-Eighth Annual Conference of the Cognitive Science Society*. [oral]
67. Schmidt, L., Kemp, C., and Tenenbaum J. B. (in press). Nonsense and sensibility: Inferring unseen possibilities. *Proceedings of the Twenty-Eighth Annual Conference of the Cognitive Science Society*. [oral]
68. Shafto, P., Kemp, C., Mansinghka, V., Gordon, M., and Tenenbaum J. B. (in press). Discovering systems of concepts. *Proceedings of the Twenty-Eighth Annual Conference of the Cognitive Science Society*. [poster]
69. Goodman, N. D., Baker, C. L., Bonawitz, E. B., Mansinghka, V. K., Gopnik, A., Wellman, H., Schulz, L., and Tenenbaum, J. B. (in press). Intuitive theories of mind: A rational approach to false belief. *Proceedings of the Twenty-Eighth Annual Conference of the Cognitive Science Society*. [poster]
70. Kemp, C., Tenenbaum J. B., Griffiths, T. L., Yamada, T., and Ueda, N. (in press). Learning systems of concepts with an infinite relational model. *Twenty-First National Conference on Artificial Intelligence (AAAI 2006)*. [oral]
71. Purver, M., Kording, K., Griffiths, T. L., and Tenenbaum, J. B. (in press). Unsupervised topic modeling for multi-party spoken discourse. *Proceedings of COLING-ACL 2006 (Joint Conference of the International Committee on Computational Linguistics and the Association for Computational Linguistics)*. [oral]
72. Mansinghka, V. K., Kemp, C., Tenenbaum, J. B., and Griffiths, T. L. (in press). Structured priors for structure learning. *Twenty-Second Conference on Uncertainty in Artificial Intelligence (UAI 2006)*. [oral]
73. Kemp, C., Shafto, P., Berke, A., and Tenenbaum, J. B. (in press). Combining causal and similarity-based reasoning. *Advances in Neural Information Processing Systems 19*. [oral]
74. Roy, D., Kemp, C., Mansinghka, V. K., and Tenenbaum, J. B. (in press). Learning annotated hierarchies from features and relations. *Advances in Neural Information Processing Systems 19*. [oral]
75. Kording, K. and Tenenbaum, J. B. (in press). Causal inference in multisensory and sensorimotor integration. *Advances in Neural Information Processing Systems 19*. [poster]
76. Kording, K., Tenenbaum, J. B., and Shadmehr, R. (in press). Multiple timescales and uncertainty in motor adaptation. *Advances in Neural Information Processing Systems 19*. [spotlight]

BOOK CHAPTERS (NOT COMPETITIVELY REFEREED)

77. Tenenbaum, J. B. and Bricolo, E. (1997). Analyzing the view-dependence of population codes in inferior temporal cortex. In J. M. Bower, (ed.) *Computational Neuroscience: Trends in Research 1997*. New York: Plenum. 867-872.

78. de Silva, V., and Tenenbaum, J. B. (2002). Unsupervised learning of curved manifolds. In D. D. Denison, M. H. Hansen, C. C. Holmes, B. Mallick and B. Yu (eds.), *Nonlinear Estimation and Classification*, Springer-Verlag, New York.
79. Tenenbaum, J.B., Griffiths, T. L., and Niyogi, S. (in press). Intuitive theories as grammars for causal inference. In A. Gopnik and L. Schulz (eds.), *Causal Learning*. Oxford University Press.
80. Griffiths, T. L. and Tenenbaum, J.B. (in press). Two proposals for causal grammar. In A. Gopnik and L. Schulz (eds.), *Causal Learning*. Oxford University Press.
81. Tenenbaum, J. B., Kemp, C., Shafto, P. (in press). Generative Bayesian models for inductive reasoning. In A. Feeney and E. Heit (eds.), *Induction*. Cambridge University Press.

