

# Michael C. Mozer

|                                   |  |
|-----------------------------------|--|
| <b>ADDRESS</b>                    | Google DeepMind<br>1600 Amphitheater Parkway<br>Mountain View, CA 94043, USA   |
|                                   | Department of Computer Science and Institute of Cognitive Science<br>University of Colorado<br>Boulder, CO 80309-0430, USA   |
| <b>EMAIL</b>                      | mcmozer@google.com<br>mozer@colorado.edu   |
| <b>PHONE</b>                      | +1 (303) 517-2777  |
| <b>WWW</b>                        | <a href="https://www.cs.colorado.edu/~mozer">https://www.cs.colorado.edu/~mozer</a>  |
| <b>BIRTHDATE</b>                  | November 20, 1958  |
| <b>CITIZENSHIP</b>                | USA  |
| <b>RESEARCH INTERESTS</b>         | <p>I am committed to <i>human-centric artificial intelligence (AI)</i>—AI that mimics and enhances human capabilities, understands and anticipates an individual’s needs, and acts in synergistic coordination with individuals. My work explores the topics of:</p> <ul style="list-style-type: none"><li>▪ <i>Cognitively informed artificial intelligence</i>: Incorporating insights from human perception and cognition into the design of AI architectures and machine learning methods. For example, I have proposed recurrent neural network models motivated by properties of human long-term memory.</li><li>▪ <i>Human optimization</i>: Developing software tools to improve how people learn, remember, and make decisions. Much of my present work is aimed at determining the most effective means of teaching and the manner in which to best present information for human consumption. For example, we created the Colorado Optimized Language Tutor, which helps students learn facts (e.g., foreign language vocabulary) by scheduling review to promote long-term retention.</li><li>▪ <i>Cognitive modeling</i>: Building psychologically grounded models of human cognition that allow us to predict and understand behavior. I have worked in the areas of selective attention, awareness, memory, learning, executive control, decision making, and neuropsychological disorders.</li><li>▪ <i>Intelligent environments</i>: Designing computer interfaces that are smarter, anticipatory, and easier to use. A past project that achieved some notoriety was the <i>adaptive house</i>, a control system that learns to manage energy resources (air heat, water heat, lighting, and ventilation) in an actual residence to to maximize the satisfaction of the inhabitants and minimize energy consumption.</li></ul> |
| <b>EDUCATION</b>                  | Ph.D. University of California, San Diego (Psychology and Cognitive Science) 1987<br>M.A. University of California, San Diego (Psychology) 1982<br>B.A. Brown University (Computer Science) 1980   |
| <b>ACADEMIC HONORS AND AWARDS</b> | Fellow, Cognitive Science Society 2017<br>Educational Data Mining Society, Best Paper Award 2016<br>Educational Data Mining Society. Best Paper Award 2014<br>Cognitive Science Society, Computational Modeling Prize 2013<br>Faculty Fellowship, University of Colorado, Boulder 2009–2010<br>Distinguished Cognitive Scientist Award, Glushko-Samuelson Foundation, UC Merced 2010<br>Faculty Fellowship, University of Colorado, Boulder 1995–1996<br>Presidential Young Investigator Award, National Science Foundation 1990<br>Junior Faculty Development Award, University of Colorado, Boulder 1989<br>IBM Graduate Fellowship 1985–1987<br>Institute for Cognitive Science SDF Graduate Fellowship, UCSD 1981–1985<br>Sigma Xi (honorary scientific society), Brown University chapter 1980  |

|                                |   |              |
|--------------------------------|---|--------------|
|                                | Phi Beta Kappa, Brown University chapter  | 1980         |
|                                | B.A. degree <i>summa cum laude</i> , Brown University   | 1980         |
| <b>PROFESSIONAL EXPERIENCE</b> | Senior Staff Research Scientist, Google Brain, Mountain View, CA  | 2019–present |
|                                | Professor, Department of Computer Science and Institute of Cognitive Science,<br>University of Colorado, Boulder  | 2001–present |
|                                | Visiting Faculty Researcher, Google Brain, Mountain View, CA  | 2018–2019    |
|                                | Associate Professor, Department of Computer Science and Institute of Cognitive Science,<br>University of Colorado, Boulder  | 1992–2001    |
|                                | Assistant Professor, Department of Computer Science and Institute of Cognitive Science,<br>University of Colorado, Boulder  | 1988–1992    |
|                                | Lecturer, Department of Psychology, University of Toronto   | 1987–1988    |
|                                | Postdoctoral Fellow, Departments of Psychology and Computer Science,<br>University of Toronto, Dr. Geoffrey Hinton, Supervisor  | 1987–1988    |
| <b>PROFESSIONAL ACTIVITIES</b> | Advisory Board, WootMath (educational software), Boulder, CO  | 2018–present |
|                                | Advisory Board, NSF/Cyberlearning project on “Software for using Collaborative,<br>Dynamic, Personalized Experimentation to Investigate & Enhance Education”,<br>Worcester Polytechnic Institute    | 2018–present |
|                                | Technical Advisory Board, Drop (AI assisted food preparation), San Francisco, CA  | 2018–present |
|                                | Advisory Board, NSF/Cyberlearning project on “Modeling perceptual fluency<br>with visual representations in an intelligent tutoring system for undergraduate chemistry”,<br>University of Wisconsin | 2016–present |
|                                | Editorial Board, Springer series on <i>Applied Machine Learning</i>   | 2016–present |
|                                | Technical Advisory Board, Imagen Technologies (medical imagery analysis), New York  | 2015–present |
|                                | Editorial Board, <i>Neural Computation</i>  | 2015–present |
|                                | Technical Advisory Board, AnswerOn (churn prediction), Longmont, CO   | 2001–present |
|                                | Board Member and Secretary, Neural Information Processing Systems Foundation  | 1995–present |
|                                | Award Committee, Cognitive Science Society  | 2019         |
|                                | Workshop Co-organizer, <i>Deep Learning for Education</i> . KDD 2019. Anchorage, AK   | Aug 2019     |
|                                | Co-Founder and Technical Advisory Board Member, Sensory Inc. (embedded<br>speech and vision technology), Sunnyvale CA   | 1992–2018    |
|                                | Advisory Board, NSF/Integrative Strategies project on “Using computational<br>cognitive neuroscience to predict and optimize memory”, New York University   | 2016–2019    |
|                                | Technical Advisory Board, Open Table, San Francisco, CA   | 2016–2018    |
|                                | Workshop Co-organizer. <i>Cognitively Informed Artificial Intelligence: Lessons from<br/>Natural Intelligence</i> . Neural Information Processing Systems. Long Beach, CA                           | Dec 2017     |
|                                | Faculty, <i>International Summer School on Deep Learning</i> , Bilbao, Spain  | Jul 2017     |
|                                | Symposium Co-organizer, <i>Enhancing Education Through Cognitive Psychology</i> .<br>Psychonomics 2015. Chicago, IL   | Nov 2015     |
|                                | Workshop Co-organizer, <i>Machine Learning For Education</i> . ICML 2015. Lille, France   | Jul 2015     |
|                                | Technical Advisory Board, Cognilytics, Inc., Denver, CO   | 2011–2015    |
|                                | Workshop Co-organizer, <i>Human Propelled Machine Learning</i> . NIPS 2014. Montreal, Canada  | Dec 2014     |
|                                | Workshop Co-organizer, <i>Approaching Twenty Years of Knowledge Tracing:<br/>Lessons Learned, Open Challenges, and Promising Developments</i> . EDM 2014. London, UK                                | Jul 2014     |
|                                | Workshop Co-organizer, <i>Personalizing Education With Machine Learning</i> , NIPS 2012.<br>Lake Tahoe, CA  | Dec 2012     |
|                                | Faculty, International Summer School in Cognitive Science, Sofia, Bulgaria  | Jul 2012     |
|                                | Technical Advisory Board, J.D. Powers and Associates, Web Intelligence Division<br>(formerly Umbria Communications), Boulder, CO  | 2003–2010    |
|                                | Editorial Board, <i>Machine Learning</i>  |              |
|                                | Chair, Finance Committee, Cognitive Science Society   | 2005–2009    |
|                                | Board of Governors, Cognitive Science Society   | 1998–2008    |
|                                | Technical Advisory Board, Green Planet Software   | 2001–2008    |
|                                | Executive Committee, Cognitive Science Society  | 2005–2008    |
|                                | Conference Liaison, Cognitive Science Society   | 2008         |
|                                | Chair, Cognitive Science Society  | 2006–2007    |

|   |           |
|---|-----------|
| Symposium Co-Organizer, <i>Emergent Cognitive Control</i> , Cognitive Neuroscience Conference   | Nov 2006  |
| Editorial Board, <i>Consciousness and Cognition</i>   | 1998–2005 |
| Advisory Board, Series on Natural Computing, Springer-Verlag  | 1998–2005 |
| Editorial Board, <i>Visual Cognition</i>  | 1992–2004 |
| Faculty, International Summer School in Cognitive Science, Sofia, Bulgaria  | Jul 2002  |
| Editorial Board, <i>Neural Networks</i>   | 1994–2001 |
| Symposium Co-organizer, <i>Computational Neuropsychology</i> , Neural Information Processing Systems Conference   | Dec 2001  |
| Chief Scientist, Athene Software, Boulder, CO   | 1998–2001 |
| Tutorials Chair, Neural Information Processing Systems Conference   | Dec 2000  |
| Symposium Co-organizer, <i>Bayesian approaches to cognitive modeling</i> , Cognitive Science Conference   | Aug 2000  |
| Editorial Board, <i>Cognitive Science</i>   | 1999–2000 |
| Advisory Board, <i>Connectionist Surveys</i>  | 1996–2000 |
| Symposium Organizer, <i>Principles of computation in the brain</i> , Cognitive Neuroscience Conference  | Apr 1998  |
| Workshop Co-organizer, <i>Interfacing models of language</i> , Neural Information Processing Systems  | Dec 1997  |
| Co-Editor, Special issue of <i>Neurocomputing</i> on recurrent networks,  | 1997      |
| Consultant, Lifestyle Technologies, Los Angeles, California,  | 1995–1997 |
| General Chair, Neural Information Processing Systems Conference   | Dec 1996  |
| Program Chair, Neural Information Processing Systems Conference   | Nov 1995  |
| Faculty, James S. McDonnell Foundation Summer Institute in Cognitive Neuroscience,  | Jul 1995  |
| Workshop Chair, Neural Information Processing Systems   | Nov 1993  |
| Co-Organizer, Connectionist Models Summer School  | Jun 1993  |
| Local Arrangements Chair, Neural Information Processing Systems Conference  | Nov 1991  |
| Faculty, James S. McDonnell Foundation Summer Institute in Cognitive Neuroscience   | Jul 1991  |
| Participant, James S. McDonnell Foundation Summer Institute in Cognitive Neuroscience   | Jun 1988  |
| Research Assistant, Cognitive Science Laboratory, UCSD  | 1981–1987 |
| Teaching Assistant, Department of Psychology, UCSD  | 1981–1987 |
| Programmer/Research Assistant, Electronic Speech Systems, Santa Clara, California   | 1975–1987 |
| Participant, Connectionist Models Summer School   | Jun 1986  |
| Visiting Scholar, Department of Computer Science, Carnegie-Mellon University  | 1984–1985 |
| Editorial Assistant to Diana Deutsch, Editor, <i>Music Perception</i>   | 1983–1984 |
| Teaching Assistant, Department of Computer Science, Brown University  | 1977–1980 |
| Occasional Reviewer for <i>Proceedings of the National Academy of Sciences</i> , <i>Neural Information Processing Systems Conference</i> , <i>Cognitive Science Society Conference</i> , <i>Journal of Cognitive Neuroscience</i> , <i>IEEE Transactions on Neural Networks</i> , <i>Neural Computation</i> , <i>Connection Science</i> , <i>Artificial Intelligence</i> , <i>Cognitive Science</i> , <i>Cognitive Psychology</i> , <i>Cognitive Neuropsychology</i> , <i>Psychological Science</i> , <i>Consciousness and Cognition</i> , <i>Neurocomputing</i> , <i>Neuropsychologia</i> , <i>Neural Networks</i> , <i>Psychological Review</i> , <i>Journal of Experimental Psychology</i> , <i>Canadian Journal of Psychology</i> , <i>Quarterly Journal of Experimental Psychology</i> , <i>Psychological Research</i> , NSF, AFOSR, NSERC |           |
| <b>SPONSORED RESEARCH</b>   |           |
| “Predictive models of human memory”, Unnamed Corporate Sponsor, \$209,000   | 2018–2019 |
| “Operationalizing students’ textbook annotations to improve comprehension and long-term retention”, NSF IIS (NCS-FO), \$1,000,000 (my share \$300,000)  | 2016–2020 |
| “Bayesian optimization for exploratory experimentation in the behavioral sciences”, NSF SES, \$400,000  | 2015–2019 |
| REU Supplement, NSF IIS, \$7,200  | 2017–2018 |
| REU Supplement, NSF SES, \$2,500  | 2016      |
| “Aphasia rehabilitation: Modulating cues, feedback, and practice conditions” (L. Cherney and S. Van Vuuren, PIs), NIH, \$3.1M (my share \$50,000)   | 2011–2016 |
| “Temporal dynamics of human learning and memory” (Garrison Cottrell, PI), TDLC Science of Learning Center, National Science Foundation, \$15M (my share ~\$500k)  | 2006–2016 |
| “Context-Award Music Recommendation”, Samsung, \$11,400   | 2014–2015 |

