

Paul Pu Liang

✉ ppliang@mit.edu • 🌐 <https://ppliang279.github.io/> • 🌐 [ppliang279](#) • 🎓 [Google Scholar](#)

Personal Information

Citizenship: US Citizen
Office: Wiesner Building E15-391, 20 Ames Street, Cambridge, MA 02139, USA

Short Biography

Paul Liang is an Assistant Professor at the MIT Media Lab and MIT EECS. His research advances the foundations of multisensory artificial intelligence to enhance the human experience. He is a recipient of the [Siebel Scholars Award](#), Waibel Presidential Fellowship, [Facebook PhD Fellowship](#), [Center for ML and Health Fellowship](#), [Rising Stars in Data Science](#), and 3 best paper awards. Outside of research, he received the [Alan J. Perlis Graduate Student Teaching Award](#) for developing new [courses](#) on multimodal machine learning.

Research Interests

- Foundations of multisensory AI: The science and engineering of AI systems that can learn, reason, and interact with the world through integrating diverse sensory channels.
- Enhancing human experiences: Designing interactive AI technologies to augment human capabilities and improve overall well-being.
- Real-world human-AI interaction: Quantifying and mitigating real-world societal concerns for responsible deployment.

Academic Career

- Massachusetts Institute of Technology 2024–Present
Sony Corporation Career Development Assistant Professor, MIT Media Lab & MIT EECS
Director, [Multisensory Intelligence](#) Research Group
Affiliate Professor, MIT Institute for Data Systems and Society & Broad Institute of MIT and Harvard
- UC Berkeley Simons Institute for the Theory of Computing Summer 2024
Visiting Researcher in the program on AI, Psychology, and Neuroscience
- Carnegie Mellon University, Ph.D. in Machine Learning 2018–2024
Committee: Louis-Philippe Morency (co-advisor), Ruslan Salakhutdinov (co-advisor), Manuel Blum, Lenore Blum, Trevor Darrell
Thesis: [Foundations of Multisensory Artificial Intelligence](#)
- Carnegie Mellon University, M.S. in Machine Learning 2017–2018
Thesis: [Computational Modeling of Human Multimodal Language](#)
- Carnegie Mellon University, B.S. in Computer Science & Neural Computation (with university honors) 2014–2017
Advisors: Roni Rosenfeld, Ryan Tibshirani, Tai Sing Lee

Professional Experiences

- Google DeepMind, Research Intern with Dani Yogatama, Aida Nematzadeh, Lisa Hendricks Summer 2021
- Facebook AI Research, Research Intern with Brandon Amos, Tim Rocktäschel, Ed Grefenstette Summer 2020
- Nvidia AI, Research Intern with Yuke Zhu, Anima Anandkumar, Sanja Fidler Spring 2020
- Google Research, Research Intern with Manzil Zaheer, Yuan Wang, Amr Ahmed Summer 2019
- Riken AI Project, Visiting Researcher with Makoto Yamada, Qibin Zhao, Masashi Sugiyama Winter 2018

Honors and Awards

- TMLR [Finalist for Outstanding Certification Paper](#) 2025
- [Forbes 30 under 30 Science](#) 2025
- MIT Sony Career Development Chair Professorship 2024
- CMU School of Computer Science Distinguished Dissertation Award 2024
- [Siebel Scholars Award](#) (\$35,000 award, 80 awardees worldwide) 2023
- Waibel Presidential Fellowship (\$100,000 funding for 1 year, 1 awardee from CMU) 2023
- [University of Chicago & UC San Diego Rising Stars in Data Science](#) 2023
- [Alan J. Perlis Graduate Student Teaching Award](#) 2023
- NeurIPS 2022 Workshop on Human in the Loop Learning [Best Paper Nomination](#) 2022
- Young Researcher Invitation to the [Heidelberg Laureate Forum](#) 2022
- [Facebook PhD Fellowship](#) (\$200,000 funding over 2 years, 26 awardees out of 2100 applicants worldwide) 2021
- [Center for Machine Learning and Health Fellowship](#) (\$100,000 funding for 1 year, 10 awardees from CMU) 2021
- MIT Fintech Innovation Prize Grant and Sandbox Innovation Program Funding 2021
- CVPR 2020 [Argoverse Competition Honorable Mention](#) 2020
- NeurIPS 2019 Workshop on Federated Learning [Distinguished Student Paper](#) 2019
- CMU Machine Learning Department Teaching Assistant Award 2019
- NSF Graduate Research Fellowship Honorable Mention 2018
- ICMI 2017 [Best Paper Honorable Mention](#) 2017