OTHER PUBLICATIONS (NOT COMPETITIVELY REFEREED)

82. Xu, F., Tenenbaum, J. B., and Sorrentino, C. (1998). Concepts are not beliefs, but having concepts is having beliefs. Comment on R. Millikan, "A common structure for concepts of individuals, stuffs, and real kinds: more mama, more milk, and more mouse." *Behavioral and Brain Sciences*, 22 (1).
83. Tenenbaum, J. B. (1999). *A Bayesian Framework for Concept Learning*. Ph.D. Thesis, Massachusetts Institute of Technology.
84. C. Kemp, T. L. Griffiths, and J. B. Tenenbaum (2004). Discovering latent classes in relational data. MIT AI Memo 2004-019.
85. Chater, N., Tenenbaum, J. B., and Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences* 10(7), 287-291..
86. T. L. Griffiths and J. B. Tenenbaum (in press). Statistics and the Mind. *Significance*. [Magazine of the Royal Statistical Society]

INVITED TALKS

1. Yale University, Department of Computer Science, Vision seminar, November 1996.
2. University of Toronto, Department of Computer Science, Neural Networks ResearchGroup, October 1997.
3. Brown University, Department of Cognitive and Linguistic Sciences, Vision and Neural Networks seminar, November 1997.
4. Schlumberger Cambridge Research, UK, November 1997.
5. NEC Research Institute, Vision Research seminar, January 1998.
6. University of California, Irvine, Department of Cognitive Sciences Colloquium, February, 1998.
7. University of Rochester, Department of Brain and Cognitive Science Colloquium, February 1998.
8. Stanford University, Department of Psychology Colloquium, February 1998.
9. Workshop on Learning from Relational Data Representations, Breckenridge, CO, December 1998.

10. Yale University, Department of Psychology, Cognitive Seminar, February 1999.
11. Invited Discussant, "Finding the Words..." workshop, Stanford University, April 2000.
12. Stanford University, Center for the Study of Language and Information, May 2000.
13. Cold Spring Harbor Laboratory, Bainbury Center Workshop on Network Models of Brain Function, September 2000.
14. Schlumberger Doll Research, Ridgefield, CT, September 2000.
15. UC Berkeley, Cognitive Science Colloquium, November 2000.
16. MIT, Brain and Cognitive Sciences Seminar, February 2001.
17. MIT, Perceptual Sciences Seminar, February 2001.
18. Stanford University, Department of Statistics Colloquium, April 2001.
19. University of Edinburgh, Division of Informatics, Colloquium, August 2001.
20. Gatsby Institute, University College London, Colloquium, August 2001.
21. Indiana University, Cognitive Science Colloquium, April 2002.
22. Indiana University, Mathematical Psychology Seminar, April 2002
23. Harvard University, Concepts and Words Seminar, April 2002
24. Brown University, Causality Seminar, April 2002.
25. New York University, Cognitive Psychology Colloquium, September 2002.
26. University of Pennsylvania, Institute for Research in Cognitive Science (IRCS) Colloquium, September 2002.
27. Cornell University, Department of Psychology Colloquium, October 2002.
28. Harvard University, Department of Psychology, Cognition, Brain and Behavior Seminar, October 2002.
29. Symposium on Causal Inference, Biennial Meeting of the Philosophy of Science Association, Milwaukee, WI, November 2002.
30. NTT Communication Sciences Research Labs, Intelligent Communication Laboratory, January 2003.
31. Yale University, Department of Psychology, Cognitive Lunch, February 2003.
32. Brown University, Brain Sciences Program Speaker Series, March 2003.
33. Northwestern University, Cognitive Science Program, Week-long Tutorial, May 2003.
34. Annual Meeting of the Society for Mathematical Psychology, Invited Tutorial, July 2003.
35. Rutgers University, Center for Cognitive Science (RUCCS), Colloquium, October 2003.
36. MIT Sloan School, Marketing Seminar, November 2003.
37. MIT Japan program, Mitsui/Toshiba Executive Mission, November 2003.
38. Air Force Office of Scientific Research, Workshop on Culture and Personality in Models of Adversarial Decision Making, McLean, VA, November 2003.
39. MIT Language, Cognition, and Computation Seminar Series, December 2003.
40. Knight Fellows Medical Evidence Bootcamp, Keynote talk, December 2003.
41. Workshop on Syntax, Semantics, and Statistics, Invited discussant, Whistler, B.C., Canada, December 2003.
42. Workshop on Infinite Models and Nonparametric Bayesian methods, Invited talk, December 2003.