“Improving memory retention via spacing of practice: Computational and empirical investigations” (Harold Pashler, co-PI), National Science Foundation, BCS, \$450,000 total (my share \$224,977) 2007–2010

“Understanding the performance of modern systems” (Amer Diwan, PI), National Science Foundation, SMA, \$400,000 total (my share \$200,000) 2005–2009

“Control and adaptation of attentional processing: Empirical and computational investigations” (Shaun Vecera, Co-PI), National Science Foundation, Human and Social Dynamics Program, \$430,000 total (my share \$240,000) 2004–2008

“Enhancing learning through testing: Theoretical and practical issues” (subcontract to University of Colorado, Michael Mozer, PI; Harold Pashler, overall PI), National Institute of Health, \$450,000 total (my share \$119,510) 2000–2004

“Discrete representations in working memory: Developmental, neuropsychological, and computational investigations” (Randy O’Reilly, Yuko Munakata, Akira Miyake, Co-PIs), National Science Foundation, Knowledge and Distributed Intelligence Program, \$800,000 total (my share \$200,000) 1998–2002

“Temporal dynamics of cognition in a modular cortical architecture”, McDonnell-Pew Program in Cognitive Neuroscience, \$105,000 1997–2000

“Artificial Intelligence and Home Automation”, Lifestyle Technologies, \$40,000 1997

“Rapid Behavioral Tuning to Task Demands: Computational Modeling of Empirical Data” (Clark Fagot, Co-PI), McDonnell-Pew Program in Cognitive Neuroscience, \$90,000 1994–1996

REU Supplement, National Science Foundation, \$23,000 1993–1995

“Connectionist Modeling and Cognitive Neuroscience”, James S. McDonnell Foundation, \$187,500 1990–1995

Presidential Young Investigator Award, National Science Foundation, \$312,500 1990–1995

“Connectionist Models Summer School”, American Association for Artificial Intelligence, National Science Foundation, and Siemens Research Center, \$45,000 1993

CRCW Grant In Aid, University of Colorado, \$3,000 1992

Digital Equipment Corporation External Research Grant, \$45,750 1991

Junior Faculty Development Award, University of Colorado, \$5,000 1989

“Connectionist Models of Selective Attention and Object Recognition”, James S. McDonnell Foundation, \$9,000 1988–1989

**U.S. PATENTS**

“A Speech Recognition Apparatus For Consumer Electronic Applications”, Forrest Mozer, Michael Mozer, and Todd Mozer. Submitted September 1994; issued August 4, 1998. US Patent 5,790,754.

“Speech Recognition in Consumer Electronic Products”, Todd Mozer, Michael Mozer, and Forrest Mozer. Issued February 1, 2000. US Patent 6,021,387.

“Parallel Cascaded Neural Networks”, Michael Mozer, Michael Iuzzolino, and Samy Bengio. August 19, 2022. US Patent Application 20220253695.

“Online training of machine learning models using Bayesian inference over noise”, M Jones, MC Mozer. US Patent App. 18/477,525

**GRADUATE STUDENTS SUPERVISED**

Aniket Didolkar, Ph.D. 2025 (expected), University of Montreal (co-advisor with Y. Bengio, Anirudh Goyal)

Camden Elliott-Williams, Ph.D. 2024. *Facilitating human perception via task-specific image modulations*

Shruthi Sukumar, Ph.D. 2023. *On the determinants of the history-dependent modulation of motor control and decision making* (co-advisor with Alaa Ahmed)

Tyler Scott, Ph.D. 2023. *Deep visual representation learning for classification and retrieval: Uncertainty, geometry, and applications* (co-advisor with Elizabeth Bradley)

Michael Iuzzolino, Ph.D. 2021. *The role of time in machine perception* (co-advisor with Clayton Lewis)

David Kim, M.S. 2021. *Using semantics of textbook highlights to predict student comprehension and knowledge retention*

Adam Winchell, M.S. 2019. *Textbook annotations as predictors of student learning*

Shirly Montero Quesada, M.S. 2019. *Modeling interactions among learned skills*

Matt Maierhofer, M. S. 2019. *Lifetime limited memory networks*

Karl Ridgeway, Ph.D. 2018. *Content-style decomposition: Representation discovery and applications*

Denis Kazakov, M.S. 2018. *State-denoised recurrent neural networks*

Brennan McConnell, M.S. 2018. *A calibration loss for neural networks*

Aditya Thyagarajan, M.S. 2018. *Convolutional attractor networks for superresolution*

Brett Roads, Ph.D., 2017. *Accelerating human visual concept learning and boosting performance via computational models of perception and cognition*

Shruthi Sukumar, M.S. 2017. *Analysis and solution of Markov decision problems with a continuous, stochastic state component*

Mohammad Khajah, Ph.D. 2017. *Optimizing game engagement via nonparametric models and manipulations of difficulty, tension, and perceived performance*

Ron Kneusel, Ph.D. 2015. *Improving hybrid human-machine search through soft highlighting*

Robert Lindsey, Ph.D. 2014. *Probabilistic models of student learning and forgetting*

Karl Ridgeway, M.S. 2014. *Forgetting of foreign language skills: A corpus based analysis of Rosetta Stone®*

Ahbishek Jainantilil, Ph.D. 2013. *Feature selection via iterative reweighting: An exploration of algorithms for linear models and random forests*

Brett Roads, M.S. 2013. *Using attentional highlighting to train visual expertise.*

Matthew Wilder, Ph.D. 2012. *Probabilistic modeling of sequential effects in human behavior: Theory and practical applications.*

Daniel Knights, Ph.D. 2012. *Predictive modeling of metagenomes* (co-advised with Robin Knights, received College of Engineering Outstanding Dissertation Award, #1 of 390)

Benjamin Link, M.S. 2011. *Modeling the effect of recent experience on judgments.*

Karthik Venkatesh, M.S. 2010 (Electrical and Computer Engineering). *Experience guided search: A theory of attentional control.*

Owen Lewis, M.S. 2010 (Applied Math). *A review of mathematical techniques in machine learning.*

Samuel Reid, Ph.D. 2010. *Model combination in multiclass classification.*

Adrian Fan, M.S. 2008. *A synthesis of theoretical and empirical perspectives on repetition suppression.*

Scott Richardson, M.S. 2007. *Discovering the runtime structure of software with probabilistic generative models.*

Thomas Borchert, M.S. 2007. *Computational correlates of access consciousness.*

Brian Loughery, M.S. 2003. *Learning working memory tasks by reward prediction in the basal ganglia and prefrontal cortex* (co-advisor with Randall O'Reilly)

Michael Colagrosso, Ph.D. 2003. *A rational theory of skilled performance and practice: Modeling long-term repetition priming.*

David Nix, Ph.D. 1998. *Machine learning methods for inferring vocal-tract articulation from speech acoustics*

Torleif Mohling, M.S., 1998. *Predicting human performance on anagram solving: A computational model*

Donald Mathis, Ph.D., 1997. *A computational theory of consciousness in cognition*

Srecko Vidmar, M.S., 1997. *Optimal control of home heating systems*

Kelvin Fedrick, M.S., 1996. *A decompositional approach to time series forecasting*

Debra Miller, M.S., 1995. *Adaptive lighting control*

Kevin Markey, Ph.D., 1994. *The sensorimotor foundations of phonology: A computational model of early childhood articulatory and phonetic development*

Sreerupa Das, Ph.D., 1994. *Connectionist models of language induction incorporating symbolic constraints*

John Allison, M.S., 1994. *Explorations of Bayesian input relevance determination for neural networks*

Jay Alexander, M.S., 1993. *Template-based procedures for neural network interpretation*

Ken Parker, M.S., 1993. *Selecting regression estimators for the generalized ensemble method*

Clayton McMillan, Ph.D., 1992. *Rule induction in a neural network through integrated symbolic and subsymbolic processing*

Stefanie Lindstaedt, M.S., 1992. *Comparison of unsupervised neural network models for redundancy reduction*

**BOOKS AND  
EDITED VOLUMES**

Mozer, M. C. (1991). *The perception of multiple objects: A connectionist approach*. Cambridge, MA: MIT Press/Bradford Books.

Mozer, M. C., Smolensky, P., Touretzky, D. S., Elman, J. L., & Weigend, A. S. (Eds.). (1994). *Proceedings of the 1993 Connectionist Models Summer School*. Hillsdale, NJ: Erlbaum.

Smolensky, P., Mozer, M. C., & Rumelhart, D. E. (Eds.). (1996). *Mathematical perspectives on neural networks*. Hillsdale, NJ: Erlbaum.

Touretzky, D. S., Mozer, M. C., & Hasselmo, M. (Eds.). (1996). *Neural Information Processing Systems 8*. Cambridge, MA: MIT Press.

- Mozer, M. C., Jordan, M. I., & Petsche, T. (Eds.). (1997). *Neural Information Processing Systems 9*. Cambridge, MA: MIT Press.
- Mozer, M. C. (1983). Letter migration in word perception. *Journal of Experimental Psychology: Human Perception and Performance*, 9, 531–546.
- McClelland, J. L., & Mozer, M. C. (1986). Perceptual interactions in multi-word displays: Familiarity and similarity effects. *Journal of Experimental Psychology: Human Perception and Performance*, 12, 18–35.
- Mozer, M. C. (1989). Types and tokens in visual letter perception. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 287–303.
- Mozer, M. C. (1989). A focused back-propagation algorithm for temporal sequence recognition. *Complex Systems*, 3, 349–381.
- Mozer, M. C., & Smolensky, P. (1989). Using relevance to reduce network size automatically. *Connection Science*, 1, 3–16.
- Mozer, M. C., & Behrmann, M. (1990). On the interaction of spatial attention and lexical knowledge: A connectionist account of neglect dyslexia. *Cognitive Neuroscience*, 2, 96–123.
- Behrmann, M., Moscovitch, M., Black, S. E., & Mozer, M. C. (1990). Perceptual and conceptual mechanisms in neglect dyslexia: Two contrasting case studies. *Brain*, 113, 1163–1183.
- Mozer, M. C., & Bachrach, J. (1990). Discovering the structure of a reactive environment by exploration. *Neural Computation*, 2, 447–457.
- Behrmann, M., Moscovitch, M., & Mozer, M. C. (1991). Directing attention to words and nonwords in normal subjects and in a computational model: Implications for neglect dyslexia. *Cognitive Neuropsychology*, 8, 213–248.
- Mozer, M. C., & Bachrach, J. (1991). SLUG: A connectionist architecture for inferring the structure of finite-state environments. *Machine Learning*, 7, 139–160.
- Behrmann, M., & Mozer, M. C. (1992). A connectionist account of neglect dyslexia. *Journal of Clinical and Experimental Neuropsychology*, 14, 48–49.
- Mozer, M. C., Zemel, R. S., Behrmann, M., & Williams, C. K. I. (1992). Learning to segment images using dynamic feature binding. *Neural Computation*, 4, 650–665.
- Dodier, R. H., Lukianow, D., Ries, J., & Mozer, M. C. (1994). Comparison of neural net and conventional techniques for lighting control. *Applied Mathematics and Computer Science*, 4, 447–462.
- Mozer, M. C. (1994). Neural network music composition by prediction: Exploring the benefits of psychophysical constraints and multiscale processing. *Connection Science*, 6, 247–280.
- Zemel, R. S., Williams, C. K. I., & Mozer, M. C. (1995). Lending direction to neural networks. *Neural Networks*, 8, 503–512.
- Mozer, M. C. (1996). Neural network speech processing for toys and consumer electronics. *IEEE Expert*, 11, 4–5.
- Calder, B., Grunwald, D., Jones, M., Lindsay, D., Martin, J., Mozer, M., & Zorn, B. (1997). Evidence-based static branch prediction using machine learning. *Transactions on Programming Languages and Systems*, 19, 188–222. [Authorship order is alphabetical.]
- Mozer, M. C., Halligan, P. W., Marshall, J. C. (1997). The end of the line for a brain-damaged model of unilateral neglect. *Journal of Cognitive Neuroscience*, 9, 171–190.
- Das, S., & Mozer, M. C. (1998). Dynamic on-line clustering and state extraction: An approach to symbolic learning. *Neural Networks*, 11, 53–64.
- Behrmann, M., Zemel, R. S., and Mozer, M. C. (1998). Object-based attention and occlusion: Evidence from normal subjects and a computational model. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 1011–1036.
- Alexander, J. A., & Mozer, M. C. (1999). Template-based procedures for neural network interpretation. *Neural Networks*, 12, 479–498.
- Mozer, M. C. (1999). An intelligent environment should be adaptive. *IEEE Intelligent Systems and their Applications*, 14(2), 11–13.
- Behrmann, M., Zemel, R. S., & Mozer, M. C. (2000). Occlusion, symmetry, and object-based attention: Reply to Saiki (1999). *Journal of Experimental Psychology: Human Perception and Performance*, 26, 1497–1505.
- Mozer, M. C., Wolniewicz, R., Grimes, D., Johnson, E., & Kaushansky, H. (2000). Maximizing revenue by predicting and addressing customer dissatisfaction. *IEEE Transactions on Neural Networks*, 11, 690–696.

