Lio Wong

liowong@stanford.edu
https://web.mit.edu/zyzzyva/www/academic.html

Research Interests: language understanding, reasoning and meta-reasoning, probabilistic and decision making models, cognitive science, program synthesis

Academic Positions

2027 *Assistant Professor,* Department of Brain and Psychological Sciences, Brown University
2025–2027 *Postdoctoral Human-Centered Artificial Intelligence Fellow,* Stanford University
2024 *Postdoctoral Research Associate,* Massachusetts Institute of Technology

Education

2018–2024 *PhD,* Brain and Cognitive Science, Massachusetts Institute of Technology
Thesis: From Words to Worlds: Bridging Language and Thought
Glushko Prize for Outstanding Dissertation from the Cognitive Science Society
Advisors: Joshua B. Tenenbaum and Jacob Andreas

2018 *MS,* Computer Science, *Stanford University*

2013–2017 *BS,* Computer Science, *Stanford University*
with honors. Academic advisor: Daniel Jurafsky. Thesis advisor: Sebastian Thrun.

Fellowships and Honors

2025 Robert J. Glushko Prize for Outstanding Doctoral Dissertation, Cognitive Science Society
2024 Stanford Human-Centered AI Postdoctoral Fellowship
2024 Sayan Gul Award for Best Undergraduate Student Paper, CogSci (Student advisee supervisor)
2023 William James Prize Honorable Mention, Society for Philosophy and Psychology
2022 Departmental Diversity, Equity, Inclusion and Justice Impact Award, MIT
2021 Walle Nauta Award for Excellence in Undergraduate Teaching, MIT
2018 Singleton Presidential Graduate Fellowship, MIT
2017 Siebel Scholar, Stanford
2017 Wegbreit Prize for Best Undergraduate Honors Thesis in Computer Science, Stanford
2017 Terman Engineering Award (Top 5% Engineering), Stanford
2016 Tau Beta Pi Engineering Honor Society
2014 Stanford President's Award for Academic Excellence

Conference and Journal Papers

*Denotes equal contribution; Underlining indicates supervised mentee during graduate work.
I publish work in venues in cognitive science / linguistics (eg, CogSci, CDPS, SPP, CAMP), AI / natural language processing (eg, NeurIPS, ICML), and programming languages (eg, POPL).

2026, Forthcoming — [1] People use fast and flat simulation to reason about new games.
K. Collins*, C. Zhang*, **L. Wong*,** M. Barba da Costa, G. Todd, A. Weller, S. Cheyette, T. Griffiths, J. Tenenbaum.
Nature.

2026 — [2] Meaningful long-term thought partnerships of minds and machines.
K. Collins*, **L. Wong*,** J. Tenenbaum, J. Fan.
Current Directions in Psychological Science (CDPS).

2026 — [3] Medical Model Synthesis Architectures: A Case Study
K. Collins, M. Berke, I. Sucholutsky, A. Ali, A. Weller, T. O'Donnell, T. Brooke-Wilson*, **L. Wong***, J. Tenenbaum.
ICML: AI4Good Workshop.

2026 — [4] MedRedFlag: Investigating how LLMs Redirect Misconceptions in Real-World Health Communication.
S. Sambara, Y. Pu, A. Ali, V. Mishra, **L. Wong**, M. Agrawal
Findings of the Association for Computational Linguistics (ACL).

2026 — [5] Evaluating language models' evaluations of games.
K. Collins, C. Zhang, G. Todd, L. Ying, M. Barba da Costa, R. Liu, P. Sharma, A. Weller, I. Kuperwajs, **L. Wong**, J. Tenenbaum*, T. Griffiths*
International Conference on Learning Representations (ICLR).

2026 — [6] Resource rational contracturalism should guide AI alignment.
S. Levine, M. Franklin, T. Zhi-Xuan, S. Yanik Guyot, **L. Wong**, D. Kilov, Y. Choi, J. Tenenbaum, N. Goodman, S. Lazar, I. Gabriel.
International Association for Safe and Ethical AI (IASEAI).

2026 — [7] Generation and evaluation in the human invention process through the lens of game design.
K. Collins, G. Todd, C. Zhang, A. Weller, J. Togelius, J. Chu, **L. Wong**, T. Griffiths, J. Tenenbaum
Machine Learning for Health Symposium.

