Brenden M. Lake

New York University · 60 5th Ave., Room 610 · New York, NY 10011 brenden@nyu.edu · (212) 998-3059 · cims.nyu.edu/~brenden/

Last updated April 2024

Positions

| 2024 – | Associate Professor of Psychology and Data Science |
|-------------|---|
| | Department of Psychology and Center for Data Science, New York University |
| 2017 — | Affiliated Faculty, Department of Computer Science and Center for Neural Science, NYU |
| 2017 - 2024 | Assistant Professor of Psychology and Data Science, New York University |
| 2018 - 2022 | Research Scientist, part time |
| | Meta Al |
| 2014 - 2017 | Moore-Sloan Data Science Fellow |
| | Center for Data Science, New York University |
| Education | |

2009 – 2014 Ph.D., Cognitive Science, Massachusetts Institute of Technology Thesis advisor: Joshua B. Tenenbaum 2008 – 2009 M.S., Symbolic Systems, Stanford University Thesis advisor: James L. McClelland 2005 – 2009 B.S., with University Distinction, Symbolic Systems, Stanford University

Publications (peer reviewed)

Preprints (trainees are underlined)

<u>Tartaglini, A. R.</u>, Feucht, S., Lepori, M. A., <u>Vong, W. V.</u>, Lovering, C., **Lake, B. M.**, Pavlick, E. (2023). Deep Neural Networks Can Learn Generalizable Same-Different Visual Relations. Preprint available on *arxiv:2310.09612*

Teehan, R., Lake, B. M., Ren, M. (2024). CoLLEGe: Concept Embedding Generation for Large Language Models. Preprint available on *arXiv:2403.15362*.

Articles – published or accepted (trainees are underlined)

Vong, W. K., Wang, W., Orhan, A. E., and Lake, B. M. (2024). Grounded language acquisition through the eyes and ears of a single child. *Science*, 383(6682), 504-511.

<u>Orhan, A. E.</u>, and Lake, B. M. (2024). Learning high-level visual representations from a child's perspective without strong inductive biases. *Nature Machine Intelligence*, 487.

Zhou, Y., Feinman, R., and Lake, B. M. (2024). Compositional diversity in visual concept learning. *Cognition*, 244, 105711.

Davidson, G., Orhan, A. E., and Lake, B. M. (2024). Spatial Relation Categorization in Infants and Deep Neural Networks. *Cognition*, 245, 105690.

- Orhan, A. E., Wang, W., Wang, A. N., Ren, M., and Lake, B. M. (2024). Self-supervised learning of video representations from a child's perspective. *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*
- Kumar, S., Marjieh, R., Zhang, B., Campbell, D., Hu, M. Y., Bhatt, U., Lake, B. M., and Griffiths, T. (2024). Comparing Abstraction in Humans and Machines Using Multimodal Serial Reproduction. *Proceedings* of the 46th Annual Conference of the Cognitive Science Society.
- Zhou, Y., Lake, B. M., and Williams, A. (2024). Compositional learning of functions in humans and machines. *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*
- <u>Qin, Y.</u>, <u>Wang, W.</u>, and **Lake, B. M.** (2024). A systematic investigation of learnability from single child linguistic input. *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*
- Li, W., Yasuda, S. C., Dillon, M. R., and Lake, B. M. (2024). A Machine Social Reasoning Benchmark Inspired by Infant Cognition. *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*
- Leong, C. and Lake, B. M. (2024). Prompting invokes expert-like downward shifts in GPT-4V's conceptual hierarchies. *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*
- Luo, K., Zhang, B., Xiao, Y., and Lake, B. M. (2024). Finding Unsupervised Alignment of Conceptual Systems in Image-Word Representations. *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*
- LeGris, S., Lake, B. M., Gureckis, T. M. (2024). Predicting Insight during Physical Reasoning. *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*
- Lake, B. M. and Baroni, M. (2023). Human-like systematic generalization through a meta-learning neural network. *Nature*, 623, 115-121.
- Wang, W., Vong, W. K., Kim, N., and Lake, B. M. (2023). Finding Structure in One Child's Linguistic Experience. *Cognitive Science*, 47, e13305.
- Lake, B. M. and Murphy, G. L. (2023). Word meaning in minds and machines. *Psychological Review*, 130, 401-431.
- Stojnic, G., <u>Gandhi, K.</u>, Yasuda, S., **Lake, B. M.**, and Dillon, M. R. (2023). Commonsense Psychology in Human Infants and Machines. *Cognition*, 235, 105406.
- Vong, W. K. and Lake, B. M. (2022). Cross-situational word learning with multimodal neural networks. *Cognitive Science*, 46, e13122.
- Tartaglini, A. R., Vong, W. K., and Lake, B. M. (2022). A Developmentally-Inspired Examination of Shape versus Texture Bias in Machines. In *Proceedings of the 44th Annual Conference of the Cognitive Science Society.*
- Davidson, G., Gureckis, T. M., and Lake, B. M. (2022). Creativity, Compositionality, and Common Sense in Human Goal Generation. *In Proceedings of the 44th Annual Conference of the Cognitive Science Society.*
- <u>Ruis, L.</u> and Lake, B. M. (2022). Improving Systematic Generalization Through Modularity and Augmentation. In *Proceedings of the 44th Annual Conference of the Cognitive Science Society.*
- <u>Feinman, R.</u> and Lake, B. M. (2021). Learning Task-General Representations with Generative Neuro-Symbolic Modeling. *International Conference on Learning Representations (ICLR)*.