- Sitton, M., Mozer, M. C., & Farah, M. J. (2000). Superadditive effects of multiple lesions in a connectionist architecture: Implications for the neuropsychology of optic aphasia. *Psychological Review*, *107*, 709–734.
- Zemel, R. S., & Mozer, M. C. (2001). Localist attractor networks. *Neural Computation*, *13*, 1045–1064.
- Mozer, M. C. (2002). Frames of reference in unilateral neglect and visual perception: A computational perspective. *Psychological Review*, *109*, 156–185.
- Pashler, H., Mozer, M. C., & Harris, C. R. (2002). Mating strategies in a Darwinian microworld: Simulating the consequences of female reproductive refractoriness. *Adaptive Behavior*, *9*, 5–15.
- Zemel, R. S., Behrmann, M., & Mozer, M. C. (2002). Experience-dependent perceptual grouping and object-based attention. *Journal of Experimental Psychology: Human Perception and Performance*, *28*, 202–217.
- Kinoshita, S., & Mozer, M. C. (2006). How lexical decision is affected by recent experience: Symmetric versus asymmetric frequency blocking effects. *Memory and Cognition*, *34*, 726–742.
- Bohte, S.M., & Mozer, M. C. (2007). Reducing the variability of neural responses: A computational theory of spike-timing dependent plasticity. *Neural Computation*, *19*, 371–403.
- Kinoshita, S., Forster, K. I., & Mozer, M. C. (2008). Unconscious cognition isn't that smart: Modulation of masked repetition priming effect in the word naming task. *Cognition*, *107*, 623–649.
- Mozer, M. C., & Fan, A. (2008). Top-down modulation of neural responses in visual perception: A computational exploration. *Natural Computing*, *7*, 45–55.
- Mozer, M. C., Pashler, H., & Homaei, H. (2008). Optimal predictions in everyday cognition: The wisdom of individuals or crowds? *Cognitive Science: A Multidisciplinary Journal*, *32*, 1133–1147.
- Cepeda, N. J., Coburn, N., Rohrer, D., Wixted, J. T., Mozer, M. C., & Pashler, H. (2009). Optimizing distributed practice: Theoretical analysis and practical implications. *Experimental Psychology*, *56*, 236–246.
- Lee, H., Mozer, M. C., & Vecera, S. (2009). Mechanisms of priming of pop out: Stored representations or feature gain modulations? *Attention, Perception, & Psychophysics*, *71*, 1059–71.
- Kang, S. H. K., Pashler, H., Cepeda, N. J., Rohrer, D., Carpenter, S. K., & Mozer, M. C. (2011). Does incorrect guessing impair fact learning? *Journal of Educational Psychology*, *103*, 48–59.
- Kinoshita, S., Mozer, M. C., & Forster, K. I. (2011). Dynamic adaptation to history of trial difficulty explains the effect of congruency proportion on masked priming. *Journal of Experimental Psychology: General*, *140*, 622–636.
- Knight, D., Kuczynski, J., Charlson, E., Zaneveld, J., Collman, R. G., Mozer, M. C., Bushman, F. D., Knight, R., & Kelley, S. T. (2011). Bayesian community-wide culture-independent microbial source tracking. *Nature Methods*, *8*, 761–763.
- Wilder, M. H., Mozer, M. C., & Wickens, C. D. (2011). An integrative experience-based theory of attentional control. *Journal of Vision*, *11*, 1–30.
- Doshi, A., Tran, C., Wilder, M., Mozer, M. C., & Trivedi, M. (2012). Sequential effects in driving. *Cognitive Science*, *36*, 948–963.
- Lee, H., Mozer, M. C., Kramer, A. F., & Vecera, S. P. (2012). Object-based control of attention is sensitive to recent experience. *Journal of Experimental Psychology: Human Perception and Performance*, *38*, 314–325.
- Chukoskie, L., Snider, J., Mozer, M. C., Krauzlis, R. J., & Sejnowski, T. J. (2013). Learning where to look: An empirical, computational, and theoretical account of hidden target search performance. *Proceedings of the National Academy of Sciences*, *110*, 10438–10445.
- Jones, M., Curran, T., Mozer, M. C., & Wilder, M. H. (2013). Sequential effects in response time reveal learning mechanisms and event representations. *Psychological Review*, *120*, 628–666.
- Pashler, H., Kang, S., & Mozer, M. C. (2013). Reviewing erroneous information facilitates memory updating. *Cognition*, *128*(3), 424–430.
- Pashler, H., & Mozer, M. C. (2013). When does fading help perceptual category learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*, 1162–73.
- Wilder, M. H., Jones, M., Ahmed A. A., Curran, T., & Mozer, M. C. (2013). The persistent impact of incidental experience. *Psychonomic Bulletin and Review*, *20*, 1221–1231.
- Khajah, M. M., Lindsey, R. V., & Mozer, M. C. (2014). Maximizing students' retention via spaced review: Practical guidance from computational models of memory. *Topics in Cognitive Science*, *6*, 157–169.
- Lindsey, R., Shroyer, J., Pashler, H., & Mozer, M. C. (2014). Improving long-term knowledge retention through personalized review. *Psychological Science*, *25*, 639–647. doi: 10.1177/0956797613504302.

- Kang, S. H. K., Lindsey, R. V., Mozer, M. C., & Pashler, H. (2014). Retrieval practice over the long term: Should spacing be expanding or equal-interval? *Psychonomic Bulletin & Review*, *21*, 1544–50.
- Roads, B. D., Mozer, M. C., & Busey, T. A. (2016). Using highlighting to train attentional expertise. *PLoS ONE* *11*(1): e0146266. doi:10.1371/journal.pone.0146266
- Ridgeway, K., Mozer, M. C., & Bowles, A. (2017). Forgetting of foreign language skills: A corpus-based analysis of online tutoring software. *Cognitive Science: A Multidisciplinary Journal*, *41*, 924–949. doi:10.1111/cogs.12385.
- Roads, B. D., & Mozer, M. C. (2017). Improving human-computer cooperative classification via cognitive theories of similarity. *Cognitive Science: A Multidisciplinary Journal*, *41*, 1394–1411. doi:10.1111/cogs.12400
- Kneusel, R. T., & Mozer, M. C. (2017). Improving human-machine cooperative visual search with soft highlighting. *ACM Transactions on Applied Perception*, *15*, 3:1–3:21. Also arXiv:1612.08117 [cs.HC]
- Vatterott, D. B., Mozer, M. C., & Vecera, S. P. (2018). Rejecting salient distractors: Generalization from experience. *Attention, Perception, & Psychophysics*, *80*, 485–499. DOI:10.3758/s13414-017-1465-8.
- Lindsey, R., Daluski, A., Chopra, S., Lachapelle, A., Mozer, M. C., Sicular, S., Hanel, D., Gardner, M., Gupta, A., Hotchkiss, R., & Potter, H. (2018). A deep neural network improves fracture detection by clinicians. *Proceedings of the National Academy of Sciences*, *115*, 11591–11596.
- Mozer, M. C., Wiseheart, M., & Novikoff, T. P. (2019). Artificial intelligence to support human instruction. *Proceedings of the National Academy of Sciences*, *116* (10), 3953–3955.
- Roads, B. D., & Mozer, M. C. (2019). Obtaining psychological embeddings through joint kernel and metric learning. *Behavioral Research Methods*, *51*, 2180–2193. doi:10.3758/s13428-019-01285-3.
- Beckage, N., Colunga, E., & Mozer, M. C. (2020). Quantifying the role of vocabulary knowledge in predicting future word learning. *IEEE Transactions on Cognitive and Developmental Systems*, *12*, 148–159. DOI:10.1109/TCDS.2019.2928023
- Winchell, A., Lan, A., & Mozer, M. C. (2020). Textbook annotations as predictors of student learning. *Cognitive Science: An Interdisciplinary Journal*, *44*: e12901. <https://doi.org/10.1111/cogs.12901>
- Kim, B., Reif, E., Watenberg, M., Bengio, S., & Mozer, M. C. (2021). Neural networks trained on natural scenes exhibit closure. *Computational Brain and Behavior*, *4*(3), 251–263.
- Roads, B. D., & Mozer, M. C. (2021). Predicting the ease of human category learning using radial basis function networks. *Neural Computation*, *33*, 376–397.
- Mozer, F. S., Bale, S. D., Bonnell, J. W., Drake, J. F., Hanson, E. L. M., & Mozer, M. C. (2021). On the origin of switchbacks observed in the solar wind. *Journal of Astrophysics*, *919*, 60–69.
- Ke, N. R., Bilaniuk, O., Goyal, A., Bauer, S., Larochelle, H., Schoelkopf, B., Mozer, M. C., Pal, C. Bengio, Y. (2023). Neural causal structure discovery from interventions. *Transactions on Machine Learning Research*.
- Veerabadran, V., Goldman, J., Shankar, S., Cheung, B., Papernot, N., Kurakin, A., Goodfellow, I., Shlens, J., Sohl-Dickstein, J., Mozer, M. C., & Elsayed, G. F. (2023). Subtle adversarial image manipulations influence both human and machine perception. *Nature Communications*, *14*, 4933. <https://doi.org/10.1038/s41467-023-40499-0>
- Sukumar, S., Elliott-Williams, C., Hakimi, S., Ward, A. F., & Mozer, M. C. (submitted). Overcoming temptation: Incentive design for intertemporal choice. arXiv.org:2203.05782 [cs.LG / q-bio.NC]. Submitted for publication.

**REFEREED  
CONFERENCE  
PUBLICATIONS**

- Mozer, M. C. (1987). Early parallel processes in reading: A connectionist approach. In M. Coltheart (Ed.), *Attention and performance XII: The psychology of reading* (pp. 83–104). Hillsdale, NJ: Erlbaum.
- Mozer, M. C. (1987). RAMBOT: A connectionist expert system that learns by example. In M. Caudillo & C. Butler (Eds.), *Proceedings of the IEEE First Annual International Conference on Neural Networks* (pp. 693–700). San Diego: IEEE Publishing Services.
- Mozer, M. C. (1988). A connectionist model of selective attention in visual perception. *Proceedings of the Tenth Annual Conference of the Cognitive Science Society* (pp. 195–201). Hillsdale, NJ: Erlbaum.
- Mozer, M. C., & Smolensky, P. (1989). Skeletonization: A technique for trimming the fat from a network via relevance assessment. In D. Touretzky (Ed.), *Advances in Neural Information Processing Systems I* (pp. 107–115). San Mateo, CA: Morgan Kaufmann.
- Mozer, M. C. (1990). Discovering faithful “Wickelfeature” representations in a connectionist network. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society* (pp. 356–363). Hillsdale, NJ: Erlbaum.