- CMU School of Computer Science University Honors
- CMU School of Computer Science Dean's List

2017

2015, 2016, 2017

Publications

(*denotes joint first-authors, [Google Scholar](#) citations: 17,000, h-index: 44)

86. [QoQ-Med: Building Multimodal Clinical Foundation Models with Domain-Aware GRPO Training](#)
Wei Dai, Peilin Chen, Chanakya Ekbote, Paul Pu Liang
NeurIPS 2025 ([oral](#)) [[code](#)]
85. [What One Cannot, Two Can: Two-Layer Transformers Provably Represent Induction Heads on Any-Order Markov Chains](#)
Chanakya Ekbote, Ashok Vardhan Makkuva, Marco Bondaschi, Nived Rajaraman, Michael Gastpar, Jason D. Lee, Paul Pu Liang
NeurIPS 2025 ([spotlight](#))
84. [Multimodal Fusion Balancing Through Game-Theoretic Regularization](#)
Konstantinos Kontras, Thomas Strypsteen, Christos Chatzichristos, Paul Pu Liang, Matthew Blaschko, Maarten De Vos
NeurIPS 2025 ([spotlight](#)) [[code](#)]
83. [Partial Information Decomposition via Normalizing Flows in Latent Gaussian Distributions](#)
Wenyuan Zhao, Adithya Balachandran, Chao Tian, Paul Pu Liang
NeurIPS 2025
82. [REGen: Multimodal Retrieval-Embedded Generation for Long-to-Short Video Editing](#)
Wei Han Xu, Yimeng Ma, Jingyue Huang, Yang Li, Wenye Ma, Taylor Berg-Kirkpatrick, Julian McAuley, Paul Liang, Hao-Wen Dong
NeurIPS 2025 [[code](#)]
81. [MimeQA: Towards Socially-Intelligent Nonverbal Foundation Models](#)
Hengzhi Li, Megan Tjandrasuwita, Yi R. Fung, Armando Solar-Lezama, Paul Pu Liang
NeurIPS 2025 Datasets and Benchmarks Track [[code](#)]
80. [Position: Simulating Society Requires Simulating Thought](#)
Chance Li, Jiayi Wu, Z. Mo, A. Qu, Y. Tang, K. Zhao, Y. Gan, J. Fan, J. Yu, Jinhua Zhao, Paul Liang, Luis Alonso, Kent Larson
NeurIPS 2025 Position Paper Track
79. [Social Genome: Grounded Social Reasoning Abilities of Multimodal Models](#)
Leena Mathur, Marian Qian, Paul Pu Liang, Louis-Philippe Morency
EMNLP 2025
78. [Deriving Strategic Market Insights with Large Language Models: A Benchmark for Forward Counterfactual Generation](#)
Keane Ong, Rui Mao, Deeksha Varshney, Paul Pu Liang, Erik Cambria, Gianmarco Mengaldo
EMNLP 2025
77. [Understanding the Emergence of Multimodal Representation Alignment](#)
Megan Tjandrasuwita, Chanakya Ekbote, Liu Ziyin, Paul Pu Liang
ICML 2025 [[code](#)]
76. [CLIMB: Data Foundations for Large-Scale Multimodal Clinical Foundation Models](#)
David Dai, Peilin Chen, Malinda Lu, Daniel Li, Haowen Wei, Hejie Cui, Paul Pu Liang
ICML 2025 [[code](#)]
75. [POET: Supporting Prompting Creativity and Personalization with Automated Expansion of Text-to-Image Generation](#)
Evans Han, Alice Zhang, Haiyi Zhu, Hong Shen, Paul Pu Liang, Jane Hsieh
UIST 2025 [[code](#)]
74. [Can A Society of Generative Agents Simulate Human Behavior and Inform Public Health Policy? A Case Study on Vaccine Hesitancy](#)
Abe Hou, Hongru Du, Yichen Wang, Jingyu Zhang, Paul Pu Liang, Daniel Khashabi, Lauren Gardner, Tianxing He
COLM 2025 [[code](#)]
73. [VLM²-Bench: A Closer Look at How Well VLMs Implicitly Link Explicit Matching Visual Cues](#)
Jianshu Zhang, Dongyu Yao, Renjie Pi, Paul Pu Liang, Yi R. Fung
ACL 2025 [[code](#)]
72. [TAMP: Token-Adaptive Layerwise Pruning in Multimodal Large Language Models](#)
Jaewoo Lee, Keyang Xuan, Chanakya Ekbote, Sandeep Polisetty, Yi R. Fung, Paul Pu Liang
ACL 2025 Findings [[code](#)]
71. [Course on Multimodal AI for Human Sensing and Interaction](#)
Paul Pu Liang, Karan Ahuja, Yiyue Luo
CHI 2025 Course [[website](#)]
70. [Interactive Sketchpad: An Interactive Multimodal System for Collaborative, Visual Problem-Solving](#)
Jimin Lee, Steven Chen, Paul Pu Liang
CHI 2025 Late Breaking Work [[code](#)]
69. [Fits like a Flex-Glove: Automatic Design of Personalized FPCB-Based Tactile Sensing Gloves](#)
Devin Murphy, Yichen Li, Crystal Owens, Layla Stanton, Paul Pu Liang, Yiyue Luo, Antonio Torralba, Wojciech Matusik
CHI 2025 Late Breaking Work [[code](#)]

