

# Mark K Ho

Assistant Professor  
New York University  
Department of Psychology

Meyer Hall  
4 Washington Place  
New York, NY 10003

[mark.ho@nyu.edu](mailto:mark.ho@nyu.edu)  
[codec-lab.github.io](https://github.com/codec-lab)

## Employment

Assistant Professor   Department of Psychology, New York University	2024 - Present
Affiliated Faculty   Center for Data Science, New York University	2024 - Present
Affiliated Faculty   Department of Computer Science, New York University	2025 - Present
Assistant Professor   Department of Computer Science, Stevens Institute of Technology	2023 - 2024
Faculty Fellow   Center for Data Science, New York University	2022 - 2023
Postdoctoral Researcher   Department of Computer Science, Princeton University	2020 - 2022
Postdoctoral Researcher   Department of Electrical Engineering and Computer Sciences, University of California, Berkeley	2018 - 2020

## Education

Ph.D., Cognitive Science, Brown University	2018
M.S., Computer Science, Brown University	2015
M.S., Cognitive Science, Brown University	2014
B.A., Philosophy, Princeton University	2011
Bronx High School of Science	2007

## Grants and Funding

CISE Research Initiation Initiative (CRII), National Science Foundation (\$175,000)	2024 - 2026
“Optimal Teaching Strategies for Humans in Hybrid Domains,” Joint Project with Toyota Research Institute (\$420,440)	2024 - 2026
“Predicting Algorithmic Trust at Scale,” Defense Advanced Research Projects Agency (DARPA) contract with Princeton, Boston University, and Cornell (total: \$1,722,222; NYU sub-award: \$887,716)	2025 - 2026

“Cognition-Aware Co-adaptation with Humans,” Joint Project with Toyota Research Institute (\$938,419) 2026 - 2029

Office of Naval Research (ONR) Young Investigator Award (\$750,000) 2026-2029

## Fellowships

National Science Foundation Graduate Research Fellowship 2014 - 2019

Diverse Intelligences Summer Institute Fellow 2019

Professor Lorrin A. Riggs Graduate Student Dissertation Fellowship 2017

Brown University Open Graduate Education Fellowship 2014 - 2015

## Journal Articles

Harootonian, S., Griffiths, T. L., Niv, Y., and Ho, M. K. (in press). Mentalizing and heuristics as distinct cognitive strategies in human teaching. *Nature Human Behavior*.

Dubey, R., Ho, M. K., Mehta, H., and Griffiths, T. (in press). Aha! moments correspond to meta-cognitive prediction errors. *Cognition*.

Schulze, C., Aka, A., Bartels, D. M., Bucher, S. F., Embrey, J. R., Gureckis, T. M., Häubl, G., Ho, M. K., Krajbich, I., and Moore, A. K. (2025). A timeline of cognitive costs in decision-making. *Trends in Cognitive Sciences*.

Correa, C. G., Sanborn, S., Ho, M. K., Callaway, F., Daw, N. D., and Griffiths, T. L. (2025). Exploring the hierarchical structure of human plans via program generation. *Cognition*, page 105990.

Arumugam\*, D., Ho\*, M. K., Goodman, N. D., and Van Roy, B. (2024). Bayesian reinforcement learning with limited cognitive load. *Open Mind*, pages 395–438.

Collins, K. M., Sucholutsky, I., Bhatt, U., Chandra, K., Wong, L., Lee, M., Zhang, C. E., Zhi-Xuan, T., Ho, M., Mansinghka, V., Weller, A., Tenenbaum, J. B., and Griffiths, T. L. (2024). Building machines that learn and think with people. *Nature Human Behavior*.

Allen, K. R., Brändle, F., Botvinick, M. M., Fan, J., Gershman, S. J., Gopnik, A., Griffiths, T. L., Hartshorne, J. K., Hauser, T. U., Ho, M. K., and , e. a. (2024). Using games to understand the mind. *Nature Human Behavior*.

Ho, M. K., Cohen, J. D., and Griffiths, T. L. (2023). Rational simplification and rigidity in human planning. *Psychological Science*.

Sumers, T., Ho, M. K., Griffiths, T. L., and Hawkins, R. (2023). Reconciling truthfulness and relevance via decision-theoretic utility. *Psychological Review*.

Correa, C. G., Ho, M. K., Callaway, F., Daw, N. D., and Griffiths, T. L. (2023). Humans decompose tasks by trading off utility and computational cost. *PLOS Computational Biology*, 19:1–31.

Sumers, T. R., Ho, M. K., Hawkins, R. D., and Griffiths, T. L. (2023). Show or tell? teaching with language outperforms demonstration but only when context is shared. *Cognition*.