- Mozer, M. C., & Bachrach, J. (1990). Discovering the structure of a reactive environment by exploration. In D. Touretzky (Ed.), *Advances in neural information processing systems II* (pp. 439–446). San Mateo, CA: Morgan Kaufmann.
- Zemel, R. S., Mozer, M. C., & Hinton G. E. (1990). TRAFFIC: Object recognition using hierarchical reference frame transformations. In D. Touretzky (Ed.), *Advances in neural information processing systems II* (pp. 266–273). San Mateo, CA: Morgan Kaufmann.
- Mozer, M. C., & Soukup, T. (1991). Algorithmic music composition with melodic and stylistic constraints. In R. P. Lippmann, J. Moody, and D. S. Touretzky (Eds.), *Advances in neural information processing systems III* (pp. 789–796). San Mateo, CA: Morgan Kaufmann.
- Mozer, M. C. (1991). Discovering discrete distributed representations with iterative competitive learning. In R. P. Lippmann, J. Moody, and D. S. Touretzky (Eds.), *Advances in neural information processing systems III* (pp. 627–634). San Mateo, CA: Morgan Kaufmann.
- McMillan, C., Mozer, M. C., & Smolensky, P. (1991). Learning explicit rules in a neural network. In *Proceedings of the International Joint Conference on Neural Networks, Volume II* (pp. 83–88). Piscataway, NJ: IEEE Publishing Services.
- McMillan, C., Mozer, M. C., & Smolensky, P. (1991). The connectionist scientist game: Rule extraction and refinement in a neural network. *Proceedings of the Thirteenth Annual Conference of the Cognitive Society* (pp. 424–430). Hillsdale, NJ: Erlbaum.
- Mozer, M. C. (1992). The induction of multiscale temporal structure. In J. E. Moody, S. J. Hanson, & R. P. Lippmann (Eds.), *Advances in neural information processing systems IV* (pp. 275–282). San Mateo, CA: Morgan Kaufmann.
- Mozer, M. C., Zemel, R. S., & Behrmann, M. (1992). Learning to segment images using dynamic feature binding. In J. E. Moody, S. J. Hanson, & R. P. Lippmann (Eds.), *Advances in neural information processing systems IV* (pp. 436–443). San Mateo, CA: Morgan Kaufmann.
- McMillan, C., Mozer, M. C., & Smolensky, P. (1992). Rule induction through integrated symbolic and subsymbolic processing. In J. E. Moody, S. J. Hanson, & R. P. Lippmann (Eds.), *Advances in neural information processing systems IV* (pp. 969–976). San Mateo, CA: Morgan Kaufmann.
- Markey, K., & Mozer, M. C. (1992). Comparison of reinforcement algorithms on learning discrete functions: Learnability, time complexity, and scaling. *Proceedings of the International Joint Conference on Neural Networks* (Volume I, pp. 853–859). San Diego, CA: IEEE Publishing Services.
- Mozer, M. C., Zemel, R. S., & Behrmann, M. (1992). Discovering and using perceptual grouping principles in visual information processing. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society* (pp. 283–288). Hillsdale, NJ: Erlbaum.
- Wagner, K., Mozer, M., Smolensky, P., Miyata, Y., & Fellows, M. (1992). Optical neural networks using a new radial nonlinear neural layer. *Proceedings of the SPIE, 1773A-10*.
- Bonnlander, B. V., & Mozer, M. C. (1993). Metamorphosis networks: An alternative to constructive methods. S. J. Hanson, J. D. Cowan, & C. L. Giles (Eds.), *Advances in Neural Information Processing Systems V* (pp. 131–138). San Mateo, CA: Morgan Kaufmann Publishers.
- Das, S., & Mozer, M. C. (1993). A connectionist symbol manipulator that induces rewrite rules in context-free grammars. *Proceedings of the first colloquium on grammatical inference: Theory, applications, and alternatives*. London: Institute of Electrical Engineers Digest 1993/092.
- Gupta, P., & Mozer, M. C. (1993). Exploring the nature and development of phonological representations. *Proceedings of the Fifteenth Annual Conference of the Cognitive Society* (pp. 516–521). Hillsdale, NJ: Erlbaum.
- McMillan, C., Mozer, M. C., & Smolensky, P. (1993). Implicit conflict resolution in a connectionist rule-based system. *Proceedings of the 13th Annual International Joint Conference on Artificial Intelligence* (pp. 1366–1371). San Mateo, CA: Morgan Kaufmann.
- Mozer, M. C., & Das, S. (1993). A connectionist symbol manipulator that discovers the structure of context-free languages. In S. J. Hanson, J. D. Cowan, & C. L. Giles (Eds.), *Advances in Neural Information Processing Systems V* (pp. 863–870). San Mateo, CA: Morgan Kaufmann Publishers.
- Schmidhuber, J. H., Mozer, M. C., & Prelinger, D. (1993). Continuous history compression. In H. Huening, S. Neuhauser, M. Raus, & W. Ritschel (Eds.), *Proceedings of the International Workshop on Neural Networks, RWTH Aachen* (pp. 87–95). Augustinus.
- Williams, C. K. I., Zemel, R. S., & Mozer, M. C. (1993). Unsupervised learning of object models. *Proceedings of the AAAI Fall 1993 Symposium on Machine Learning in Computer Vision* (pp. 20–24). Proceedings available as AAAI Tech Report FSS-93-04.

- Zemel, R. S., Williams, C. K. I., & Mozer, M. C. (1993). Directional-unit Boltzmann machines. In S. J. Hanson, J. D. Cowan, & C. L. Giles (Eds.), *Advances in Neural Information Processing Systems V* (pp. 172–179). San Mateo, CA: Morgan Kaufmann Publishers.
- Das, S., & Mozer, M. C. (1994). A unified gradient-descent/clustering architecture for finite-state machine induction. In J. D. Cowan, G. Tesauro, & J. Alspecter (Eds.), *Advances in Neural Information Processing Systems VI* (pp. 19–26). San Mateo, CA: Morgan Kaufmann Publishers.
- Mozer, M. C. (1994). Computational approaches to functional feature learning [Extended abstract]. *Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society* (pp. 975–976). Hillsdale, NJ: Erlbaum.
- Alexander, J. A., & Mozer, M. C. (1995). Template-based algorithms for connectionist rule extraction. In G. Tesauro, D. S. Touretzky, & T. K. Leen (Eds.), *Advances in Neural Information Processing Systems 7* (pp. 609–616). Cambridge, MA: MIT Press.
- Calder, B., Grunwald, D., Lindsay, D., Martin, J., Mozer, M., Zorn, B. (1995). Corpus-based static branch prediction. *ACM SIGPLAN '95 Conference on Programming Language Design and Implementation* (pp. 79–92). La Jolla, CA. June, 1995. [alphabetical authorship order]
- Mathis, D. A., & Mozer, M. C. (1995). On the computational utility of consciousness. In G. Tesauro, D. S. Touretzky, & T. K. Leen (Eds.), *Advances in Neural Information Processing Systems 7* (pp. 10–18). Cambridge, MA: MIT Press.
- Markey, K. L., Menn, L., and Mozer, M. C. (1995). A developmental model of the sensorimotor foundations of child phonology. In D. MacLaughlin & S. McEwen (Eds.), *Proceedings of the Nineteenth Boston University Conference on Language Development, 19*, 367–378. Somerville, MA: Cascadilla Press.
- Mathis, D., & Mozer, M. C. (1996). Conscious and unconscious perception: A computational theory. In G. Cottrell (Ed.), *Proceedings of the Eighteenth Annual Conference of the Cognitive Science Society* (pp. 324–328). Hillsdale, NJ: Erlbaum.
- Mozer, M. C., Vidmar, L., & Dodier, R. H. (1997). The Neurothermostat: Adaptive control of residential heating systems. In M. C. Mozer, M. I. Jordan, & T. Petsche (Eds.), *Advances in Neural Information Processing Systems 9* (pp. 953–959). Cambridge, MA: MIT Press.
- Uno, Y., & Mozer, M. C. (1997). Neural net architectures in modeling compositional syntax: Prediction and perception of continuity in minimalist works by Phillip Glass and Louis Andriessen. *Proceedings of the International Computer Music Conference, Greece*.
- Mozer, M. C., Sitton, M., & Farah, M. J. (1998). A superadditive-impairment theory of optic aphasia. In M. I. Jordan, M. Kearns, & S. A. Solla (Eds.), *Advances in Neural Information Processing Systems 10* (pp. 66–72). Cambridge, MA: MIT Press.
- Mozer, M. C. (1998). The neural network house: An environment that adapts to its inhabitants. In M. Coen (Ed.), *Proceedings of the American Association for Artificial Intelligence Spring Symposium on Intelligent Environments* (pp. 110–114; Technical Report SS-98-02). Menlo Park, CA: AAAI Press.
- Mozer, M. C. (1999). A principle for the unsupervised hierarchical decomposition of visual scenes. In M. S. Kearns, S. A. Solla, & D. Cohn (Eds.), *Advances in Neural Information Processing Systems 11* (pp. 52–58). Cambridge, MA: MIT Press.
- Mozer, M. C. (1999). Do attention and perception require multiple reference frames? Evidence from a computational model of unilateral neglect. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the Twenty First Annual Conference of the Cognitive Science Society* (pp. 456–461). Hillsdale, NJ: Lawrence Erlbaum Associates.
- O'Reilly, R. C., Mozer, M. C., Munakata, Y., & Miyake, A. (1999). Discrete representations in working memory: A hypothesis and computational investigations. In *Proceedings of the Second International Conference on Cognitive Science* (pp. 183–188). Tokyo, Japan: Japanese Cognitive Science Society.
- Lee, S.-Y., & Mozer, M. C. (2000). Robust recognition of noisy and superimposed patterns via selective attention. In S. A. Solla, T. K. Leen & K.-R. Mueller (Eds.), *Advances in Neural Information Processing Systems 12* (pp. 31–37). Cambridge, MA: MIT Press.
- Mozer, M. C., Wolniewicz, R., Grimes, D. B., Johnson, E., & Kaushansky, H. (2000). Churn reduction in the wireless industry. In S. A. Solla, T. K. Leen & K.-R. Mueller (Eds.), *Advances in Neural Information Processing Systems 12* (pp. 935–941). Cambridge, MA: MIT Press.
- Hochreiter, S., & Mozer, M. C. (2000). An electric field approach to independent component analysis. In P. Pajunen & J. Karhunen (Eds.), *Proceedings of the Second International Workshop on Independent Component Analysis and Blind Signal Separation* (pp. 145–150). Espoo, Finland: Otamedia.