2026 — [8] What patients really ask: exploring the effect of false assumptions in patient information seeking.
R. Xiong, F. Jia, **L. Wong**, M. Agrawal.
Proceedings of the Annual Meeting of the Cognitive Science Society.

2025 — [9] Modeling open-world cognition as on-demand synthesis of probabilistic models.
L. Wong*, K. Collins*, L. Ying, C. Zhang, A. Weller, T. Gerstenberg, T. O'Donnell, A. Lew, J. Andreas, J. Tenenbaum, T. Brooke-Wilson.
Proceedings of the Annual Meeting of the Cognitive Science Society, also presented at California Meeting on Psycholinguistics and Society for Philosophy and Psychology.

2025 [10] Language-informed synthesis of rational agent models for grounded theory-of-mind reasoning
on-the-fly.
L. Ying, R. Truong, K. Collins, C. Zhang, M. Wei. Tyler Brooke-Wilson, T. Zhi-Xuan, **L.Wong***, J. Tenenbaum*
Findings of the Association for Computational Linguistics (ACL).

2025 [11] Retrieval-augmented systems can be dangerous medical communicators.
L. Wong, A. Ali, R. Xiong, S. Shen, Y. Kim, M. Agrawal.
International Conference on Machine Learning (ICML).

2025 [12] Consequences of prior experience on visual problem solving.
S. Anderson, **L. Wong,** M. Bowers, J. Fan.
Proceedings of the Annual Meeting of the Cognitive Science Society.

2025 [13] Finding structure in logographic writing with library learning II:
Grapheme, sound, and meaning systematicity.
G. Jiang, M. Hofer, J. Mao, **L. Wong**, J. Tenenbaum, R. Levy.
Proceedings of the Annual Meeting of the Cognitive Science Society.

2025 [14] Generation and evaluation in the human invention process through the lens of game design.
K. Collins, G. Todd, C. Zhang, A. Weller, J. Togelius, J. Chu, **L. Wong**, T. Griffiths, J. Tenenbaum
Proceedings of the Annual Meeting of the Cognitive Science Society.

2025 [15] Understanding epistemic language with a Bayesian Theory-of-Mind.
L. Ying, T. Zhi-Xuan, **L. Wong**, V. Mansinghka, J. Tenenbaum.
Transactions of the Association for Computational Linguistics (TACL).

2024 [16] Building machines that learn and think with people.
K. Collins* , I. Sucholutsky*, U. Bhatt*, K. Chandra*, **L. Wong*** , M. Lee, C. Zhang, T. Zhi-Xuan, M. Ho, V. Mansinghka, A. Weller, J. Tenenbaum, T. Griffiths.
Nature Human Behavior.

2024 [17] Evaluating language models for mathematics through interactions.
K. Collins, A. Jiang, S. Frieder, **L. Wong**, M. Zilka, U. Bhatt, T. Lukasiewicz, Y. Wu, J. Tenenbaum, W. Hart, T. Gowers, W. Li, A. Weller, M. Jamnik.
Proceedings of the National Academy of Sciences (PNAS).

2024 [18] People use fast, goal-directed simulation to reason about novel games.
C. Zhang*, K. Collins*, **L. Wong*,** A. Weller, J. Tenenbaum.
Proceedings of the Annual Meeting of the Cognitive Science Society.

2024 [19] Finding structure in logographic writing with library learning.
G. Jiang, M. Hofer, J. Mao, **L. Wong**, J. Tenenbaum, R. Levy.
Proceedings of the Annual Meeting of the Cognitive Science Society.
(**Sayan Gul Award for Best Undergraduate Student Paper, Diversity and Inclusion Award**)

2024 [20] Learning grounded action abstractions from language.
L. Wong*, J. Mao*, P. Sharma*, Z. Siegel, J. Feng, N. Korneev, J. Tenenbaum, J. Andreas.
International Conference on Learning Representations (ICLR).

2024 [21] LILO: Learning interpretable libraries by compressing and documenting code.
G. Grand, **L. Wong,** M. Bowers, T. Olausson, M. Liu, J. Tenenbaum, J. Andreas.
International Conference on Learning Representations (ICLR).