Ho, M. K., Abel, D., Correa, C. G., Littman, M. L., Cohen, J. D., and Griffiths, T. L. (2022). People construct simplified mental representations to plan. *Nature*.

Ho, M. K., Saxe, R., and Cushman, F. (2022). Planning with theory of mind. *Trends in Cognitive Sciences*.

Ho, M. K. and Griffiths, T. L. (2022). Cognitive science as a source of forward and inverse models of human decisions for robotics and control. *Annual review of Control, Robotics, and Autonomous Systems*.

Ho, M. K., Cushman, F., Littman, M. L., and Austerweil, J. L. (2021). Communication in action: Planning and interpreting communicative demonstrations. *Journal of Experimental Psychology: General*.

Gates, V., Callaway, F., Ho, M. K., and Griffiths, T. (2021). A rational model of people’s inferences about others’ preferences based on response times. *Cognition*, 217:104885.

Sarin, A., Ho, M. K., Martin, J. W., and Cushman, F. A. (2021). Punishment is organized around principles of communicative inference. *Cognition*, 208:104544.

Ho, M. K., Abel, D., Griffiths, T. L., and Littman, M. L. (2019). The value of abstraction. *Current Opinion in Behavioral Sciences*, 29:111–116.

Ho, M. K., Cushman, F., Littman, M. L., and Austerweil, J. L. (2019). People teach with rewards and punishments as communication, not reinforcements. *Journal of Experimental Psychology: General*, 148:520–549.

Ho, M. K., MacGlashan, J., Littman, M. L., and Cushman, F. (2017). Social is special: A normative framework for teaching with and learning from evaluative feedback. *Cognition*, 167:91–106.

## Articles Under Review

Sucholutsky, I., Collins, K. M., Malaviya, M., Jacoby, N., Liu, W., Sumers, T. R., Korakakis, M., Bhatt, U., Ho, M. K., Tenenbaum, J. B., Love, B., Pardos, Z. A., Weller, A., and Griffiths, T. L. (under review). Representational alignment supports effective machine teaching

Kuperwajs, I., Ho, M. K., and Ma, W. J. (under review). Heuristics for meta-planning from a normative model of information search

## Refereed Conference Papers<sup>1</sup>

Banerjee, S., Cornelisse, D., Gopinath, D., Sumner, E., DeCastro, J., Rosman, G., Vinitzky, E., and Ho, M. K. (2025). Estimating cognitive biases with attention-aware inverse planning. In *Advances in Neural*

---

<sup>1</sup>Note that in computer science, conference papers are the primary form of publication.

*Information Processing Systems 38*, pages XX–XX. Curran Associates, Inc. **\*\*Selected for spotlight presentation\*\* (688 out of 21,575 full paper submissions)**

Zhang, M. Y., Leslie, S.-J., Rhodes, M., and Ho, M. K. (2025). Learning about inductive potential from generic statements. In *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

Su, E., Ho, M. K., and Gureckis, T. M. (2025). Integration of language and experience via the instructed bandit task. In *Proceedings of the 47th Annual Conference of the Cognitive Science Society*

Abel, D., Ho, M. K., and Harutyunyan, A. (2024). Three dogmas of reinforcement learning. In *Proceedings of the Reinforcement Learning Conference*

Rane, S., Ho, M. K., Sucholutsky, I., and Griffiths, T. L. (2024). Concept alignment as a prerequisite for value alignment. In *Proceedings of the 46rd Annual Conference of the Cognitive Science Society*

Pennisi, B., Budiono, R., Gureckis, T. M., and Ho, M. K. (2024). Investigating flexible role binding in ai agents. In *Proceedings of the 46rd Annual Conference of the Cognitive Science Society*

Peng, A., Netanyahu, A., Ho, M. K., Shu, T., Bobu, A., Shah, J., and Agrawal, P. (2023). Diagnosis, feedback, adaptation: A human-in-the-loop framework for test-time policy adaptation. In *International conference on machine learning*. PMLR. **\*Less than 28% of submissions accepted\***

Sumers, T. R., Hawkins, R. D., Ho, M. K., Griffiths, T. L., and Hadfield-Menell, D. (2022). How to talk so your robot will learn: Instructions, descriptions, and pragmatics. In *Advances in Neural Information Processing Systems 35*, pages XX–XX. Curran Associates, Inc