43. The Johns Hopkins University, Center for Language and Speech Processing, Invited talk, February 2004.
44. MIT Stochastic Systems Seminar, Laboratory for Information and Decision Systems, April 2004.
45. Carnegie Mellon University, Department of Psychology, Colloquium, April 2004.
46. Carnegie Mellon University, Department of Psychology, Brownbag lunch speaker, April 2004.
47. Center for Advanced Study in the Behavioral Sciences, Stanford, CA, Invited Symposium on Causation, Theory Formation, and Explanation, April 2004.
48. International workshop on “Theories of Causality in Cognition: Frameworks, Frailties and Future Directions”, Invited speaker, July 2004.
49. International Conference on Thinking, Invited Symposium on Induction, July 2004.
50. Workshop on Probabilistic Models of Cognition: The Mathematics of Mind, Institute for Pure and Applied Mathematics, UCLA, January 2005.
51. University of Rochester, Department of Brain and Cognitive Sciences, Colloquium, February, 2005.
52. Brown University, Symposium for Undergraduates in the Mathematical Sciences, February 2005.
53. Stanford University, Department of Computer Science, Broad Area Colloquium for Artificial Intelligence, Graphics, Geometry and Computer Vision, March 2005.
54. California Institute of Technology, Sloan-Swartz Center for Theoretical Neuroscience, Colloquium, March 2005.
55. Rensselaer Polytechnic Institute, Cognitive Science Colloquium, September 2005.
56. Harvard University, Applied Statistics Workshop, October 2005.
57. International workshop on Bayesian Cognition, College de France, Paris, France, January 2006.
58. University of Arizona, Cognitive Science Colloquium, March 2006.
59. University of Southern California/Information Sciences Institute, “Young Stars” Colloquium Series, March 2006.
60. Harvard University, Computer Science Colloquium, May 2006.
61. University of Rochester, Center for Visual Science symposium on “Statistical learning and brain plasticity”, May 2006.
62. New York University Conference on Causation, Villa La Pietra, Florence, Italy, June 2006.
63. University of Edinburgh, Distinguished Lecture in Informatics, June 2006.
64. Rational Models of Cognition, Invited Conference Address, University College London, June 2006.
65. DARPA ISAT/IXO Workshop on “Adaptive and Interactive Representations”, Cambridge, MA, July 2006.
66. Society for Mathematical Psychology, Plenary address, July 2006.
67. NSF Workshop on “Mind and Brain: Strategies and Directions for Future Research”, Arlington, VA, July 2006.

68. Symposium on “Exploring Mind and World” (in honor of Roger Shepard), Annual Conference of the Cognitive Science Society, July 2006.
69. Symposium on “Building and Evaluating Models of Human-Level Intelligence”, Annual Conference of the Cognitive Science Society, July 2006.
70. Invited Tutorial, Neural Information Processing Systems Conference, Vancouver, B.C., Canada, December 2006.
71. University of British Columbia, Distinguished Lecture, Institute for Computing, Information, and Cognitive Systems, March 2007.
72. University of Washington, Center for Statistics in the Social Sciences (CSSS) seminar, March 2007.