2023 [22] From word models to world models: translating from natural language to the language of thought.
L. Wong*, G. Grand*, A. Lew, N. Goodman, V. Mansinghka, J. Andreas, J. Tenenbaum.
Society for Philosophy and Psychology, in prep for *Psychological Review.*
(**Glushko Dissertation Prize, William James Prize Honorable Mention)**

2023 [23] The neuro-symbolic inverse planning engine (NIPE): modeling probabilistic social inferences from linguistic inputs.
L. Ying, K. Collins, M. Wei, C. Zhang, Zhi-Xuan T., A. Weller, J. Tenenbaum, **L. Wong.**
ICML Workshop on Theory of Mind in Communicating Agents.

2023 [24] Grounded physical language understanding with probabilistic programs and simulated worlds.
C. Zhang, **L. Wong,** G. Grand, J. Tenenbaum.
Proceedings of the Annual Meeting of the Cognitive Science Society.

2023 [25] Evaluating statistical language models as pragmatic reasoners.
B. Lipkin, **L. Wong**, G. Grand, J. Tenenbaum.
Proceedings of the Annual Meeting of the Cognitive Science Society.

2023 [26] DreamCoder: growing generalizable, interpretable knowledge with wake–sleep Bayesian program learning.
K. Ellis, **L. Wong**, M. Nye, M. Sable-Meyer, L. Cary, L. Pozo, L. Hewitt, A. Solar-Lezama, J. Tenenbaum.
Philosophical Transactions of the Royal Society A.

2023 [27] Top-down synthesis for library learning.
M. Bowers, T. Olausson, **L. Wong**, G. Grand, J. Tenenbaum, K. Ellis, A. Solar-Lezama.
Proceedings of the ACM on Programming Languages (POPL).

2023 [28] Communicating natural programs to humans and machines.
S. Acquaviva, Y. Pu, M. Kryven, T. Sechopoulos, **L. Wong**, G. Ecanow, M. Nye, M. Tessler, J. Tenenbaum.
Advances in Neural Information Processing Systems (NeurIPS).

2022 [29] Identifying concept libraries from language about object structure.
L. Wong*, W. McCarthy*, G. Grand*, Y. Friedman, J. Tenenbaum, J. Andreas, R. Hawkins, J. Fan.
Proceedings of the Annual Meeting of the Cognitive Science Society.

2022 [30] Structured, flexible, and robust: benchmarking and improving large language models towards more human-like behavior in out-of-distribution reasoning task.
K. Collins*, **L. Wong***, J. Feng, M. Wei, J. Tenenbaum.
Proceedings of the Annual Meeting of the Cognitive Science Society.
(**Student Travel Grant Award)**

2021 [31] Leveraging language to learn program abstractions and search heuristics.
L. Wong, K. Ellis, J. Tenenbaum, J. Andreas.

International Conference on Machine Learning (ICML).

2021 [32] DreamCoder: bootstrapping inductive program synthesis with wake-sleep library learning.
K. Ellis, **L. Wong**, M. Nye, M. Sable-Meyer, L. Cary, L. Pozo, L. Hewitt, A. Solar-Lezama, J. Tenenbaum.
Proceedings of the ACM SIGPLAN International Conference on Programming Language Design and Implementation (*PLDI).*

2019 [33] Modeling expertise with neurally-guided Bayesian program induction.
L. Wong, K. Ellis, M. Sablé-Meyer, J. Tenenbaum.
Proceedings of the Annual Meeting of the Cognitive Science Society.

2019 [34] From mental representations to neural codes: a multi-level approach.
J. Gauthier*, J. Loula*, E. Pollock*, T. Wilson*, **L. Wong***,
Commentary on *Behavioral Brain Sciences (BBS).*

2018 [35] Transfer learning with AutoML.
L. Wong, N. Houlsby, Y. Lu, A. Gesmundo.
Advances in Neural Information Processing Systems (NeurIPS).

2015 [36] A rapid molecular approach for chromosomal phasing.
J. Regan, N. Kamitaki, T. Legler, S. Cooper, N. Klitgord, G. Karlin-Neumann, **L. Wong**, S. Hodges, R. Koehler, S. Tzonev, and S. McCarroll. *Public Library of Science* (*PLOS One).*

2012 [37] Cell-phone compatible wireless stethoscope.
L. Wong.
U.S. Patent Application No. 13/326,927.