68. [Progressive Compositionality In Text-to-Image Generative Models](#)
Evans Han, Linghao Jin, Xiaofeng Liu, Paul Pu Liang
ICLR 2025 ([spotlight](#)) [[code](#)]
67. [OS-ATLAS: A Foundation Action Model for Generalist GUI Agents](#)
Zhiyong Wu, Zhenyu Wu, Fangzhi Xu, Y. Wang, Q. Sun, C. Jia, K. Cheng, Z. Ding, L. Chen, Paul Pu Liang, Yu Qiao
ICLR 2025 ([spotlight](#)) [[code](#)]
66. [VideoWebArena: Evaluating Long Context Multimodal Agents with Video Understanding Web Tasks](#)
Lawrence Jang, Yinheng Li, Charles Ding, Justin Lin, Paul Pu Liang, Dan Zhao, Rogerio Bonatti, Kazuhito Koishida
ICLR 2025 [[code](#)]
65. [TeaserGen: Generating Teasers for Long Documentaries](#)
Weihan Xu, Paul Pu Liang, Haven Kim, Julian McAuley, Taylor Berg-Kirkpatrick, Hao-Wen Dong
ICLR 2025 [[code](#)]
64. [Beyond the Imitation Game: Quantifying and Extrapolating the Capabilities of Language Models](#)
442 authors including Paul Pu Liang
ICLR 2025, TMLR 2023 ([finalist for outstanding certification](#)) [[code](#)]
63. [OpenFace 3.0: A Lightweight Toolkit for Comprehensive Facial Behavior Analysis](#)
Jiewen Hu, Leena Mathur, Paul Pu Liang, Louis-Philippe Morency
FG 2025 ([oral](#)) [[code](#)]
62. [Comparative Knowledge Distillation](#)
Alex Xu, Alex Wilf, Paul Pu Liang, Alexander Obolenskiy, Daniel Fried, Louis-Philippe Morency
WACV 2025 [[code](#)]
61. [HEMM: Holistic Evaluation of Multimodal Foundation Models](#)
Paul Pu Liang, Akshay Goindani, Talha Chafekar, Leena Mathur, Haofei Yu, Ruslan Salakhutdinov, Louis-Philippe Morency
NeurIPS 2024 Datasets and Benchmarks Track [[code](#)]
60. [MMoE: Enhancing Multimodal Language Models with Mixtures of Multimodal Interaction Experts](#)
Haofei Yu, Jason Qi, Lawrence Jang, Ruslan Salakhutdinov, Louis-Philippe Morency, Paul Pu Liang
EMNLP 2024 [[code](#)]
59. [Advancing Social Intelligence in AI Agents: Technical Challenges and Open Questions](#)
Leena Mathur, Paul Pu Liang, Louis-Philippe Morency
EMNLP 2024 [[code](#)]
58. [Semantically Corrected Amharic Automatic Speech Recognition](#)
Samuel Adnew, Paul Pu Liang
Deep Learning Indaba 2024 [[code](#)]
57. [Foundations of Multisensory Artificial Intelligence](#)
Paul Pu Liang
PhD Thesis 2024 ([CMU School of Computer Science distinguished dissertation award](#))
56. [Foundations & Trends in Multimodal Machine Learning: Principles, Challenges, and Open Questions](#)
Paul Pu Liang, Amir Zadeh, Louis-Philippe Morency
ACM Computing Surveys 2024, Tutorials at ICML & ICMJ 2023, CVPR & NAACL 2022 [[website](#)] [[videos 2023](#)] [[videos 2022](#)]
55. [Think Twice: Perspective-Taking Improves Large Language Models' Theory-of-Mind Capabilities](#)
Alex Wilf, Sihyun Shawn Lee, Paul Pu Liang, Louis-Philippe Morency
ACL 2024 [[code](#)]
54. [Modeling Dense Multimodal Interactions Between Biological Pathways and Histology for Survival Prediction](#)
Guillaume Jaume, Anurag Vaidya, Richard Chen, Drew Williamson, Paul Pu Liang, Faisal Mahmood
CVPR 2024 [[code](#)]
53. [Multimodal Learning Without Labeled Multimodal Data: Guarantees and Applications](#)
Paul Pu Liang, Chun Kai Ling, Yun Cheng, Alex Obolenskiy, Yudong Liu, Rohan Pandey, Alex Wilf, L.P. Morency, R. Salakhutdinov
ICLR 2024 [[code](#)]
52. [Quantifying & Modeling Multimodal Interactions: An Information Decomposition Framework](#)
Paul Pu Liang, Y. Cheng, X. Fan, C.K. Ling, S. Nie, R. Chen, Z. Deng, N. Allen, R. Auerbach, F. Mahmood, R. Salakhutdinov, L.P. Morency
NeurIPS 2023 [[code](#)]
51. [Factorized Contrastive Learning: Going Beyond Multi-view Redundancy](#)
Paul Pu Liang*, Zihao Deng*, Martin Ma*, James Zou, Louis-Philippe Morency, Ruslan Salakhutdinov
NeurIPS 2023 [[code](#)]
50. [Localized Symbolic Knowledge Distillation for Visual Commonsense Models](#)
Jae Sung Park, Jack Hessel, Khyathi Chandu, Paul Pu Liang, Ximing Lu, Qiuyuan Huang, Peter West, Jianfeng Gao, Ali Farhadi, Yejin Choi
NeurIPS 2023 [[code](#)]
49. [Read and Reap the Rewards: Learning to Play Atari with the Help of Instruction Manuals](#)
Yue Wu, Yewen Fan, Paul Pu Liang, Amos Azaria, Yuanzhi Li, Tom Mitchell
NeurIPS 2023, ICLR 2023 Workshop on Reincarnating RL ([oral](#)) [[code](#)]