OTHER CONFERENCE PRESENTATIONS

1. Shepard, R. N. and Tenenbaum, J. B. (1991). Multidimensional models of generalization and categorization. Annual Meeting of the Psychonomic Society.
2. Tenenbaum, J. B. (1997). Making sense of typicality: What makes good examples? Nineteenth Annual Conference of the Cognitive Science Society.
3. Tenenbaum, J. B. (1997). A Bayesian theory of concept learning. Workshop on Neural Models of Concept Learning, Breckenridge, CO, December 6, 1997.
4. Freeman, W. T. and Tenenbaum, J. B. (1998). Learning bilinear models for two-factor problems in vision. Investigative Ophthalmology and Visual Science (ARVO), Ft. Lauderdale, FL.
5. Xu, F. and Tenenbaum, J. B. (2000). Word learning as Bayesian inference: The case of nested categories. 25th Boston University Conference on Child Language.
6. Tenenbaum, J. B. and Xu, F. (2002). Bayesian inference in learning words with overlapping extensions. 27th Boston University Conference on Child Language.
7. Sobel, D. M. & Tenenbaum, J. B. (2003). The development of causal inferences based on indirect evidence.. Society for Research on Child Development.
8. Saxe, R., Tenenbaum, J. B., & Carey, S. (2004). Inferring the presence of hidden agents. International Society for Infant Studies.
9. Tenenbaum, J. B. and Kemp, C. (2004). Causality and prior knowledge in rational models of inductive reasoning. Annual Meeting of the Psychonomic Society.

TEACHING

Stanford University

Fall 1999: Psychology 205, Foundations of Cognition, sole instructor.

Spring 2000: Psychology 224, Learning and Inference in Humans and Machines, sole instructor.

Fall 2000: Psychology 205, Foundations of Cognition, sole instructor.

Spring 2001: Psychology 224, Learning and Inference in Humans and Machines, sole instructor.

Fall 2001: Psychology 205, Foundations of Cognition, sole instructor.

MIT

Spring 2002: 9.012, The Brain and Cognitive Sciences II (18 units), one of several co-instructors, responsible for four lectures.

Spring 2003: 9.69/9.690, Foundations of Cognition (9 units), one of two co-instructors, responsible for 50% of lectures.

Spring 2003: 9.012, The Brain and Cognitive Sciences II (18 units, partially overlapping with 9.69/9.690), one of five co-instructors and two co-organizers, responsible for 25% of lectures.

Fall 2003: 9.52/9.914, Computational Models of Concepts and Theories (9 units), one of two co-instructors, responsible for 50% of lectures.

IAP 2004: 9.9X, From Vision to Cognition (3 units, credit/no-credit), co-organizer, responsible for one lecture.

Spring 2004: 9.012, The Brain and Cognitive Sciences II (18 units), one of four co-instructors and two co-organizers, responsible for 30% of lectures.

Fall 2004: 9.66/9.660, Computational Cognitive Science (12 units), sole lecturer.

Spring 2005: 9.012, The Brain and Cognitive Sciences II (18 units), organizer and one of several co-instructors, responsible for 30% of lectures.

Fall 2005: 9.66/9.660/6.804, Computational Cognitive Science (12 units), sole lecturer.

IAP 2006: 9.94, The Cognitive Science of Intuitive Theories (3 units, credit/no-credit), organizer, responsible for one lecture.

Spring 2006: 9.012, Cognitive Science (18 units), one of three co-instructors, responsible for 35% of lectures.

External Tutorials on Bayesian Models in Cognitive Science

May 2003: Northwestern University, twelve hours, sole instructor.

July 2003: Society for Mathematical Psychology, 90 minutes, sole instructor.

July 2004: Annual Conference of the Cognitive Science Society, six hours, co-taught (responsibility shared equally) with Tom Griffiths.

July 2006: Annual Conference of the Cognitive Science Society, six hours, co-taught

(responsibility shared equally) with Tom Griffiths and Charles Kemp.

December 2006: Neural Information Processing Systems (NIPS) Conference, two hours, sole instructor.

RESEARCH SUPERVISED

Postdoctoral

Mark Steyvers, 2000-2002, Models of causal learning and semantic memory. Currently Associate Professor at UC Irvine, Department of Cognitive Sciences.

Vin de Silva, 2000-2004, Nonlinear dimensionality reduction. Currently Assistant Professor at Pomona College, Department of Mathematics.

Sean Stromsten, 2003-2004, Models of concept learning. Currently Member of Research Staff at BAE Systems Advanced Information Technologies.

Konrad Koerding, 2004-2006, Probabilistic models for motor learning, sensory integration, topic detection, and computational neuroscience. Currently Assistant Professor at Northwestern University, Department of Physiology.

Pat Shafto, 2004-present, Models and experiments on inductive reasoning.