- Wolniewicz, R., Mozer, M. C., Dodier, R. H., Yan, L., and Guerra, C. (2000). Enhancing customer relationship management solutions through choice of data representation. In *Proceedings of CRM Infrastructure*. IBC Global Conferences, September 20-21, 2000. London, England. Available at [www.ibctelecoms.com/crminfra.s](http://www.ibctelecoms.com/crminfra.s)
- Zemel, R. S., & Mozer, M. C. (2000). A generative model for attractor dynamics. In S. A. Solla, T. K. Leen & K.-R. Mueller (Eds.), *Advances in Neural Information Processing Systems 12* (pp. 80–86). Cambridge, MA: MIT Press.
- Grimes, D., & Mozer, M. C. (2001). The interplay of symbolic and subsymbolic processes in anagram problem solving. In T. K. Leen, T. Dietterich, & V. Tresp (Eds.), *Advances in Neural Information Processing Systems 13* (pp. 17–23). Cambridge, MA: MIT Press.
- Hochreiter, S., & Mozer, M. C. (2001). Beyond maximum likelihood and density estimation: A sample-based criterion for unsupervised learning of complex models. In T. K. Leen, T. Dietterich, & V. Tresp, (Eds.), *Advances in Neural Information Processing Systems 13* (pp. 535–541). Cambridge, MA: MIT Press.
- Hochreiter, S., & Mozer, M. C. (2001). A discrete probabilistic memory model for discovering dependencies in time. In G. Dorffner, H. Bischof, & K. Hornig (Eds.), *Proceedings of the International Conference on Artificial Neural Networks (ICANN) 2001* (pp. 661–668). Berlin: Springer-Verlag.
- Yan, L., Miller, D., Mozer, M. C., & Wolniewicz, R. (2001). Improving prediction of customer behavior in nonstationary environments. *Proceedings of the International Joint Conference on Neural Networks* (pp. 2258–2263). Piscataway, NJ: IEEE Press.
- Hochreiter, S., & Mozer, M. C. (2001). Monaural separation and classification of mixed signals: A support-vector regression perspective. *Proceedings of the Third International Conference on Independent Component Analysis and Blind Signal Separation*, San Diego, CA.
- Mozer, M. C., Colagrosso, M. D., & Huber, D. H. (2002). A rational analysis of cognitive control in a speeded discrimination task. In T. Dietterich, S. Becker, & Ghahramani, Z. (Eds.) *Advances in Neural Information Processing Systems 14* (pp. 51–57). Cambridge, MA: MIT Press.
- Mozer, M. C., Dodier, R., Colagrosso, M. D., Guerra-Salcedo, C., & Wolniewicz, R. (2002). Prodding the ROC curve: Constrained optimization of classifier performance. In T. Dietterich, S. Becker, & Ghahramani, Z. (Eds.) *Advances in Neural Information Processing Systems 14* (pp. 1409–1415). Cambridge, MA: MIT Press.
- Hochreiter, S., Mozer, M. C., & Obermayer, K. (2003). Coulomb classifiers: Generalizing support vector machines via an analogy to electrostatic systems. In S. Becker, S. Thrun, & K. Obermayer (Eds.), *Advances in Neural Information Processing Systems 15* (pp. 545–552). Cambridge, MA: MIT Press.
- Mozer, M. C., Colagrosso, M. D., & Huber, D. E. (2003). Mechanisms of long-term repetition priming and skill refinement: A probabilistic pathway model. In *Proceedings of the Twenty Fifth Conference of the Cognitive Science Society* (pp. 828–833). Hillsdale, NJ: Erlbaum Associates.
- Yan, L., Dodier, R., Mozer, M. C., & Wolniewicz, R. (2003). Optimizing classifier performance via the Wilcoxon-Mann-Whitney statistic. In T. Fawcett & N. Mishra (Eds.), *Proceedings of the Twentieth International Conference on Machine Learning (ICML-2003)* (pp. 848–855). Menlo Park, CA: AAAI Press.
- Mozer, M. C., Howe, M., & Pashler, H. (2004). Using testing to enhance learning: A comparison of two hypotheses. In K. Forbus, D. Gentner, T. Regier (Eds.), *Proceedings of the Twenty Sixth Conference of the Cognitive Science Society* (pp. 975–980). Hillsdale, NJ: Erlbaum Associates.
- Mozer, M. C., Kinoshita, S., & Davis, C. (2004). Control of response initiation: Mechanisms of adaptation to recent experience. In K. Forbus, D. Gentner, T. Regier (Eds.), *Proceedings of the Twenty Sixth Conference of the Cognitive Science Society* (pp. 981–986). Hillsdale, NJ: Erlbaum Associates.
- Bohte, S., & Mozer, M. C. (2005). Reducing spike train variability: A computational theory of spike-timing dependent plasticity. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17* (pp. 201–208). Cambridge, MA: MIT Press.
- Colagrosso, M. D., & Mozer, M. C. (2005). Theories of access consciousness. In L. K. Saul, Y. Weiss, & L. Bottou (Eds.), *Advances in Neural Information Processing Systems 17* (pp. 289–296). Cambridge, MA: MIT Press.
- Hauswirth, M., Diwan, A., Sweeney, P. F., & Mozer, M. C. (2005). Automated vertical profiling. In *Proceedings of Object Oriented Programming, Languages, and Applications 2005 (OOPSLA'05)*.
- Mozer, M. C., Shettel, M., & Vecera, S. P. (2006). Control of visual attention: A rational account. In Y. Weiss, B. Schölkopf, & J. Platt (Eds.), *Neural Information Processing Systems 18* (pp. 923–930). Cambridge, MA: MIT Press.

- Baldwin, D., & Mozer, M. C. (2006). Controlling attention with noise: The cue-combination model of visual search. In R. Sun & N. Miyake (Eds.), *Proceedings of the Twenty Eighth Conference of the Cognitive Science Society* (pp. 42–47). Hillsdale, NJ: Erlbaum Associates.
- Mozer, M. C., & Fan, A. (2007). Top-down modulation of neural responses in visual perception: A computational exploration. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Cognitive Science Society* (pp. 491–496). Austin, TX: Cognitive Science Society.
- Mozer, M. C., Jones, M., & Shettel, M. (2007). Context effects in category learning: An investigation of four probabilistic models. In B. Schölkopf, J. Platt, & T. Hoffmann (Eds.), *Neural Information Processing Systems 19* (pp. 993–1000). Cambridge, MA: MIT Press.
- Mozer, M. C., & Baldwin, D. (2008). Experience-guided search: A theory of attentional control. In J. Platt, D. Koller, Y. Singer, & Roweis, S. (Eds.), *Advances in Neural Information Processing Systems 20* (pp. 1033–1040). Cambridge, MA: MIT Press.
- Mozer, M. C., Pashler, H., & Homaei, H. (2008). Optimal predictions in everyday cognition: The wisdom of individuals or crowds? In B. C. Love, K. McRae, & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1051–1056). Austin, TX: Cognitive Science Society.
- Jones, M., Mozer, M. C., & Kinoshita, S. (2009). Optimal response initiation: Why recent experience matters. In D. Koller, D. Schuurmans, Y. Bengio, & L. Bottou (Eds.), *Advances in Neural Information Processing Systems 21* (pp. 785–792). La Jolla, CA: NIPS Foundation.
- Reynolds, J., & Mozer, M. C. (2009). Temporal dynamics of cognitive control. In D. Koller, D. Schuurmans, Y. Bengio, & L. Bottou (Eds.), *Advances in Neural Information Processing Systems 21* (pp. 1353–1360). La Jolla, CA: NIPS Foundation.
- Lindsey, R., Mozer, M. C., Cepeda, N. J., & Pashler, H. (2009). Optimizing memory retention with cognitive models. In A. Howes, D. Peebles, & R. Cooper (Eds.), *Proceedings of the Ninth International Conference on Cognitive Modeling (ICCM)*. Manchester, UK.
- Knights, D., Mytkowicz, T., Sweeney, P. F., Mozer, M. C., & Diwan, A. (2009). Blind optimization for exploiting hardware features. In O. de Moor & M. I. Schwartzbach (Eds.), *Lecture Notes in Computer Science, v. 5501: Compiler Construction 2009* (pp. 251–265). New York: Springer
- Knights, D., Mozer, M. C., Nicolov, N. (2009). Detecting topic drift with compound topic models. In *Proceedings of the Third International AAAI Conference on Weblogs and Social Media*. AAAI Press.
- Mozer, M. C., Pashler, H., Cepeda, N., Lindsey, R., & Vul, E. (2009). Predicting the optimal spacing of study: A multiscale context model of memory. In Y. Bengio, D. Schuurmans, J. Lafferty, C.K.I. Williams, & A. Culotta (Eds.), *Advances in Neural Information Processing Systems 22* (pp. 1321–1329). La Jolla, CA: NIPS Foundation.
- Wilder, M., Jones, M., & Mozer, M. C. (2009). Sequential effects reflect parallel learning of multiple environmental regularities. In Y. Bengio, D. Schuurmans, J. Lafferty, C.K.I. Williams, & A. Culotta (Eds.), *Advances in Neural Information Processing Systems 22* (pp. 2053–2061). La Jolla, CA: NIPS Foundation.
- Lindsey, R., Lewis, O., Pashler, H., & Mozer, M. C. (2010). Predicting students' retention of facts from feedback during training. In S. Ohlsson & R. Catrambone (Eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 2332–2337). Austin, TX: Cognitive Science Society.
- Mozer, M. C., Pashler, H., Wilder, M., Lindsey, R., Jones, M., & Jones, M. (2010). Decontaminating human judgments by removing sequential dependencies. In J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. S. Zemel, & A. Culota (Eds.), *Advances in Neural Information Processing Systems 23* (pp. 1705–1713). La Jolla, CA: NIPS Foundation.
- Link, B. V., Kos, B., Wager, T. D., & Mozer M. C. (2011). Past experience influences judgment of pain: Prediction of sequential dependencies. In L. Carlson, C. Hoelscher, & T. F. Shipley (Eds.), *Proceedings of the 33d Annual Conference of the Cognitive Science Society* (pp. 1248–1253). Austin, TX: Cognitive Science Society.
- Mozer, M. C., Link, B. V., & Pashler, H. (2011) An unsupervised decontamination procedure for improving the reliability of human judgments. In Shawe-Taylor, J., Zemel, R. S., Bartlett, P., Pereira, F., & Weinberger, K. Q. (Eds.), *Advances in Neural Information Processing Systems 24* (pp. 1791–1799). La Jolla, CA: NIPS Foundation.

- Khajah, M., Lindsey, R., & Mozer, M. C. (2013). Maximizing students' retention via spaced review: Practical guidance from computational models of memory. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 758–763). Austin, TX: Cognitive Science Society. [Awarded the Cognitive Science Society Computational Modeling Prize]
- Lindsey, R. V., Mozer, M. C., Huggins, W. J., & Pashler, H. (2013). Optimizing instructional policies. In C. J. C. Burges et al. (Eds.), *Advances in Neural Information Processing Systems 26* (pp. 2778–2786). Curran Associates, Inc. [Among the 20 of 1420 submissions accepted for oral presentation]
- Khajah, M., Huang, Y., Gonzales-Brenes, J. P., Mozer, M. C., & Brusilovsky, P. (2014). Integrating knowledge tracing and item response theory: A tale of two frameworks. In M. Kravcik, O. C. Santos, J. G. Boticario (Eds.), *Proceedings of the 4th International Workshop on Personalization Approaches in Learning Environments* (pp. 7–15). CEUR Workshop Proceedings, ISSN 1613-0073.
- Khajah, M., Wing, R. M., Lindsey, R. V., & Mozer, M. C. (2014) Incorporating latent factors into knowledge tracing to predict individual differences in learning. In J. Stamper, Z. Pardos, M. Mavrikis, & B. M. McLaren (Eds.), *Proceedings of the Seventh International Conference on Educational Data Mining* (pp. 99–106). Educational Data Mining Society Press. [Awarded Best Paper of Conference]
- Lindsey, R. V., Khajah, M., & Mozer, M. C. (2014). Automatic discovery of cognitive skills to improve the prediction of student learning. In Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, & K. Q. Weinberge (Eds.), *Advances in Neural Information Processing Systems 27* (pp. 1386–1394). La Jolla, CA: Curran Associates Inc. [Acceptance rate 24.7%]
- Beckage, N., Mozer, M. C., & Colunga, E. (2015). Predicting a child's trajectory of lexical acquisition. In D. Noelle et al. (Eds.), *Proceedings of the 37th Annual Conference of the Cognitive Science Society* (pp. 196–201). Austin, TX: Cognitive Science Society.
- Khajah, M., Lindsey, R. V., & Mozer, M. C. (2016). How deep is knowledge tracing? In T. Barnes, M. Chi, & M. Feng (Eds.), *Proceedings of the Ninth International Conference on Educational Data Mining* (pp. 94–101). Educational Data Mining Society. [Awarded Best Paper of Conference]
- Khajah, M., Roads, B. D., Lindsey, R. V., Liu, Y.-E., & Mozer, M. C. (2016). Designing engaging games using Bayesian optimization. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems* (pp. 5571–5582). New York: ACM. [Acceptance rate 23.4%]
- Wilson, K.H., Xiong, X., Khajah, M., Lindsey, R.V., Zhao, S., Karklin, Y., Van Inwegen, E.G., Han, B., Ekanadham, C., Beck, J.E., Heffernan, N., & Mozer, M. C. (2016). Estimating student proficiency: Deep learning is not the panacea. In R. G. Baraniak, J. Ngiam, C. Studer, P. Grimaldi, & A. S. Lan (Eds.), *Proceedings of the 2016 NIPS Workshop on Machine Learning for Education*. Barcelona, Spain.
- Mozer, M. C., Kazakov, D., & Lindsey, R. V. (2017). Continuous-time, discrete-event memory nets. Submitted for publication. arXiv:1710.04110 [cs.NE]
- Snell, J., Ridgeway, K. Liao, R., Roads, B. D., Mozer, M. C., & Zemel, R. S. (2017). Learning to generate images with perceptual similarity metrics. *IEEE International Conference on Image Processing*. Beijing, China: IEEE Press.
- Ke, N. R., Goyal, A., Bilaniuk, O., Binas, J., Charlin, L. Mozer, M. C., Pal, C., & Bengio, Y. (2018). Sparse attentive backtracking: Temporal credit assignment through reminding. In S. Bengio et al. (Eds.), *Advances in Neural Information Processing Systems 32* (pp. 7651–7662). [Acceptance rate 4.1% for spotlight and orals]
- Khajah, M. M., Mozer, M. C., Kelly, S., & Milne, B. (2018). Boosting engagement with educational software using near wins. In C. P. Rosé et al. (Eds.), *Nineteenth Conference on Artificial Intelligence for Education (LNCS 10948)* (pp. 171–175). Springer. <https://doi.org/10.1007/978-3-319-93846-2>
- Montero, S., Arora, A., Kelly, S., Milne, B., & Mozer, M. C. (2018). Does deep knowledge tracing model interactions among skills? In K. E. Boyer & M. Yudelson (Eds.), *Proceedings of the 11th International Conference on Educational Data Mining* (pp. 462–466). EDM Society Press.
- Ridgeway, K., & Mozer, M. C. (2018). Learning deep disentangled representations with the  $F$  statistic loss. In S. Bengio et al. (Eds.), *Advances in Neural Information Processing Systems 32* (pp. 185–194). [Acceptance rate 20.1%]
- Scott, T., Ridgeway, K., & Mozer, M. C. (2018). Adapted deep embeddings: A synthesis of methods for  $k$ -shot inductive transfer learning. In S. Bengio et al. (Eds.), *Advances in Neural Information Processing Systems 32* (pp. 76–85). [Acceptance rate 4.1% for spotlight and orals]
- Winchell, A., Lan, A., Grimaldi, P., Pashler, H., & Mozer, M. C. (2018). Can student annotations serve as an early predictor of student learning? In K. E. Boyer & M. Yudelson (Eds.), *Proceedings of the 11th International Conference on Educational Data Mining* (pp. 431–437). EDM Society Press.