Abel, D., Dabney, W., Harutyunyan, A., Ho, M. K., Littman, M., Precup, D., and Singh, S. (2021). On the expressivity of markov reward. In *Advances in Neural Information Processing Systems 34*, pages XX–XX. Curran Associates, Inc. **\*Outstanding Paper Award (6 out of 9122 full paper submissions)\***

Sumers, T. R., Hawkins, R. D., Ho, M. K., and Griffiths, T. L. (2021). Extending rational models of communication from beliefs to actions. In *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*

Wu, C. M., Ho, M. K., Kahl, B., Leuker, C., Meder, B., and Kurvers, R. H. (2021). Specialization and selective social attention establishes the balance between individual and social learning. In *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*

Sumers, T. R., Ho, M. K., Hawkins, R. D., Narasimhan, K., and Griffiths, T. L. (2021). Learning rewards from linguistic feedback. In *Proceedings of the aai conference on artificial intelligence*, volume 35

Ho, M. K., Abel, D., Cohen, J., Littman, M., and Griffiths, T. (2020). The efficiency of human cognition reflects planned information processing. In *Proceedings of the aai conference on artificial intelligence*, volume 34, pages 1300–1307. **\*Selected for oral presentation (less than 5% of papers)\***

Correa\*, C. G., Ho\*, M. K., Callaway, F., and Griffiths, T. L. (2020). Resource-rational task decomposition to minimize planning costs. In Denison., S., Mack, M., Xu, Y., and Armstrong, B., editors, *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*, pages 2974–2980. Cognitive Science Society

Sumers, T. R., Ho, M. K., and Griffiths, T. L. (2020). Show or tell? demonstration is more robust to changes in shared perception than explanation. In Denison, S., Mack, M., Xu, Y., and Armstrong, B., editors, *Proceedings of the 42nd Annual Conference of the Cognitive Science Society*, pages 3073–3079. Cognitive Science Society

Wang, G., Trimbach, C., Lee, J. K., Ho, M. K., and Littman, M. L. (2020). Teaching a robot tasks of arbitrary complexity via human feedback. In *Proceedings of the 2020 acm/ieee international conference on human-robot interaction*, pages 649–657

Ho, M. K., Korman, J., and Griffiths, T. L. (2019). The computational structure of unintentional meaning. In Goel, A., Seifert, C., and Freksa, C., editors, *Proceedings of the 41st Annual Conference of the Cognitive Science Society*, pages 1915–1921. Cognitive Science Society

Carroll, M., Shah, R., Ho, M. K., Griffiths, T., Seshia, S., Abbeel, P., and Dragan, A. (2019). On the utility of learning about humans for human-ai coordination. In Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché Buc, F., Fox, E., and Garnett, R., editors, *Advances in neural information processing systems*, volume 32. Curran Associates, Inc

Vazquez-Chanlatte, M., Jha, S., Tiwari, A., Ho, M. K., and Seshia, S. (2018). Learning task specifications from demonstrations. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R., editors, *Advances in neural information processing systems*, volume 31. Curran Associates, Inc

Ho, M. K., Littman, M. L., Cushman, F., and Austerweil, J. L. (2018). Effectively learning from pedagogical demonstrations. In Kalish, C., Rau, M., Rogers, T., and Zhu, J., editors, *Proceedings of the 40th Annual Conference of the Cognitive Science Society*, pages 505–510, Austin, TX. Cognitive Science Society

Ho, M. K., Littman, M. L., and Austerweil, J. L. (2017). Teaching by intervention: Working backwards, undoing mistakes, or correcting mistakes? In Gunzelmann, G., Howes, A., Tenbrink, T., and Davelaar, J., editors, *Proceedings of the 39th Annual Conference of the Cognitive Science Society*, pages 526–531, Austin, TX. Cognitive Science Society

MacGlashan, J., Ho, M. K., Loftin, R., Peng, B., Wang, G., Roberts, D. L., Taylor, M. E., and Littman, M. L. (2017). Interactive learning from policy-dependent human feedback. In *International conference on machine learning*, pages 2285–2294. PMLR

Ho, M. K., Littman, M., MacGlashan, J., Cushman, F., and Austerweil, J. L. (2016). Showing versus doing: Teaching by demonstration. In Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems 29*, pages 3027–3035. Curran Associates, Inc. **\*Selected for oral presentation (less than 2% of papers)\***

Ho, M. K., MacGlashan, J., Greenwald, A., Littman, M. L., Hilliard, E., Trimbach, C., Brawner, S., Tenenbaum, J. B., Kleiman-Weiner, M., and Austerweil, J. L. (2016). Feature-based joint planning and norm learning in collaborative games. In Papafragou, A., Grodner, D., Mirman, D., and Trueswell, J. C., editors, *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, pages 1158–1163, Austin, TX. Cognitive Science Society