- <u>Gandhi, K.</u>, Stojnic, G., Lake, B. M. and Dillon, M. R. (2021). Baby Intuitions Benchmark (BIB): Discerning the goals, preferences, and actions of others. *Advances in Neural Information Processing Systems* (*NeurIPS*) 34.
- <u>Nye, M.</u>, Tessler, M. H., Tenenbaum, J. B., and **Lake, B. M.** (2021). Improving Coherence and Consistency in Neural Sequence Models with Dual-System, Neuro-Symbolic Reasoning. *Advances in Neural Information Processing Systems (NeurIPS) 34.*
- Vedantam, R., Szlam, A., Nickel M., Morcos, A., and Lake, B. M. (2021). CURI: A Benchmark for Productive Concept Learning Under Uncertainty. *International Conference on Machine Learning (ICML)*.
- Zhou, Y. and Lake, B. M. (2021). Flexible compositional learning of structured visual concepts. In *Proceedings of the 43rd Annual Conference of the Cognitive Science Society.*
- <u>Tartaglini, A. R.</u>, <u>Vong, W. K.</u>, and **Lake, B. M.** (2021). Modeling artificial category learning from pixels: Revisiting Shepard, Hovland, and Jenkins (1961) with deep neural networks. In *Proceedings of the 43rd Annual Conference of the Cognitive Science Society.*
- Davidson, G. and Lake, B. M. (2021). Examining Infant Relation Categorization Through Deep Neural Networks. In *Proceedings of the 43rd Annual Conference of the Cognitive Science Society.*
- Johnson, A., <u>Vong, W. K.</u>, **Lake, B. M.** and Gureckis, T. M. (2021). Fast and flexible: Human program induction in abstract reasoning tasks. In *Proceedings of the 43rd Annual Conference of the Cognitive Science Society.*
- <u>Wang, Z.</u> and Lake, B. M. (2021). Modeling question asking using neural program generation. In *Proceedings of the 43rd Annual Conference of the Cognitive Science Society.*
- Lake, B. M. and Piantadosi, S. T. (2020). People infer recursive visual concepts from just a few examples. *Computational Brain & Behavior*, 3(1), 54-65.
- Lewis, M., Cristiano, V., Lake, B. M., Kwan, T., and Frank, M. C. (2020). The role of developmental change and linguistic experience in the mutual exclusivity effect. *Cognition*, 198.
- <u>Orhan, A. E.</u>, <u>Gupta, V. B.</u>, and **Lake, B. M.** (2020). Self-supervised learning through the eyes of a child. *Advances in Neural Information Processing Systems (NeurIPS)* 33.
- <u>Ruis, L.</u>, Andreas, J., Baroni, M. Bouchacourt, D., and Lake, B. M. (2020). A Benchmark for Systematic Generalization in Grounded Language Understanding. *Advances in Neural Information Processing Systems (NeurIPS)* 33.
- <u>Nye, M.</u>, Solar-Lezama, A., Tenenbaum, J. B., and **Lake, B. M.** (2020). Learning Compositional Rules via Neural Program Synthesis. *Advances in Neural Information Processing Systems (NeurIPS)* 33.
- <u>Gandhi, K.</u> and Lake, B. M. (2020). Mutual exclusivity as a challenge for deep neural networks. *Advances in Neural Information Processing Systems (NeurIPS)* 33.
- <u>Feinman, R.</u> and Lake, B. M. (2020). Generating new concepts with hybrid neuro-symbolic models. In *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*.
- Davidson, G. and Lake, B. M. (2020). Investigating simple object representations in model-free deep reinforcement learning. In *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2019). The Omniglot challenge: a 3-year progress report. *Current Opinion in Behavioral Sciences*, 29, 97-104.