- Mozer, M. C., Kazakov, D., & Lindsey, R. V. (2018). State-denoised recurrent neural networks. arXiv:1805.08394 [cs.NE]
- Lamb, A., Binas, J., Goyal, A., Subramanian, S., Mitliagkas, I., Kazakov, D., Bengio, Y., & Mozer, M. C. (2019). State-reification networks: Improving generalization by modeling the distribution of hidden representations. *Proceedings of the 36th International Conference on Machine Learning*, 97, 3622–3631.
- Ridgeway, K., & Mozer, M. C. (2019). Open-ended content-style recombination via leakage filtering. Submitted for publication. arXiv:1810.00110 [cs.LG]
- Sense, F., Jastrzembski, T., Mozer, M. C., Krusmark, M., and van Rijn, H. (2019). Perspectives on computational models of learning and forgetting. *Proceedings of the Seventeenth International Conference on Cognitive Modeling* (53–58). State College, PA: Applied Cognitive Science Lab.
- Scott, T., Ridgeway, K., & Mozer, M. C. (2019). Stochastic prototype embeddings. arXiv:1909.11702 [stat.ML]
- Iuzzolino, M., Singer, Y., & Mozer, M. C. (2019). Convolutional bipartite attractor networks. arXiv:1906.03504 [cs.LG]
- Attarian, M., Roads, B. D., & Mozer, M. C. (2020). Transforming neural network representations to predict human judgments of similarity. *Workshop on Shared Visual Representations in Human and Machine Intelligence (SVRHM 2020)*. Also arXiv:2010.06512 [cs.NE]
- Davidson, G., & Mozer, M. C. (2020). Sequential mastery of multiple tasks: Networks naturally learn to learn and forget to forget. *IEEE Conference on Computer Vision and Pattern Recognition* (pp. 9282–9293). Also arXiv:1905.10837 [cs.LG]
- Kim, D. Y. J., Winchell, A., Waters, A. E., Grimaldi, P. J., Baraniuk, R., & Mozer, M. C. (2020). Inferring student comprehension from highlighting patterns in digital textbooks: An exploration in an authentic learning platform. In S. Sosnovsky, P. Brusilovsky, R. G. Baraniuk, & A. S. Lan (Eds.), *Proceedings of the Second International Workshop on Intelligent Textbooks*. Springer.
- Mittal, S., Lamb, A., Goyal, A., Voleti, V., Shanahan, M., Lajoie, G., Mozer, M. C., & Bengio, Y. (2020). Learning to combine top-down and bottom-up signals in recurrent neural networks with attention over modules. *International Conference on Machine Learning*.
- Zhang, C., Bengio, S., Hardt, M., Mozer, M. C., & Singer, Y. (2020). Identity crisis: Memorization and generalization under extreme overparameterization. *Proceedings of the International Conference on Learning Representations*. Also arXiv:1902.04698 [stat.ML]
- Didolkar, A., Goyal, A., Ke, N. R., Blundell, C., Beaudoin, P., Heess, N., Mozer, M. C., & Bengio, Y. (2021). Neural production systems. In *Advances in Neural Information Processing Systems 34*. Also arXiv:2103.01937 [cs.AI]
- Goyal, A., Lamb, A., Gampa, P., Beaudoin, P., Levine, S., Blundell, C., Bengio, Y., & Mozer, M. C. (2021). Object files and schemata: Factorizing declarative and procedural knowledge in dynamical systems. *International Conference on Learning Representations*. Also arXiv:2006.16225 [cs.LG]
- Iuzzolino, M. L., Mozer, M. C., & Bengio, S. (2021). Improving anytime prediction with parallel, cascaded networks and a temporal-difference loss. In *Advances in Neural Information Processing Systems 34*. Also arXiv:2102.09808 [cs.LG]
- Jiang, Z., Zhang, C., Talwar, K., & Mozer, M. C. (2021). Characterizing structural regularities of labeled data in overparameterized models. In *Proceedings of the 38th International Conference on Machine Learning*, PMLR 139:5034–5044.
- Karandikar, A., Cain, N., Tran, D., Lakshminarayanan, B., Shlens, J., Mozer, M. C., & Roelofs, B. (2021). Soft calibration objectives for neural networks. In *Advances in Neural Information Processing Systems 34*. Also arXiv.org:2018.00106 [cs.LG]
- Ke, N. R., Didolkar, A., Mittal, S., Goyal, A., Lajoie, G., Bauer, S., Rezende, D., Mozer, M. C., Bengio, Y., & Pal, C. (2021). Systematic evaluation of causal discovery in visual model-based reinforcement learning. Submitted for publication.
- Kim, D. Y. J., Scott, T. R., Mallick, D., & Mozer, M. C. (2021). Using semantics of textbook highlights to predict student comprehension and knowledge retention. In S. Sosnovsky, P. Brusilovsky, R. G. Baraniuk, & A. S. Lan (Eds.), *Proceedings of the Third International Workshop on Intelligent Textbooks (iTextbooks)* (pp. 108–120). Springer.
- Lamb, A., Goyal, A., Stowik, A., Mozer, M. C., Beaudoin, P., & Bengio, Y. (2021). Neural function modules with sparse arguments: A dynamic approach to integrating information across layers. *AISTATS*. Also arXiv:2010.08012 [cs.LG]
- Li, Z., Mozer, M. C., & Whitehill, J. (2021). Compositional embeddings for multi-label one-shot learning. *IEEE Winter Conference on Applications of Computer Vision*. Also arXiv:2002.04193 [cs.LG]

- Liu, D., Lamb, A., Kawaguchi, K., Goyal, A., Sun, C., Mozer, M. C., & Bengio, Y. (2021). Discrete-valued neural communication in structured models enhances generalization. In *Advances in Neural Information Processing Systems 34*. Also arXiv.org:2107.02367 [cs.LG]
- Ren, M., Iuzzolino, M. L., Mozer, M. C., & Zemel, R. S. (2021). Wandering within a world: Online contextualized few-shot learning. *International Conference on Learning Representations*. Also arXiv:2007.04546 [cs.LG]
- Ren, M., Scott, T. R., Iuzzolino, M. L., Mozer, M. C., & Zemel, R. S. (2021). Online unsupervised learning of visual representations and categories. arXiv.org:2109.05675 [cs.LG]
- Scherrer, N., Bilaniuk, O., Annadani, Y., Goyal, A., Schwab, P., Schölkopf, B., Mozer, M. C., Bengio, Y., & Ke, N. R. (2021). Learning neural causal models with active interventions. *NeurIPS Workshop on Causal Inference and Machine Learning (WHY-21)*. Also arXiv.org:2019.02429 [stat.ML]
- Scott, T. R., Gallagher, A. C., & Mozer, M. C. (2021). Von Mises-Fisher loss: An exploration of embedding geometries for supervised learning. *Proceedings of the IEEE/CVF International Conference on Computer Vision*. Also arXiv:2103.15718 [cs.LG]
- Teterwak, P., Zhang, C., Krishnan, D., & Mozer, M. C. (2021). Understanding invariance via feedforward inversion of discriminatively trained classifiers. In *Proceedings of the 38th International Conference on Machine Learning*. PMLR 139:10225-10235. Also arXiv: 2103.07470 [cs.LG]
- Aithal, S. K., Goyal, A., Lamb, A., Bengio, Y., & Mozer, M. C. (2022). Leveraging the third dimension in contrastive learning. *NeurIPS Workshop on Self-Supervised Learning: Theory and Practice*. New Orleans, LA. Also arXiv.org:2301.11790 [cs.CV].
- Elsayed, G. F., Mahendran, A., van Steenkiste, S., Greff, K., Mozer, M. C., & Kipf, T. (2022). SAVi++: Towards End-to-End Object-Centric Learning from Real-World Videos. *Advances in Neural Information Processing Systems 35*. Also arXiv.org:2206.07764 [cs.CV]
- Evcı, U., Dumoulin, V., Larochelle, H., & Mozer, M. C. (2022). Head2Toe: Utilizing intermediate representations for better transfer learning. *Proceedings of the 39th Annual Conference on Machine Learning*. Also arXiv.org:2201.03529 [cs.LG]
- Goyal, A., Didolkar, A., Lamb, A., Badola, K., Ke, N. R., Rahaman, N., Binas, J., Blundell, C. Mozer, M. C., & Bengio, Y. (2022). Coordination among neural modules through a shared global workspace. *International Conference on Learning Representations (ICLR)*. Also arXiv:2103.01197 [cs.Lg]
- Jones, M., Mayo, D., Scott, T., Ren, M., ElSayed, G., Hermann, K., & Mozer, M. C. (2022). Neural network online training with sensitivity to multiscale temporal structure. *Proceedings of the NeurIPS 2022 Workshop on Memory in Artificial and Real Intelligence*.
- Khalifa, A., Mozer, M. C., Sedghi, H., Neyshabur, B., & Alabdulmohsin, I. (2022). Layer-stack temperature scaling. arXiv:2211:10193 [cs.LG]
- Liu, D., Shah, V., Boussif, O. Meo, C., Goyal, A., Shu, T., Mozer, M. C., Heess, N., & Bengio, Y. (2022). Stateful active facilitator: Coordination and environmental heterogeneity in cooperative multi-agent reinforcement learning. arxiv.org:2210.03022 [cs.AI]
- Roelofs, R., Cain, N., Shlens, J., & Mozer, M. C. (2022). Mitigating bias in calibration error estimation. *Twenty Fifth International Conference on Artificial Intelligence and Statistics (AISTats)*. Also arXiv:2012.08668 [cs.LG]
- Chakravarthy, A., Nguyen, T., Goyal, A., Bengio, Y., & Mozer, M. C. (2023). Spotlight attention: Robust object-centric learning with a spatial locality prior. arXiv.org:2305.19550 [cs.CV]
- Hermann, K. L., Mobahi, H., Fel, T., & Mozer, M. C. (2023). On the foundations of shortcut learning. arXiv.org:2310.16228 [cs.LG]
- Jones, M., Scott, T. R., Ren, M., Elsayed, G. F., Hermann, K., Mayo, D., & Mozer, M. C. (2023). Learning in temporally structured environments. *Proceedings of the International Conference on Learning Representations (ICLR 2023)*.
- Ke, N. R., Chiappa, S., Wang, J., Bornschein, J., Weber, T., Goyal, A., Botvinick, M., Mozer, M. C., & Rezende, D. J. (2023). Learning to induce causal structure. *Proceedings of the International Conference on Learning Representations (ICLR 2023)*. Also arXiv.org:2204.04875 [stat.ML]
- Liu, D., Lamb, A., Ji, X., Notsawo, P., Mozer, M. C., Bengio, Y., & Kawaguchi, K. (2023). Adaptive discrete communication bottlenecks with dynamic vector quantization for heterogeneous representational coarseness. *Thirty-Seventh AAAI Conference on Artificial Intelligence*. Also arXiv.org:2202.01334 [cs.LG]