Kleiman-Weiner, M., Ho, M. K., Austerweil, J. L., Littman, M. L., and Tenenbaum, J. B. (2016). Coordinate to cooperate or compete: Abstract goals and joint intentions in social interaction. In Papafragou,

A., Grodner, D., Mirman, D., and Trueswell, J. C., editors, *Proceedings of the 38th Annual Conference of the Cognitive Science Society*, pages 1679–1684, Austin, TX. Cognitive Science Society

Ho, M. K., Littman, M. L., Cushman, F., and Austerweil, J. L. (2015). Teaching with rewards and punishments: Reinforcement or communication? In Noelle, D., Dale, R., Warlaumont, A. S., Yoshimi, J., Matlock, T., Jennings, C. D., and Maglio, P. P., editors, *Proceedings of the 37th Annual Conference of the Cognitive Science Society*, pages 920–925, Austin, TX. Cognitive Science Society

## Book Chapters

Chater, N., Griffiths, T. L., and Ho, M. K. (2024). Chapter 7: From probabilities to actions. In Griffiths, T. L., Chater, N., and Tenenbaum, J., editors, *Bayesian Models of Cognition: Reverse Engineering the Mind*. MIT Press, Cambridge

Cushman, F. A., Sarin, A., and Ho, M. K. (2021). Punishment as communication. In Doris, J. and Vargas, M., editors, *Oxford handbook of moral psychology*. Oxford University Press, Oxford

## Workshop Papers, Posters, and Extended Abstracts

Ho, M. K. & Gureckis, T. (2023, August) Learning from Language and Experience. Extended abstract and poster presented at Cognitive Computational Neuroscience. Oxford, England.

Rane, S., Ho, M. K., Sucholustsky, I., & Griffiths, T. (2023, July) Concept Alignment as a Prerequisite for Value Alignment. Workshop paper at Workshop on Social Intelligence in Humans and Robots. Robotics: Science and Society (RSS).

Arumugam, D., Ho, M. K., Goodman, N. D., and Van Roy, B. (2022). On rate-distortion theory in capacity-limited cognition and reinforcement learning. In *NeurIPS workshop on information-theoretic principles in cognitive systems (infocog)*

Abel, D., Barreto, A., Bowling, M., Dabney, W., Hansen, S., Harutyunyan, A., Ho, M. K., Kumar, R., Littman, M. L., Precup, D., and Singh, S. (2022). Expressing non-markov reward to a markov agent. In *Proceedings of the conference on reinforcement learning and decision making*

Ho, M. K., Cohen, J.D. & Griffiths, T. (2022, August) Construal Set Selection and Rigidity in Planning. Extended Abstract presented at Cognitive Computational Neuroscience. San Francisco, California.

Harootyan S., Ho, M. K., Klevak, N., & Niv, Y (2022, August) The best advice you can give. Extended Abstract presented at Cognitive Computational Neuroscience. San Francisco, California.

Ho, M. K., Abel, D., Cohen, J. D., Littman, M. L. & Griffiths, T. L. (2019, September) Optimal planning to plan: People partially plan based on plan specificity. Extended Abstract presented at Cognitive Computational Neuroscience. Berlin, Germany.

Seshia, S, Griffiths, T., Ho, M. K. & Vazquez-Chanlatte, Marcell (2019, November) Learning and Teaching Task Specifications from Demonstrations. Poster presented at NSF Cyber-Physical Systems PI Meeting. Washington, DC.

Vazquez-Chanlatte, M, Ho, M. K., Griffiths, T., Seshia, S. (2018, December) Communicating Compositional and Temporal Specifications by Demonstration. Poster presented at the 2nd IFAC Conference on Cyber-Physical and Human-Systems. Miami, FL.

Ho, M. K., Sanborn, S., Callaway, F., Bourgin, D., & Griffiths, T. (2018, September). Human Priors in Hierarchical Program Induction. Extended Abstract presented at Cognitive Computational Neuroscience. Philadelphia, Pa.

Ho, M. K., Littman, M., MacGlashan, J., Cushman, F., & Austerweil, J. L. (2017, September) Showing versus Doing: Teaching by Demonstration. Poster presented at the Inaugural Conference on Cognitive Computational Neuroscience, New York City, New York.