- Lake, B. M. (2019). Compositional generalization through meta sequence-to-sequence learning. Advances in Neural Information Processing Systems (NeurIPS) 32.
- <u>Feinman, R.</u> and Lake, B. M. (2019). Learning a smooth kernel regularizer for convolutional neural networks. In *Proceedings of the 41st Annual Conference of the Cognitive Science Society*.
- Lake, B. M., Linzen, T., and Baroni, M. (2019). Human few-shot learning of compositional instructions. In *Proceedings of the 41st Annual Conference of the Cognitive Science Society.*
- <u>Rothe, A.</u>, Lake, B. M., and Gureckis, T. M. (2019). Asking goal-oriented questions and learning from answers. In *Proceedings of the 41st Annual Conference of the Cognitive Science Society.*
- Rothe, A., Lake, B. M., and Gureckis, T. M. (2018). Do people ask good questions? *Computational Brain & Behavior*, 1(1), 69-89.
- Lake, B. M., Lawrence, N. D., and Tenenbaum, J. B. (2018). The emergence of organizing structure in conceptual representation. *Cognitive Science*, 42(S3), 809-832.Info.]
- Lake, B. M. and Baroni, M. (2018). Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks. *International Conference on Machine Learning (ICML)*.
- <u>Feinman, R.</u> and Lake, B. M. (2018). Learning inductive biases with simple neural networks. In *Proceedings* of the 40th Annual Conference of the Cognitive Science Society.
- Loula, J., Baroni, M., and Lake, B. M. (2018). Rearranging the familiar: Testing compositional generalization in recurrent networks. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP.*
- Lake, B. M., Ullman, T. D., Tenenbaum, J. B., and Gershman, S. J. (2017). Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40, E253.
- <u>Rothe, A.</u>, Lake, B. M., and Gureckis, T. M. (2017). Question asking as program generation. *Advances in Neural Information Processing Systems (NeurIPS)* 30.
- <u>Rothe, A.</u>, Lake, B. M., and Gureckis, T. M. (2016). Asking and evaluating natural language questions. In *Proceedings of the 38th Annual Conference of the Cognitive Science Society.*
- <u>Cohen, A.</u> and Lake, B. M. (2016). Searching large hypothesis spaces by asking questions. In *Proceedings* of the 38th Annual Conference of the Cognitive Science Society.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266), 1332-1338.
- Monfort, M., Lake, B. M., Ziebart, B. D., Lucey, P., and Tenenbaum, J. B. (2015). Softstar: Heuristic-Guided Probabilistic Inference. Advances in Neural Information Processing Systems (NeurIPS) 28.
- Lake, B. M., Zaremba, W., Fergus, R. and Gureckis, T. M. (2015). Deep Neural Networks Predict Category Typicality Ratings for Images. In *Proceedings of the 37th Annual Conference of the Cognitive Science Society*.
- Lake, B. M., Lee, C.-y., Glass, J. R., and Tenenbaum, J. B. (2014). One-shot learning of generative speech concepts. In *Proceedings of the 36th Annual Conference of the Cognitive Science Society.*
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2013). One-Shot Learning by Inverting a Compositional Causal Process. *Advances in Neural Information Processing Systems (NeurIPS)* 26.

- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2012). Concept learning as motor program induction: A large-scale empirical study. In *Proceedings of the 34th Annual Conference of the Cognitive Science Society.*
- Lake, B. M., Salakhutdinov, R., Gross, J., and Tenenbaum, J. B. (2011). One shot learning of simple visual concepts. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*.
- Lake, B. M. and McClelland, J. L. (2011). Estimating the strength of unlabeled information during semisupervised learning. In *Proceedings of the 33rd Annual Conference of the Cognitive Science Society.*
- Lake, B. M. and Tenenbaum, J. B. (2010). Discovering Structure by Learning Sparse Graphs. In *Proceedings* of the 32nd Annual Conference of the Cognitive Science Society.
- Lake, B. M., Vallabha, G. K., and McClelland, J. L. (2009). Modeling unsupervised perceptual category learning. *IEEE Transactions on Autonomous Mental Development*, 1(1), 35-43.
- Lake, B. M., Vallabha, G. K., and McClelland, J. L. (2008). Modeling unsupervised perceptual category learning. In *Proceedings of the 7th International Conference on Development and Learning*.
- Lake, B. M. and Cottrell, G.W. (2005). Age of acquisition in facial identification: A connectionist approach. In *Proceedings of the 27th Annual Conference of the Cognitive Science Society*.

Funding

Research contracts, grants, and gifts

Lockheed Martin, \$110,000, "Neuro-symbolic reasoning," 2021-2022, PI.

Google gift, \$30,000, "Object-based representations for more human-like machine learning and reasoning," PI.

Meta, \$667,000, "Towards machines that actually help," 2020-2022, PI.

- DARPA, \$1,703,000, "Cognitive milestones for machine common sense," as Co-PI. \$836,000 in support of Lake's research.
- Meta, \$586,000, "Automatic concept learning from first person video and audio streams," 2018-2020, Pl.

Huawei gift, \$243,000, "Human-like question asking and learning for machine intelligence," 2018, PI.

NSF I-Corps for Learning, \$50,000, "Learning-to-learn with touchscreen technology," 2014-2015, PI.

Training grants

NSF NRT, \$3,000,000, "FUTURE Foundations, translation, and responsibility for data science impact," 2019-2024, as Co-PI.

Mentoring

Ph.D. theses supervised

Wentao Wang (Data Science, 2023-current, Co-supervised with Tal Linzen)

Solim LeGris (Psychology, 2022-current, Co-supervised with Todd Gureckis)

Guy Davidson (Data Science, 2019-current)

Yanli Zhou (Data Science, 2018-current, Meta Al Fellowship)

Reuben Feinman (Neural Science, 2017-2023, Google Fellowship, subsequently at Common Sense Machines)

Anselm Roth (Psychology, 2014-2019, Co-supervised with Todd Gureckis, subsequently a postdoc at Max Planck Institute for Human Development)

Postdoctoral research supervised Wai Keen Vong (2019-present)

Emin Orhan (2019-present)

Undergraduate and Masters research supervised Yulu Qin (2023-present) Wentao Wang (2021-2023, subsequently Ph.D. student at NYU Data Science) Wenjie Li (2021-present) Alexa Tartaglini (2019-present, Glushko Undergraduate Thesis Prize) Vaibhav Gupta (2019-2020, subsequently Amazon) Ziyun Wang (2018-2020, subsequently Tencent) Kanishk Gandhi (2018-2020, subsequently Ph.D. student at Stanford CS) Alexander Cohen (2015-2017, high school mentee, subsequently Ph.D. student at MIT Math)

Honors

Outstanding Paper Award, *Computational Brain & Behavior*, 2019. Top 35 Innovators Under 35 (TR35), *MIT Technology Review*, 2018. Top 10 World Changing Ideas, *Scientific American*, 2016. Robert J. Glushko Prize for Outstanding Doctoral Dissertation in Cognitive Science, 2015. National Science Foundation (NSF) Graduate Research Fellowship, 2011-2014. Angus MacDonald Award for Excellence in Teaching Undergraduate Students, 2010. Singleton Presidential Fellowship, Massachusetts Institute of Technology, 2009-2010. Best Paper Award, *International Conference on Development and Learning (ICDL)*, 2008.