- Liu, D., Shah, V., Boussif, O., Meo, C., Goyal, A., Shu, T., Mozer, M. C., Heess, N., & Bengio, Y. (2023). Stateful active facilitator: Coordination and environmental heterogeneity in cooperative multi-agent reinforcement learning. *Proceedings of the International Conference on Learning Representations (ICLR 2023)*.
- Maini, P., Mozer, M. C., Sedghi, H., Lipton, Z. C., Kolter, J. Z., Zhang, C. (2023). Can neural network memorization be localized? *Proceedings of the International Conference on Learning Representations (ICLR 2023)*.
- Mayo, D., Scott, T. R., Ren, M., Elsayed, G., Hermann, K., Jones, M., & Mozer, M. C. (2023). Multitask learning via interleaving: A neural network investigation. In *Proceedings of the 45th Annual Conference of the Cognitive Science Society (COGSCI 2023)*.
- Shah, V., Träuble, F., Malik, A., Larochelle, H., Mozer, M., Arora, S., Bengio, Y., & Goyal, A. (2023). Unlearning via sparse representations. [arXiv.org:2311.15268](https://arxiv.org/abs/2311.15268) [cs.LG]
- Träuble, F., Goyal, A., Rahaman, N., Mozer, M. C., Kawaguchi, K., Bengio, Y., & Schölkopf, B. (2023). Discrete key-value bottleneck. *Proceedings of the 40th International Conference on Machine Learning*. Also [arXiv.org:2207.11240](https://arxiv.org/abs/2207.11240) [cs.LG]
- Wu, Y.-F., Greff, K., Elsayed, G., Mozer, M., Kipf, T., & van Steenkiste S. (2023). Inverted-attention transformers can learn object representations: Insights from slot attention. *NeurIPS Workshop on Causal Representation Learning*.
- Didolkar, A., Goyal, A., Ke, N. R., Guo, S., Valko, M., Lillicrap, T., Rezende, D., Bengio, Y., Mozer, M. C., & Arora, S. (2024). Metacognitive capabilities of LLMs: An exploration in mathematical problem solving. [arXiv.org:2405.12205](https://arxiv.org/abs/2405.12205) [cs.AI]
- Didolkar, A., Zadaianchuk, A., Goyal, A., Mozer, M. C., Bengio, Y., Martius, G., & Seitzer, M. (2024) Zero-shot object-centric learning. [arXiv.org:2408.09162](https://arxiv.org/abs/2408.09162) [cs.CV]
- Gopalakrishnan, A., Stanić, A., Schmidhuber, J., & Mozer, M. C. (2024). Recurrent complex-weighted autoencoders for unsupervised object discovery. [arXiv.org:2405.17283](https://arxiv.org/abs/2405.17283) [cs.LG]
- Muttenthaler, L., Greff, K., Born, F., Spitzer, B., Kornblith, S., Mozer, M. C., Mueller, K.-R., Unterthiner, T., & Lampinen, A. K. (2024). Aligning machine and human visual representations across abstraction levels. [arXiv.org:2409.06509](https://arxiv.org/abs/2409.06509) [cs.CV]
- Shah, V., Goyal, A., Yu, D., Lyu, K., Park, S., Ke, N. R., McClelland, J. L., Bengio, Y., Arora, S., & Mozer, M. C. (2024). AI-assisted generation of difficult math questions. In ICML AI4MATH Workshop.
- Wang, H., Zou, J., Mozer, M., Zhang, L., Goyal, A., Lamb, A., Deng, Z., Xie, M. Q., Brown, H., & Kawaguchi, K. (2024). Can AI be as creative as humans? [arXiv.org:2401.01623](https://arxiv.org/abs/2401.01623) [cs.AI]
- Yang, Y., Jones, M., Mozer, M. C., & Ren, M. (2024). Reawakening knowledge: Anticipatory recovery from catastrophic interference via structured training. [arXiv.org:2403.09613](https://arxiv.org/abs/2403.09613) [cs.LG]

**REFEREED BOOK CHAPTERS**

- Zemel, R. S., Mozer, M. C., & Hinton, G. E. (1989). TRAFFIC: A model of object recognition based on transformations of feature instances. In D. S. Touretzky, G. E. Hinton, and T. J. Sejnowski (Eds.), *Proceedings of the 1988 Connectionist Models Summer School* (pp. 452–461). San Mateo, CA: Morgan Kaufmann.
- Menn, L., Markey, K., Mozer, M. C., & Lewis, C. (1993). Connectionist modeling and the microstructure of phonological development: A progress report. In B. de Boysson-Bardies, S. de Schonen, P. Jusczyk, P. MacNeilage, and J. Morton (Eds.), *Developmental neurocognition: Speech and face processing in the first year of life* (pp. 421–433). Boston: Kluwer Academic Publishers.
- Mozer, M. C. (1993). Neural network architectures for temporal pattern processing. In A. S. Weigend & N. A. Gershenfeld (Eds.), *Time series prediction: Forecasting the future and understanding the past* (pp. 243–264). Redwood City, CA: Santa Fe Institute Studies in the Sciences of Complexity, Proceedings Volume XVII, Addison-Wesley Publishing.
- Mozer, M. C. (2004). Lessons from an adaptive house. In D. Cook & S. Das (Eds.), *Smart environments: Technologies, protocols, and applications* (pp. 273–294). J. Wiley & Sons.
- Mozer, M. C., Kinoshita, S., & Shettel, M. (2007). Sequential dependencies in human behavior offer insights into cognitive control. In W. Gray (Ed.), *Integrated Models of Cognitive Systems* (pp. 180–193). Oxford UK: Oxford University Press.
- Mozer, M. C. (2009). Attractor networks. In T. Bayne, A. Cleeremans, & P. Wilken (Eds.), *Oxford Companion to Consciousness* (pp. 86–89). Oxford, UK: Oxford U. Press.
- Mozer, M. C., & Lindsey, R. V. (2017). Predicting and improving memory retention: Psychological theory matters in the big data era. In M. Jones (Ed.), *Big Data in Cognitive Science* (pp. 34–64). New York: Routledge.



**OTHER BOOK  
CHAPTERS**

- Anderson, J. A., & Mozer, M. C. (1981). Categorization and selective neurons. In G. E. Hinton & J. A. Anderson (Eds.), *Parallel models of associative memory* (pp. 213–236). Hillsdale, NJ: Erlbaum Associates.
- Mozer, M. C., & Gross, H. (1986). An architecture for experiential learning. In T. M. Mitchell, J. G. Carbonell, & R. S. Michalski (Eds.), *Machine learning: A guide to current research* (pp. 219–226). Boston: Kluwer Academic Publishers.
- Mozer, M. C. (1991). Connectionist music composition based on melodic, stylistic, and psychophysical constraints. In P. M. Todd & D. G. Loy (Eds.), *Music and connectionism* (pp. 195–211). Cambridge, MA: MIT Press/Bradford Books.
- Mozer, M. C., & Behrmann, M. (1992). Reading with attentional impairments: A brain-damaged model of neglect and attentional dyslexias. In R. G. Reilly & N. E. Sharkey (Eds.), *Connectionist approaches to natural language processing* (pp. 409–460). Hillsdale, NJ: Erlbaum Associates.
- Bachrach, J., & Mozer, M. C. (1995). Modeling and controlling finite-state systems. In Y. Chauvin & D. E. Rumelhart (Eds.), *Back propagation: Theory, architectures, and applications* (pp. 351–388). Hillsdale, NJ: Erlbaum.
- Mozer, M. C. (1995). A focused back-propagation algorithm for temporal sequence recognition. In Y. Chauvin & D. E. Rumelhart (Eds.), *Back propagation: Theory, architectures, and applications* (pp. 137–170). Hillsdale, NJ: Erlbaum.
- Mozer, M. C., Dodier, R. H., Anderson, M., Vidmar, L., Cruickshank III, R. F., Miller, D. (1995). The neural network house: An overview. In L. Niklasson & M. Boden (Eds.), *Current trends in connectionism* (pp. 371–380). Hillsdale, NJ: Erlbaum.
- Mozer, M. C., & Sitton, M. (1998). Computational modeling of spatial attention. In H. Pashler (Ed.), *Attention* (pp. 341–393). London: UCL Press.
- Mozer, M. C., & Miller, D. (1998). Parsing the stream of time: The value of event-based segmentation in a complex, real-world control problem. In C. L. Giles & M. Gori (Eds.), *Adaptive processing of temporal sequences and data structures* (pp. 370–388). Berlin: Springer Verlag.
- Mozer, M. C. (1999). Explaining object-based deficits in unilateral neglect without object-based frames of reference. In J. Reggia, E. Ruppini, & D. Glanzman (Eds.), *Disorders of brain, behavior, and cognition: The neurocomputational perspective* (pp. 99–119). New York: Elsevier.
- Tsoi, A. C., Back, A., Principe, J., & Mozer, M. C. (2001). Memory kernels. In J. F. Kolen & S. C. Kremer (Eds.), *A field guide to dynamical recurrent networks* (pp. 39–54). Piscataway, NJ: IEEE Press.
- Mozer, M. C. (2001). Object recognition. In N. J. Smelser & P. B. Baltes (Eds.), *Encyclopedia of the Social and Behavioral Sciences, Vol. 16* (pp. 10781–10784). Oxford, UK: Elsevier Science.
- Mozer, M. C., & Vecera, S. P. (2005). Space- and object-based attention. In L. Itti, G. Rees, & J. Tsotsos (Eds.), *Neurobiology of attention* (pp. 130–134). Academic Press / Elsevier.
- Mozer, M. C. (2006). Rational models of cognitive control. In C. Calude et al. (Eds.), *Unconventional Computation: Fifth International Conference* (pp. 20–25). Berlin: Springer-Verlag
- Hochreiter, S., & Mozer, M. C. (2007). Monaural speech separation by support vector machines: Bridging the divide between supervised and unsupervised learning methods. In S. Makino, T.-W. Lee, & H. Sawada (Eds.), *Blind signal separation* (pp. 405–421). Berlin: Springer-Verlag.
- Mozer, M. C., & Wilder, M. H. (2009). A unified theory of exogenous and endogenous attentional control. In D. Heinke & E. Mavritsaki (Eds.), *Computational modeling in behavioral neuroscience: Closing the gap between neurophysiology and behaviour* (pp. 245–265). London: Psychology Press.
- Knights, D., Mozer, M. C., & Nicolov, N. (2009). Detecting topic drift. *Recent Advances in Natural Language Processing V (a.k.a. Current Issues in Linguistic Theory)* (pp. 113–130). Amsterdam & Philadelphia: John Benjamins.

**TECHNICAL  
REPORTS**

- Mozer, M. C. (1983). *Principal component analysis using parallel computation* (ICS Progress Report, First Quarter). La Jolla: University of California, San Diego, Institute for Cognitive Science.
- Mozer, M. C. (1984). *Inductive information retrieval using parallel distributed computation* (ICS Technical Report 8406). La Jolla: University of California, San Diego, Institute for Cognitive Science.
- Mozer, M. C., Pashler, H., & Miyata, Y. (1991). *Recovering the “where” from the “what”: A connectionist mechanism to direct attention to objects*. Unpublished manuscript.
- Rickard, T. C., Mozer, M. C., & Bourne, L. E. Jr. (1993). An interactive activation model of arithmetic fact retrieval. Technical Report 92-15. Boulder, Colorado: Institute of Cognitive Science, University of Colorado.

Mozer, M. C., Zemel, R. S., & Hungerford, M. (2003). Optimal adaptation of neural codes. An account of repetition suppression. Unpublished Manuscript.

Mozer, M. C., Mytkowicz, T., & Zemel, R. S. (2004). Achieving robust neural representations: An account of repetition suppression.

Richardson, S., Otte, M., Mozer, M. C., Diwan, A., Sweeney, P., & Connors, D. (2009). Discovering the runtime structure of software with probabilistic generative models.

Reid, S., & Mozer, M. C. (2011). Probabilistic pairwise classification.

Lindsey, R. V., & Mozer, M. C. (2013). Predicting individual differences in student learning via collaborative filtering.