Kleiman-Weiner, M., Ho, M. K., Austerweil, J. L., Littman, M. L., & Tenenbaum, J. (2017, June). Learning to Cooperate and Compete. Poster presented at the Reinforcement Learning and Decision Making conference, Ann Arbor, Michigan, USA. **\*\*Best paper award\*\***

Ho, M. K., Littman, M. L., MacGlashan J., Cushman F., & Austerweil J. L.. (2017, March). Human Teaching by Demonstration: Showing versus Doing Reinforcement Learning Tasks. Poster and talk presented at the 11th Annual Machine Learning Symposium of the New York Academy of Sciences, New York City, New York. **\*\*Top presentation award\*\***

Ho, M. K., Littman, M.L., Cushman, F. & Austerweil, J. L. (2016, January). Generous Teachers: Evaluative Feedback as Communication. Poster presented at the Annual meeting of the Society for Personality and Social Psychology, San Diego, California.

Ho, M. K., Littman, M. L., Cushman, F. & Austerweil, J. L. (2015, June). Evaluative Feedback: Reinforcement or Communication? Poster selected for a 2 minute spotlight data blitz and presented at the Reinforcement Learning and Decision Making conference, Edmonton, Alberta, Canada.

Ho, M. K. & Cushman, F. (2013, August). Modeling Social Learning and Working Memory Use. Poster presented at the 35th Annual Conference of the Cognitive Science Society. Berlin, Germany.

## Talks and Symposia

Ho, M. K. (2026, April). “The construction of mental representations in human planning” Invited talk at UT Southwestern Computational Biology Seminar. Dallas, Texas.

Ho, M. K. (2025, November). “The construction of mental representations in human planning” Invited talk at UC Berkeley Cognition colloquium. Berkeley, California.

Ho, M. K. (2025, September). “The construction of mental representations in human planning” Invited talk at Indiana University CogSci colloquium. Bloomington, Indiana.

Ho, M. K. (2025, August). “Why might having a mind be useful?” Invited talk at the *Finding the frame* workshop. Reinforcement Learning Conference (RLC). Edmonton, Alberta.

Ho, M. K. (2025, July). “Reasoning and representations in human planning” Invited talk at the *Reasoning across minds and machines* workshop. Annual meeting of the Cognitive Science Society. San Francisco, California.

Ho, M. K. (2025, July). “Meta-reasoning about problem spaces” Invited talk at the *Meta-reasoning: deciding which game to play, which problem to solve, and when to quit* workshop. Annual meeting of the Cognitive Science Society. San Francisco, California.

Ho, M. K. (2025, June). “Flexibility, Rigidity, and Representational Alignment” . Invited talk at the Representational Alignment and Aging workshop (<https://sites.google.com/brown.edu/raa/>). Reinforcement Learning and Decision-Making Conference 2025. Dublin, Ireland.

Ho, M. K. (2024, November). “The computational structure of social cognition” . Invited talk at the Social and Affective Processes Seminar. Columbia University.

Ho, M. K. (2024, October). “Mentalizing and Mental Effort in Social Cognition” . Invited talk at the Consortium for Interacting Minds. Dartmouth College.

Ho, M. K. (2024, June). “Abstraction in Humans and Machines,” Multimodal Algorithmic Reasoning Workshop Keynote Talk (alternate for Tom Griffiths). Conference on Computer Vision and Pattern Recognition (CVPR).

Ho, M. K. (2024, June). “Representation, Computation, and Costs in Human Planning” Invited virtual talk to the Center for Cognitive Computation of the Central European University.

Ho, M. K. (2024, May). “Construction of Mental Representations in Human Planning,” Stanford Vision and Learning Lab (PIs: Fei-Fei Li and Jiajun Wu).

Ho, M. K. (2024, January). “The computational structure of social cognition” Department of Psychology, New York University.

Ho, M. K. (2023, December). “Construction of Mental Representations in Human Planning,” Emory University (PIs: Peter Hitchcock, Michael Treadway).

Ho, M. K. (2023, November). “Motivation, Cognition, and Computation,” Motivation Lab, New York University (PIs: Gabriele Oettingen, Peter Gollwitzer).

Ho, M. K. (2023, August). “Algorithmic, Representational, and Information Theoretic Costs in Decision-Making” Invited Talk at Cognitive Costs in Decision Making Session. The 12th Triennial Invitational Choice Symposium.

Ho, M. K. (2023, July). “Theory of Mind, Models of Cognition, and Social Interaction” Invited Talk at Workshop on Theory-of-Mind. International Conference On Machine Learning (ICML).

Ho, M. K. (2023, July). “Rationality, Computation, and Making Sense of Intelligence” Invited Talk at Diverse Intelligences Summer Institute (DISI). St. Andrews, Scotland.