48. [Difference-Masking: Choosing What to Mask in Continued Pretraining](#)
Alex Wilf, Syeda Akter, Leena Mathur, Paul Pu Liang, Sheryl Mathew, Mengrou Shou, Eric Nyberg, Louis-Philippe Morency
EMNLP 2023 Findings [\[code\]](#)
47. [Multimodal Fusion Interactions: A Study of Human and Automatic Quantification](#)
Paul Pu Liang, Yun Cheng, Ruslan Salakhutdinov, Louis-Philippe Morency
ICMI 2023 [\[code\]](#)
46. [HIINT: Historical, Intra- and Inter- personal Dynamics Modeling with Cross-person Memory Transformer](#)
Yubin Kim, Dong Won Lee, Paul Pu Liang, Sharifa Alghowinem, Cynthia Breazeal, Hae Won Park
ICMI 2023
45. [Lecture Presentations Multimodal Dataset: Towards Understanding Multimodality in Educational Videos](#)
Dong Won Lee, Chaitanya Ahuja, Paul Pu Liang, Sanika Natu, Louis-Philippe Morency
ICCV 2023 [\[code\]](#)
44. [Cross-modal Attention Congruence Regularization for Vision-Language Relation Alignment](#)
Rohan Pandey, Rulin Shao, Paul Pu Liang, Ruslan Salakhutdinov, Louis-Philippe Morency
ACL 2023 [\[code\]](#)
43. [Language Models Get a Gender Makeover: Mitigating Gender Bias with Few-Shot Data Interventions](#)
Himanshu Thakur, Atishay Jain, Praneetha Vaddamanu, Paul Pu Liang, Louis-Philippe Morency
ACL 2023 [\[code\]](#)
42. [Nano: Nested Human-in-the-Loop Reward Learning for Few-shot Language Model Control](#)
Xiang Fan, Yiwei Lyu, Paul Pu Liang, Ruslan Salakhutdinov, Louis-Philippe Morency
ACL Findings 2023, NeurIPS 2022 Workshop on Human in the Loop Learning ([oral](#), [best paper nomination](#)) [\[code\]](#)
41. [MultiViz: Towards Visualizing and Understanding Multimodal Models](#)
Paul Pu Liang, Yiwei Lyu, Gunjan Chhablani, Nihal Jain, Zihao Deng, Xingbo Wang, Louis-Philippe Morency, Ruslan Salakhutdinov
ICLR 2023, CHI 2023 Late Breaking Work [\[code\]](#)
40. [FindAdaptNet: Find and Insert Adapters by Learned Layer Importance](#)
Junwei Huang, Karthik Ganesan, Soumi Maiti, Young Min Kim, Xuankai Chang, Paul Liang, Shinji Watanabe
ICASSP 2023
39. [Face-to-Face Contrastive Learning for Social Intelligence Question-Answering](#)
Alex Wilf, Martin Ma, Paul Pu Liang, Amir Zadeh, Louis-Philippe Morency
FG 2023
38. [High-Modality Multimodal Transformer: Quantifying Modality & Interaction Heterogeneity for High-Modality Representation Learning](#)
Paul Pu Liang, Yiwei Lyu, Xiang Fan, Jeffrey Tsaw, Yudong Liu, Dani Yogatama, Louis-Philippe Morency, Ruslan Salakhutdinov
TMLR 2022 [\[code\]](#)