Lindsey, R., Polsdofer, E., Mozer, M.C., Kang, S., H., K., & Pashler, H. (2013). Long-term recency is nothing more than ordinary forgetting.

**SELECTED  
INVITED  
PRESENTATIONS**

Invited Participant, Attention and Performance XII, England. July, 1986.

Colloquium, Department of Psychology, University of Guelph. Guelph, Ontario. March 1988.

Tutorial Speaker, Rocky Mountain Artificial Intelligence Conference, Denver, CO. June 1989.

Invited Participant, Symposium on Cognitive Neuroscience. The Eleventh Annual Conference of the Cognitive Science Society. Ann Arbor, MI. August 1989.

Colloquium, Department of Psychology, University of Michigan. Ann Arbor, MI. January 1991.

Colloquium, Department of Music, University of Washington. Seattle, WA. February 1991.

Colloquium, Department of Computer Science, University of Massachusetts at Amherst. May 1991.

Colloquium, International Computer Science Institute. Berkeley, CA. June 1991.

Faculty Lecturer, James S. McDonnell Summer Institute in Cognitive Neuroscience. Dartmouth, NH. July 1991.

Colloquium, Program in Cognitive Science, Princeton University. Princeton, NJ. September 1991.

Invited Speaker, Fourth International Conference of the Society for Computer Science (Gesellschaft fuer Informatik). Munich, Germany. October 1991.

Colloquium, Siemens Nixdorf. Munich, Germany. October 1991.

Colloquium, Department of Psychology, University of Braunschweig. Braunschweig, Germany. October 1991.

Invited Speaker, Workshop in Time Series Analysis and Prediction. Santa Fe Institute for Nonlinear Studies, Santa Fe, NM. May 1992.

Colloquium, Xerox Palo Alto Research Center, September 1992.

Colloquium, Psychology Department, University of Denver, December 1992.

Colloquium, Department of Psychology, Carnegie-Mellon University. Pittsburgh, PA. February 1993.

Colloquium, Department of Computer Science and Engineering, Oregon Graduate Institute. Portland, OR. March 1993.

Colloquium, Department of Mathematics and Computer Science, Colorado School of Mines. Golden, CO. April 1993.

Invited Speaker, Connectionist Models Summer School. Boulder, CO. June 1993.

Colloquium, Department of Computer Science, University of Hamburg. Hamburg Germany, July 1993.

Faculty Lecturer, Summer School in Cognitive Neuroscience. Trieste, Italy. July 1993.

Colloquium, Department of Psychology, Oxford University. Oxford, England. February 1994.

Colloquium, Department of Psychology, University of California at San Diego. La Jolla, CA. June 1994.

Invited Speaker, Neural Information Processing Systems. November 1994.

Colloquium, Department of Computer Science, Colorado State University. Fort Collins, CO. January 1995.

Invited Speaker, Swedish Conference on Connectionism. Skovde, Sweden. February 1995.

Invited Symposium Speaker, Cognitive Neuroscience Society. San Francisco, CA. March 1995.

Colloquium, Santa Fe Institute. Santa Fe, NM. May 1995.

Faculty Lecturer, James S. McDonnell Summer Institute in Cognitive Neuroscience. Davis, CA. July 1995.

Invited Speaker, Lifestyle Technologies. Los Angeles, CA. August 1995.

Seminar, Department of Psychology, University of Toronto. Toronto, Ontario. October 1995.

Colloquium, Department of Psychology, McMaster University. Hamilton, Ontario. October 1995.

Colloquium, Department of Computer Science, Oregon Graduate Institute. Portland, OR. March 1996.

Invited Speaker, Apple Computer. Cupertino, CA. March 1996.

Invited Speaker, Conference on Neural Networks for Computing. Snowbird, UT. April 1996.

Invited Speaker, Montreal Workshop on Neural Networks. Montreal, Quebec. April, 1996.

Invited Speaker, Interval Research. San Jose, CA. May 1996.

Invited Speaker, Siemens Corporate Research. Princeton, NJ. June 1996.

Colloquium, Department of Cognitive Science, Johns Hopkins University. Baltimore, MD. June 1996.

Colloquium, Center for the Neural Bases of Cognition, Carnegie Mellon University. Pittsburgh, PA. March 1997.

Invited Speaker, Summer School on Adaptive Processing of Temporal Information. Vietri sul Mar, Italy. September 1997.

Colloquium, Institute for Research in Cognitive Science, University of Pennsylvania. Philadelphia, PA. October 1997.

Colloquium, Department of Psychology, University of Arizona. October 1997.

Colloquium, Systems Engineering, University of Pennsylvania. Philadelphia, PA. February 1998.

Invited Speaker, Neural Modeling of Brain and Cognitive Disorders Workshop, College Park, MD. June 1998.

Invited Participant. McDonnell Pew Program in Cognitive Neuroscience Annual Meeting, Montreal, PQ. June 1998.

Colloquium, Broadband Telecommunications Center, Georgia Institute of Technology. January 1999.

Colloquium, Department of Computer Science, University of Arizona. January 1999.

Colloquium, Department of Psychology, University of Iowa. March, 1999.

Colloquium, Department of Cognitive Science, University of California, Irvine. April, 1999.

Colloquium, AT&T Research Labs, Florham Park, NJ. June 1999.

Invited Participant. McDonnell-Pew Program in Cognitive Neuroscience Annual Meeting, San Diego, CA. June 1999.

Invited Speaker, International Joint Conference on Neural Networks. Washington, DC. July, 1999.

Colloquium, Department of Psychology, University of Pennsylvania. October, 1999.

Colloquium, Santa Fe Institute. Santa Fe, NM. February, 2000.

Colloquium, Department of Computer Science, University of Toronto. March, 2000.

Colloquium, Lucent Laboratories, Murray Hill, NJ. March 2000.

Invited Speaker, Fourth International Conference on Cognitive and Neural Systems, Boston, MA. May 2000.

Invited Speaker, Symposium on *Bayesian Models of Human Cognition*, Cognitive Science Society Conference, Philadelphia, PA. August 2000.

Invited Speaker, Workshop on *Network Models of Brain Function*, Banbury Center, NY. September 2000.

Invited Speaker, ESource Members' Forum (Energy Industry Conference), Colorado Springs, November 2000.

Colloquium, Department of Psychology, McMaster University. November, 2000.

Colloquium, Microsoft Research, Seattle. January, 2001.

Lecturer, *Complex Systems Summer School*, Santa Fe Institute. June, 2001.

Invited Participant, NSF KDI Workshop, New Orleans, LA. April 2002.

Colloquium, Department of Computer Science, UC San Diego, June 2002.

Lecturer, *Ninth International Summer School in Cognitive Science*, New Bulgaria University, Sofia. July, 2002.

Invited Visitor, Center for Cognitive Science, Macquarie University, Sydney, Australia. September-October 2002.

Colloquium, Department of Psychology, University of New South Wales, October 2002.

Invited Speaker, ESource Members' Forum (Energy Industry Conference), Colorado Springs, November 2002.

Invited Speaker, *International Neuroscience Summit 2002*, Berlin, Germany. November 2002.

Invited Speaker, *American Neuropsychiatric Association*, Bal Harbor, FL. February 2004.

Keynote Speaker, *International Conference on Cognitive Modeling*, Pittsburgh, PA. July 2004.

Colloquium, Intel Research, Berkeley, CA. February 2005.

Invited Speaker, *Modeling Integrated Cognitive Systems* (AFOSR workshop), Troy, NY. March 2005.

Invited Speaker, *Computation in Neural and Machine Vision Systems*, Toronto, ON. June 2005.

Keynote Speaker, *Intelligent Environments '05*. Colchester, UK. June 2005.

Colloquium, Department of Psychology, Macquarie University, Sydney. July 2005.

Invited Speaker, Psychology Department, UCSD. January 2006.

Keynote Speaker, *Unconventional Computing '06*. York University, UK. September 2006.

Invited Speaker, Department of Cognitive Science (COGS200). University of California, San Diego. May 2007.

Invited Speaker, Workshop on *Closing the gap between neurophysiology and behavior: A computational modeling approach*. University of Birmingham, UK. June 2007.

Colloquium, Department of Computer Science, University of Nevada, Reno. October 2007.

Invited Speaker, Department of Cognitive Science (COGS200), University of California, San Diego. November 2007.

Invited Speaker, Temporal Dynamics of Learning Center Annual Meeting, University of California, San Diego, February 2009.

Colloquium, Department of Psychology, Indiana University, October 2009.

Colloquium, School of Informatics, Indiana University, October 2009.

Colloquium, Department of Brain and Cognitive Sciences, University of Rochester, March 2010.

Colloquium, Department of Cognitive Science, University of California Merced, March 2010.

Colloquium, Department of Cognitive Science, University of California Irvine, April 2010.

Invited Speaker, Temporal Dynamics of Learning Center Annual Meeting, University of California, San Diego, January 2011.

Invited Speaker, Department of Cognitive Science (COGS200), University of California, San Diego, April 2011.

Invited Speaker, Temporal Dynamics of Learning Center Annual Meeting, University of California, San Diego, January 2012.

Invited Speaker, *Workshop on Optimal Teaching*, San Diego, May 2012.

Invited Lecturer, *European Summer School in Cognitive Science*, Sofia, Bulgaria, July 2012.

Invited Speaker, *Summer Symposium on Visual Search and Selective Attention*, Munich, Germany, July 2012.

Invited Speaker, *NSF Workshop on Computational Cognitive Modeling*, Arlington, VA, May 2013.

Cognitive Brownbag, Department of Psychology, UCSD, May 2013.

Colloquium, Google Brain, Mountain View, CA, October 2013.

Invited Speaker, Temporal Dynamics of Learning Center Annual Meeting, University of California, San Diego, February 2014.

Invited Speaker, Reasoning Minds, Houston TX, February 2014.

Invited Speaker, *Personalized Learning Workshop*, Houston TX, April 2014.

Invited Speaker, Temporal Dynamics of Learning Center Annual Meeting, University of California, San Diego, February 2015.

Invited Speaker, Machine Learning Group, Department of Computer Science, University of Toronto, June 2015.

Invited Speaker, *NIPS Workshop on Reasoning, Attention, and Memory*. Montreal, December 2015.

Invited Speaker, ICML Workshop on *Machine Learning for Digital Education and Assessment Systems*. New York, NY, June 2016.

Invited Speaker, *NIPS Symposium on Recurrent Neural Networks*, December 2016.

Invited Speaker, *NIPS Workshop on Machine Learning for Education*, December 2016.

Invited Speaker, *NIPS Workshop on Future of Interactive Learning Machines*, December 2016.

Invited Speaker, CogSci 200, Department of Cognitive Science, UCSD, February 2017.

Invited Speaker, Openstax Foundation, February 2017.

Invited Speaker, ECE Seminar Series, Rice University, February 2017.

Invited Speaker, Intelligent Systems Program, University of Pittsburgh, March 2017.

Cognitive Brownbag, Department of Psychology, UCSD, May 2017.

Keynote Speaker, *Learning Understanding Cognition Intelligence Data Science (LUCID) Conference*, Madison, WI, August 2017.

Invited Speaker, *Symposium on Deep Learning and Big Data, Society for Computers in Psychology*, Vancouver, BC, November 2017.

Invited Speaker, Oculus Research, Seattle WA, June 2018.

Invited Speaker, MPI-SWS Distinguished Lecture Series, Max Planck Institute, Kaiserslautern, Germany, June 2018.

Colloquium, Department of Computer Science, University of Montreal, Quebec, September 2018.

Colloquium, Institute for Intelligent Systems, University of Memphis, Memphis, TN, March 2019.

Keynote Speaker, Conference on Educational Data Mining, Montreal Quebec, July 2019.

Invited Speaker, *Exploring Consciousness Symposium*, McMaster University, Hamilton, Ontario, October 2019.

Invited Speaker, *Workshop on Memory, Modularity, and Attention*, Cosyne, Breckenridge, CO, February 2020.

Colloquium Speaker, Computer Science Department, Worcester Polytechnic Institute, November 2020.

Invited Speaker, *From Cells to Societies: Collective Learning Across Scales*, ICLR, April 2022.

Invited Speaker, *NSF Workshop on Augmenting Individual Intelligence*, October 2022.

Invited Speaker, *Workshop on Learning to Reason*, Bernoulli Center for Fundamental Studies, EPFL, December 2022.

[2024-09-13]