Ho, M. K. (2023, July). “Cognitive Science as a Source of Design Principles for Interactive Machine Learning” Invited Talk at Social Intelligence in Humans and Robots Workshop. Robotics: Science and Systems Conference (RSS).

Ho, M. K. (2023, July). “Abstraction in Perception and Action” Invited Panelist at Abstractions Workshop. Annual Meeting of the Cognitive Science Society.

Ho, M. K. (2023, February). “Construction of Mental Representations in Human Planning” Consciousness Club seminar series. Wellcome Centre for Human Neuroimaging, University College London.

Ho, M. K. (2022, December). “Artificial Intelligence, Natural Stupidity, and Resource Rational Cognition” Invited Talk at Social Intelligence in Humans and Robots. Conference On Robot Learning (CORL).

Ho, M. K. (2022, November). “Construction of Mental Representations in Human Planning” Center for Computational Psychiatry Speaker Series. Icahn School of Medicine at Mount Sinai.

Ho, M. K. (2022, October). “World Models, Mind Models, and Planning” The Building Blocks of Human World Knowledge workshop. MIT.

Ho, M. K. (2022, September). “Construction of Mental Representations in Human Planning” Harvard Business School (PIs: Amit Goldenberg & Jillian Jordan).

Ho, M. K. (2022, September). “Construction of Mental Representations in Human Planning” Social and Cognitive Computational Neuroscience Lab (PI: Stefano Anzellotti). Boston College.

Ho, M. K. (2022, September). “Construction of Mental Representations in Human Planning” Basis Research Institute (PI: Zenna Tavares). Zuckerman Institute and Data Science Institute of Columbia University.

Ho, M. K. (2022, September). “Construction of Mental Representations in Human Planning” Concepts and Categories (ConCats) Seminar. New York University.

Ho, M. K. (2022, September). “Construction of Mental Representations in Human Planning” Center for Data Science Lunch Seminar. New York University.

Ho, M. K. (2022, August). “Construal Set Selection and Rigidity in Planning” Cognitive Computational Neuroscience (CCN) Conference, San Francisco, CA.

Harootyan, S., Niv, Y., Ho, M. K. (2022, August). “The Best Advice You Can Give” Cognitive Computational Neuroscience (CCN) Conference, San Francisco, CA.

Ho, M. K. (2022, July). “Construction of Mental Representations in Human Planning.” Causal Cognition Lab (PI: David Lagnado). University College London.

Ho, M. K. (2022, July). “Construction of Mental Representations in Human Planning.” DeepMind. London, UK.

Ho, M. K. (2022, July). “Construction of Mental Representations in Human Planning.” Max Planck Institute for Biological Cybernetics. Tübingen, Germany.

Ho, M. K. (2022, July). “Communicative Decision-Making and Interactive Teaching.” Invited Talk at the Computational Summer school on Modeling Social and collective behavior (COSMOS). Konstanz, Germany.

Ho, M. K. (2022, June). “Construction of Mental Representations in Human Planning.” Invited talk at the Neurosymbolic Expeditions Project Virtual Seminar (PI: Armando Solar-Lezama). MIT.

Ho, M. K. (2022, June). “Construction of Mental Representations in Human Planning.” Facebook Reality Labs.

Ho, M. K. (2022, June). “Construction of Mental Representations in Human Planning.” Microsoft Research. New York City.

Abel, David, Barreto, André, Bowling, Michael, Dabney, Will, Hansen, Steven, Harutyunyan, Anna, Ho, M. K., Kumar, Ramana, Littman, Michael L, Precup, Doina and others (2022, June). “Expressing Non-Markov Reward to a Markov Agent.” Oral Presentation at the Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM).

Ho, M. K. (2022, March). “Bridging Reinforcement Learning and Intuitive Pedagogy.” Invited Talk at RL4ED Workshop at AAAI 2022.

Ho, M. K. (2022, February). “Cognitive Science as a Source of Design Principles for Interactive Machine Learning.” New Jersey Institute of Technology.

Ho, M. K. (2022, February). “Cognitive Science as a Source of Design Principles for Interactive Machine Learning.” Steven’s Institute of Technology, New Jersey.

Ho, M. K. (2022, February). “Planning and Social Interaction as Meta-Computation.” Department of Psychology, University of Southern California.

Ho, M. K. (2021, November). “Control of mental representations in human planning.” Shenhav Lab (PI: Amitai Shenhav), Brown University.

Ho, M. K. (2021, November). “Control of mental representations in human planning.” Parallel Distributed Processing Seminar, Princeton University.

Ho, M. K. (2021, October). “Control of mental representations in human cognition.” Cognitive Development and Cognitive Science Brown Bag, Arizona State University.