Noah Goodman, 2005-present, Models of causal learning, concept learning, and cognitive development.

Kobi Gal, 2006-present, Theory of mind and models of belief structure in adversarial decision making.

Ph. D. Students

Thomas L. Griffiths, 1999-2004, Models of causal learning, semantic knowledge, and intuitive theories. Currently Assistant Professor of Psychology and Cognitive Science, University of California, Berkeley.

Tevye R. Kryniski, 2003-present (NSF Fellowship), Probabilistic reasoning and causal inference.

Charles Kemp, 2002-present (Albert Memorial Fellowship, MIT), Modeling the structure of mental representations.

Lauren Schmidt, 2004-present (NSF Fellowship), The acquisition of semantic knowledge.

Amy Perfors, 2004-present (NSF Fellowship, NDSEG Fellowship), Models of cognitive development and language acquisition.

Chris Baker, 2004-present (NSF Fellowship, DHS Fellowship), Theory of mind and

activity recognition.

Vikash Mansinghka, 2005-present (NSF Fellowship), Models for learning intuitive theories.

Masters Students

Lauren Schmidt, 2001-2002, Studies of word learning. Now Ph.D. student at MIT, Cognitive Science.

Gautam K. Jayarama, 2002-2003, Fast nearest-neighbor computations for nonlinear dimensionality reduction. Now at Amazon.com.

PhD Committees

Timothy Paek, Stanford University, Psychology, PhD 2000.

Kalina Christoff, Stanford University, Psychology, PhD 2001.

Karl Pflieger, Stanford University, Computer Science, PhD 2002. .

Rebecca Saxe, MIT, Brain and Cognitive Sciences, PhD 2004.

Rajesh Kasturirangan, MIT, Brain and Cognitive Sciences, PhD 2004.

Mike Dowman, University of Sydney, Computer Science, PhD 2004.

Nathan Srebro, MIT, Computer Science, PhD 2004.

Daniel Casasanto, MIT, Brain and Cognitive Sciences, PhD 2005.

Florian Wolf, MIT, Brain and Cognitive Sciences, PhD 2005.

Ashish Kapoor, MIT, Media Lab, PhD 2006.

Erik Sudderth, MIT, Computer Science, PhD 2006.

Tania Lombrozo, Harvard, Psychology, PhD 2006.

Mark Finlayson, MIT, Computer Science, PhD Expected 2007.

Bryan Russell, MIT, Computer Science, PhD Expected 2007.

Undergraduate Research Students

Mark Pearson, 2000, now Ph.D. student at Berkeley, Computer Science.

Neville Sanjana, 2000-2001, now Ph.D. student at MIT, Brain and Cognitive Science.

Khara Ramos, 2000-2001, now Ph.D. student at UCSD, Neuroscience.

Matt Bell, 2000-2001, now Chief Scientist, Reactrix, Inc.

Eric Berger, 2000-2001, now Ph.D. student at Stanford, Robotics.

John Paul D'India, 2001.

Ben Blum, 2001-2002, now Ph.D. student at Berkeley, Computer Science.

Ronnie Bryan, 2002-2004, now Ph.D. student at Caltech, Computation and Neural Systems.

Anne Chen, 2002-2004.

Brigid Dwyer, 2003-2004.

Carrie Niziolek, 2003-2005, now Ph.D. student at MIT, Speech and Hearing Sciences.

Izzat Jarudi, 2004, now Ph.D. student at Yale University, Psychology.

George Marzloff, 2004.

Suzanne Luther, 2004.

Ben Yu, 2004.

Allison Berke, 2005.
Gabe Chan, 2006.
Laila Shabir, 2006.
Lacey Kitch, 2006.
Matt Gordon, 2006.
Sadik Antwi-Boampong, 2006.
Stephanie Brenman, 2006.
Mark Tobenkin, 2006.
Pooja Jotwani, 2006.
Anna Schapiro, 2006.

High School Research Students

Rishi Mukhopadhyay, 1999-2001.
Genya Frenkel, 2004.
Aaron Bernstein, 2004.