Selected Invited Talks

2026 CogSci Workshop, Framing the Problem: Abstraction and Task Representations in Human Problem Solving
2026 Berkeley NLP Lab Meeting
2025 Computational Summer School on Modeling Social and Collective Behavior (COSMOS, October)
2025 CogSci Workshop, Putting It Together: Interactions Between Domains of Cognition
2025 Thompson Computational Cognitive Science Lab, UC Berkeley
2025 Stanford Psychology Friday Seminar
2024 Simons Center for the Social Brain Lunch Series
2024 Dartmouth Summer Workshop on Artificial Intelligence, Language, and Cognition

2023 *Integrating Language into Intelligent Architectures (with Alex Lew)*
Simons Institute Workshop on Large Language Models and Transformers

2023 *From Word Models to World Models (with Gabriel Grand)*
CogSci Workshop, Large Language Models Meet Cognitive Science

2023 CompLing Graduate Seminar, UT Austin

2023 Gershman Lab, Harvard University

2023 *From images to a lexicon: getting human-like meaning out of drawings*
CogSci Workshop, Images to Symbols

2022 *Bridging language and thought with a language of thought*
ConCats Seminar, NYU

2022 McDonnell Foundation Plenary Meeting on Concepts
2022 Embodied Intelligence Seminar, MIT

2022 *Discovering concept libraries from language*
CALLAX Seminar, UCSD

2022 RLDM Seminar, University of Tubingen
2021 ACL NL4Programs Workshop
2021 Topics in Natural Language Processing and Machine Learning Guest Lecture, Cornell University
2020 AAAI Conceptual Abstraction and Analogy in Natural and Artificial Intelligence

2020 *The Game Engineer's Challenge: Building Representations That Span Different Sentences in Different Worlds*
NeurIPS Workshop on Language Meets Games

Professional Service

Workshops and Symposiums Organized

2026 CogSci Workshop, *Unconventional Communication: making meaning across modalities.*
2025 ICML Workshop, *Assessing world models.*
2025 CogSci Workshop, *Meta-reasoning: Deciding which game to play, which problem to solve, and when to quit.*
2024 CogSci Workshop, *Compositionality in minds, brains and machines: a unifying goal that cuts across cognitive sciences.*
2024 CogSci Symposium, *Who is responsible for collective action?*
2023 ACL Workshop, *Natural language reasoning and structured explanations.*
2023 CogSci Workshop, *How does the mind discover useful abstractions?*

Service to the University

2024 Visiting Committee on Gender Identity in the Lab, MIT
2020 Institute Committee on Race and Diversity, MIT
2019 Brain and Cognitive Sciences Philosophy Circle, MIT
2015 Engineers for a Sustainable World, Stanford

Service to the Field

Reviewer for conferences including: NeurIPS, ICML, ICLR, CogSci.

Program Committee: NeurIPS AI Meets Moral Philosophy and Moral Psychology, NeurIPS Visually Grounded Interaction and Language

Other Employment

2018 Research Assistant, Computational Vision and Geometry Lab, Stanford
2017 Research Intern, Google Research
2017 Research Assistant, Thrun Lab, Stanford
2016 Software Engineering Intern, Flatiron Health
2015 Software Engineering Intern, Palantir Philanthropy
2015 Software Engineering Intern, Google Classroom
2014 Software Engineering Intern, Google[x]

Fiction and Sound Experience

2023- *Founding Editor,* Signal Hill Audio Magazine
2025 *The Unknown Unknowns*, Signal Hill Audio Magazine (Tribeca Festival, 2026)
2024 *Review, All of Us Strangers,* The Yale Review
2022 *with all the force of a great typhoon,* Split Lip Magazine
2021 *Posing,* Gasher Magazine
2020 *Clear,* Bayou Magazine
2020 *Polaroid,* Wildness Journal
2019 *Feeding Jobs,* The Cimarron Review
2019 *In Residence,* Shenandoah
2019 *Tomorrow Never Knows,* Chicago Quarterly Review
2017 Audio Production Intern, WNYC Radio
2016 Senior Audio Producer, Stanford Storytelling Project