Teaching

Computational cognitive modeling, PSYCH-GA 3405.002 / DS-GA 1016, NYU (Spring 2018-2023). Lab in Cognition and Perception, PSYCH-UG (Fall 2021, 2023). Categories and concepts, PSYCH-GA 2207 (Fall 2019, 2021, 2023). Advancing AI through cognitive science, PSYCH-GA 3405.001 / DS-GA 3001.014 (Spring 2018, 2019). Practical Training for Data Science, DS-GA 1009 (Fall 2015, 2016).

Service

University

Co-Director, NYU Minds, Brains, and Machines Initiative, 2021-present.

Joint Research Committee, Neural Science, Technology, & AI, KAIST-NYU Partnership, 2022-present.

Committee on Minds and Machines, NYU Faculty of Arts and Sciences, 2019-2020.

Committee on Information, Technology, and Library Services, NYU Faculty of Arts and Sciences, 2019-2022.

Departmental

Search committee, Center for Data Science, Artificial Intelligence, NYU, 2022-2023. Search committee (Co-chair), Psychology Department and Center for Data Science, NYU, 2021-2022. Search committee, Psychology Department, Cognitive science of language, NYU, 2021-2022. Search committee, Center for Data Science and Center for Neural Science, NYU, 2019-2020. Search committee, Center for Data Science and Center for Neural Science, NYU, 2018-2019. Search committee, Center for Data Science and Computer Science Department, NYU, 2017-2018. Organizer, Data Science Lunch Seminar Series, Fall 2015. Organizer, NYU Concepts and Categories (ConCats) seminars, Fall 2014-Spring 2015.

External

- Reviewing: Nature, Science, PNAS, Nature Machine Intelligence, Nature Human Behavior, Nature Communications, Behavioral and Brain Sciences, Psychological Review, Journal of Experimental Psychology: General, Cognition, Cognitive Science, Trends in Cognitive Science, Scientific Reports, Psychonomic Bulletin & Review, Developmental Psychology, Infant Behavior & Development, Current Directions in Psychological Science, Memory & Cognition, Vision Research, Current Opinion in Behavioral Sciences, Big Data, Frontiers in Psychology, Advances in Neural Information Processing Systems (NeurIPS), International Conference on Learning Representations (ICLR), International Conference on Artificial Intelligence and Statistics (AISTATS), Annual Conference of the Cognitive Science Society (CogSci), and National Science Foundation (NSF).
- Co-Organizer, Understanding Human and Machine Intelligence: A Workshop on Cognitive Science and AI, May 2019.

Co-Organizer, NIPS Workshop on Cognitively Informed Artificial Intelligence, Dec. 2017.

Mentor, Artificial Intelligence NexusLab Incubator, New York City, 2016-2018.

Selected invited talks (2017-onwards)