Ho, M. K. (2021, September). “The role of context in human communication and its importance for transparent agency.” Invited talk at CINEMENTAS Workshop on Transparent Agency and Learning.

Ho, M. K. & Griffiths, T. L. (2021, July). “Rationally Representing Games.” Using Games to Understand Intelligence Workshop. Annual Meeting of the Cognitive Science Society.

Ho, M. K. (2021, March). “Resource-Rational Planning Representations” Moral Psychology Research Lab (PIs: Fiery Cushman and Joshua Greene), Harvard University.

Ho, M. K. (2020, October). “Models of Multi-agent Action and Inference.” Computational Approaches to Social Cognition Talk Series (Harvard).

Ho, M. K. (2020, October). “Communicative Decision-Making and Interactive Teaching.” Dartmouth Social Brain Brown Bag.

Ho, M. K. (2020, September). “Communication in Action: Planning and Interpreting Communicative Demonstrations.” Social Learning Lab (PI: Hyowon Gweon), Stanford University.

Ho, M. K. (2020, August). “Communicative Decision-Making and Interactive Teaching.” Computational Cognitive Neuroscience Lab (PI: Anne Collins), University of California Berkeley.

Ho, M. K. (2020, May). “Communication, Planning, and Meta-Reasoning.” Causality in Cognition Lab (PI: Tobias Gerstenberg), Stanford University.

Ho, M. K. (2020, March). “Meta-Reasoning about Partial Plans.” Parallel Distributed Processing Seminar, Princeton University.

Ho, M. K. (2019, October). “Interactive Communication and Miscommunication in Humans.” Center for Human-Compatible AI, University of California, Berkeley.

Ho, M. K. (2019, October). “Human-machine collaboration and information processing limitations” Ve-HiCal Project Annual Meeting, University of California, Berkeley.

Ho, M. K. (2019, September). “Communication, Coordination, and Computation in Human Interaction.” Department of Cognitive Science Seminar Series, Central European University.

Ho, M. K., Korman., J. & Griffiths T. L. (2019, July) “A computational account of unintentional speech acts.” Annual Meeting of the Society for Philosophy and Psychology, San Diego, California.

Ho, M. K. (2019, June). “Communication in Interactive Settings.” Department of Computer Science, University of North Carolina, Chapel Hill.

Ho, M. K. (2019, April). “Communicative Intentions in an Interactive World.” Social-Ecological and Environmental Lab (PI: Alexandra Paxton), University of Connecticut.

Ho, M. K. (2019, April). “Communicative Intentions in an Interactive World.” Yale Cognitive Development Laboratory (PI: Julian Jara-Ettinger).

Ho, M. K. (2019, April). “Communicative Intentions in Demonstrations and Rewards.” Neuroscience of Social Decision-Making Seminar, Princeton University.

Ho, M. K. (2019, March). “Communicative and Pedagogical Intentions in an Interactive World.” Concepts and Categories Seminar, New York University.

Ho, M. K. (2019, January). “Communicative Intent and Interactive Teaching.” Project 6 Meeting (PI: Jonathan Cohen), Princeton University.

Ho, M. K. (2018, October). “Communicative Intent and Interactive Teaching.” Cognition/Neuroscience Seminar Series, Stanford University.

Ho, M. K. (2017, May). “How People Intentionally Teach Agents in Interactive Settings.” DREAM Seminar, University of California, Berkeley.

Ho, M. K. (2017, April). “Teaching with Communicative Intent in Interactive Settings.” Computational Cognitive Science Group (PI: Josh Tenenbaum), MIT.

Ho, M. K. (2017, April). “Teaching by Demonstration: Showing vs. Doing.” Brown Robotics Group Meeting, Brown University.

Ho, M. K., Littman, M. L., Cushman, F., & Austerweil, J. L. (2016, August). “Not Quite Intuitive Behaviorists: Teachers use Rewards and Punishments Communicatively and not as Reinforcement.” Abstract presented at the 49th Annual Meeting of the Society for Mathematical Psychology, New Brunswick, NJ.

Ho, M. K. (2016, July). “Teaching with Evaluative Feedback (and by Demonstration), Communicatively.” Moral Psychology Research Lab, Harvard University.

Ho, M. K. (2016, May). “Teachers use rewards and punishments communicatively and not as reinforcement” presentation given to HAMLET (Human and Machine Learning: Experiments and Theory) seminar series at University of Wisconsin, Madison.