Invited talk, Cognitive Control of Action Workshop, Princeton University, Mar. 2024. Harvard University, Department of Psychology and Kempner Institute, Mar. 2024. Princeton University, Center for Statistics and Machine Learning, Feb. 2024. Johns Hopkins, Computer Science Seminar Series, Feb. 2024. Vassar College, Cognitive Science Department and the Data Science and Society Program, Feb. 2024. Invited Speaker, Machine Learning Meetup, London, Jan., 2024. Developing Minds global lecture, Developmental AI Task Force, IEEE, Nov. 2023. KAIST, Symposium for Launching the Department of Brain and Cognitive Sciences, South Korea, Nov. 2023. New York University, Swartz Seminar on Computational and Theoretical Neuroscience, Apr. 2023. TU Darmstadt, Center for Cognitive Science Colloquium, Apr. 2023. Columbia University, University Seminar on Cognitive and Behavioral Neuroscience, Feb. 2023. McGill University, Cognitive Informatics Seminar, Feb. 2023. Invited Speaker, IBM Neuro-Symbolic AI Workshop, Jan. 2023. Invited Speaker, The Challenge of Compositionality for AI, June 2022. Carnegie Mellon and U. Pittsburg, Center for Neural Basis of Cognition Colloquium, Feb. 2022. Stanford University, FriSem Cognitive Area Seminar, Feb 2022. Tufts University, Cognitive and Brain Science Lecture, November 2021. Facebook Reality Labs, AI Seminar Series, Sept. 2021. Stanford University, Natural Language Processing Seminar, July 2021. Invited Speaker, Annual Meeting of the Cognitive Science Society (CogSci), Workshop on Symbolic and sub-symbolic systems in people and machines, July 2021. Hebrew University of Jerusalem, Cognitive Science Departmental Seminar, June 2021. Invited Speaker, International Conference on Learning Representations (ICLR), Workshop on generalization beyond the training distribution, May 2021. MIT, NeuroSym Seminar, April 2021. Invited Speaker, European Chapter of the Association for Computational Linguistics (EACL), LANTERN workshop, April 2021. Invited Speaker, Society for Research in Child Development (SRCD), Symposium on Reorienting the study of language learning using naturalistic data, April 2021.

Johns Hopkins, Cognitive Science Colloquium, April 2021.

Columbia University, Seminars in Society and Neuroscience, Mar. 2021.

Invited Speaker, ELLIS workshop in Meta-learning in AI and cognitive science, Mar. 2021.

Brown University, Data Science Initiative, Dec. 2020.

Invited Speaker, Machine Learning Meetup, London, Sept. 2020.

- Max Planck Institute, Tübingen, RLDM meeting, Aug. 2020.
- University of Bristol, Seminar on Generalization in Mind and Machine, July 2020.
- Princeton University, Institute for Advanced Study / Neuroscience Institute, Seminar on ML and Neuroscience, Feb. 2020.
- Invited Speaker, Neural Information Processing Systems (NeurIPS), Workshop on Meta Learning, Dec. 2019.
- Columbia University, Center for Theoretical Neuroscience Seminar Series, Nov. 2019.
- Harvard University, Cognition, Brain, & Behavior Research Seminar, Oct. 2019.
- Featured Speaker, NYU Trustee Spouses Dinner, Feb. 2019.
- Invited Speaker, Intelligent Computing and Communications Technologies Workshop, New Brunswick, NJ, Nov. 2018.
- Invited Speaker, Artificial Intelligence and the Barrier of Meaning, Santa Fe Institute, Oct. 2018.
- Invited Speaker, Symposium in honor of James L. McClelland: Parallel Distributed Processing and the Emergence of an Understanding of Mind, Princeton University, Sept. 2018.
- Invited Speaker, EmTech MIT, MIT Technology Review, Sept. 2018.
- Invited Speaker, Brain Visions: geometry, art, and symbolism, Le Plan-de-la-Tour, France, Sept. 2018.
- Invited Speaker, Track on the Future of AI, Joint Multi-Conference on Human-Level Artificial Intelligence, Aug. 2018.
- Invited Speaker, International Conference on Machine Learning (ICML), Workshop on Neural Abstract Machines & Program Induction, July 2018.
- Invited Speaker, Workshop on Deep, fast and shallow learning in humans and machines, Indiana University, Bloomington, May 2018.
- Invited Speaker, EmTech Digital, MIT Technology Review, March 2018.
- Princeton University, PDP Group seminar, Feb. 2018.
- Thomas J. Watson Research Center, Yorktown Heights, NY, Oct. 2017.
- Keynote Speaker, NYU Computational Neuroscience Symposium, June 2017.