Austerweil, J. L., Brawner, S., Greenwald, A., Hilliard, E., Ho, M. K., Littman, M. L., MacGlashan, J., & Trimbach, C. (2016, March). “The Impact of Outcome Preferences in a Collection of Non-Zero-Sum Grid Games.” AAAI Spring Symposium 2016 on Challenges and Opportunities in Multiagent Learning for the Real World.

Ho, M. K., Fernbach, P.M & Sloman, S. A. (2015, May). “Opening minds by exposing the illusion of explanatory depth.” Talk given at the annual meeting of the Association for Psychological Science, New York, NY.

Ho, M. K. (2013, July). “Causal Self-Deception”. 5-minute data blitz presentation given at the Moral Psychology Research Group annual meeting.

## Workshops

MacIver, M., Daw, N. D., Espinosa, G., Hamrick, J. B., Ho, M. K., Redish, A. D., Stadie, B. C., Wang, J. X. (2022) What is the place of planning? Generative Adversarial Collaboration at the Cognitive Computational Neuroscience (CCN) conference. Organizer. Website: <https://gac.ccneuro.org/gacs-by-year/2022-gacs/2022-3>.

Abel, D., Harutyunyan, A., Ho, M. K. (2022) Reinforcement Learning as a Model of Agency: Perspectives, Limitations, and Possibilities. Workshop at the Multi-disciplinary Conference on Reinforcement Learning and Decision Making (RLDM). Organizer. Workshop website: <https://sites.google.com/view/rl-as-agency/>

Wu, C. M., Vélez N., Ho, M. K., & Goldstone, R. L. (2020) Cognition, Collectives, and Human Culture. Workshop at the 42nd Annual Conference of the Cognitive Science Society. Organizing committee member and presenter. Workshop website: <https://cognitioncollectivesandculture.github.io/>

Hamrick, J, Nematzadeh, A, Burns, K. Dupoux, E. Gopnik, A, & Tenenbaum, J. (2020) Bridging AI And Cognitive Science (BAICS). Workshop at the International Conference on Learning Representations. Program committee member. Website: <https://baicsworkshop.github.io>

Brys, T., Harutyunyan A., Mannion, P & Subramanian, K. (2017) Adaptive Learning Agents. Workshop at the International Conference on Autonomous Agents and Multiagent Systems. Program committee member. Website: <http://ala2017.it.nuigalway.ie/>

Mathewson, K., Subramanian, K., Ho, M. K., Loftin, R., Austerweil, J.L., Harutyunyan, A., Precup, D., El Asri, L., Gombolay, M., Zhu, X., Chernova, S., Isbell, C. L., Pilarski, P. M., Wong, W. K., Veloso, M., Shah, J.A., Taylor, M., Argall, B., & Littman, M. L. (2016) Future of Interactive Learning Machines. Workshop at the 30th Conference on Neural Information Processing Systems. Organizing Committee and Programming Committee member.

## Teaching

Computational Cognitive Science (PSYCH-UA.300-011) New York University, Department of Psychology	Spring 2026
Lab in Cognition and Perception (PSYCH-UA.46) New York University, Department of Psychology	Fall 2025
Computational Social Psychology (PSYCH-GA.3404) New York University, Department of Psychology	Spring 2025
Lab in Cognition and Perception (PSYCH-UA.46) New York University, Department of Psychology	Spring 2025
Readings in Reinforcement Learning and Sequential Decision-Making (CS-810-C) Stevens Institute of Technology, Department of Computer Science	Spring 2024
Computational Cognitive Science (CS-810-C) Stevens Institute of Technology, Department of Computer Science	Fall 2023
Capstone Project and Presentation (DS-GA 1006) New York University, Center for Data Science	Fall 2022

## Advising

Maya Malaviya (current Ph.D. student, NYU Department of Psychology)  
 Peiyao Hu (current Ph.D. student, NYU Department of Psychology)  
 Brian Pennisi (Master's student project, NYU Center for Data Science)  
 Ruiqi He (visiting Ph.D. student, Max Planck Institute for Intelligent Systems)  
 Daniel Ritter (undergrad thesis, Brown University)  
 Albert Lin (undergrad junior project, Princeton; now Ph.D. student in Electrical Engineering at USC)

## Ad hoc reviewer for

Current Directions in Psychological Science  
 Trends in Cognitive Science  
 PLOS Computational Biology  
 Cognitive Science  
 Cognition

Biology Letters

Computational Brain & Behavior

Annual Meeting of Cognitive Science Society

NeurIPS

Topics in Cognitive Science

Social Cognition

Journal of Experimental Psychology: General

Nature Human Behavior

Proceedings of the National Academy of Sciences (PNAS)