37. [MultiZoo & MultiBench: A Standardized Toolkit for Multimodal Deep Learning](#)
Paul Pu Liang, Yiwei Lyu, Xiang Fan, Arav Agarwal, Yun Cheng, Louis-Philippe Morency, Ruslan Salakhutdinov
JMLR Open Source Software 2022 [\[website\]](#) [\[code\]](#)
36. [Brainish: Formalizing A Multimodal Language for Intelligence and Consciousness](#)
Paul Pu Liang
Association for the Scientific Study of Consciousness 2022, Models of Consciousness 2022 ([oral](#))
35. [Uncertainty Quantification with Pre-trained Language Models: A Large-scale Empirical Analysis](#)
Yuxin Xiao, Paul Pu Liang, Umang Bhatt, Willie Neiswanger, Ruslan Salakhutdinov, Louis-Philippe Morency
EMNLP Findings 2022 [\[code\]](#)
34. [GEMv2: Multilingual NLG Benchmarking in a Single Line of Code](#)
77 authors including Paul Pu Liang
EMNLP Demo Track 2022 [\[code\]](#)
33. [PACS: Physical Audiovisual Commonsense Reasoning](#)
Samuel Yu, Peter Wu, Paul Pu Liang, Ruslan Salakhutdinov, Louis-Philippe Morency
ECCV 2022 [\[code\]](#)
32. [DIME: Fine-grained Interpretations of Multimodal Models via Disentangled Local Explanations](#)
Yiwei Lyu, Paul Pu Liang, Zihao Deng, Ruslan Salakhutdinov, Louis-Philippe Morency
AIES 2022 [\[code\]](#)
31. [Rethinking Architecture Design for Tackling Data Heterogeneity in Federated Learning](#)
Liangqiong Qu*, Yuyin Zhou*, Paul Pu Liang*, Yingda Xia, Feifei Wang, Li Fei-Fei, Ehsan Adeli, Daniel Rubin
CVPR 2022 [\[code\]](#)
30. [MultiBench: Multiscale Benchmarks for Multimodal Representation Learning](#)
Paul Pu Liang, Yiwei Lyu, Xiang Fan, Z. Wu, Y. Cheng, J. Wu, L. Chen, P. Wu, M. Lee, Y. Zhu, R. Salakhutdinov, L.P. Morency
NeurIPS 2021 Datasets and Benchmarks Track [\[website\]](#) [\[code\]](#)
29. [Understanding the Tradeoffs in Client-side Privacy for Speech Recognition](#)
Peter Wu, Paul Pu Liang, Jiatong Shi, Ruslan Salakhutdinov, Shinji Watanabe, Louis-Philippe Morency
Asia Pacific Signal and Information Processing Association Annual Summit and Conference 2021 [\[code\]](#)

28. [Towards Understanding and Mitigating Social Biases in Language Models](#)
Paul Pu Liang, Chiyu Wu, Louis-Philippe Morency, Ruslan Salakhutdinov
ICML 2021 [\[code\]](#)
27. [Learning Language and Multimodal Privacy Preserving Markers of Mood from Mobile Data](#)
Paul Pu Liang*, Terrance Liu*, A. Cai, M. Muszynski, R. Ishii, N. Allen, R. Auerbach, D. Brent, R. Salakhutdinov, L.P. Morency
ACL 2021 (oral)
26. [Cross-Modal Generalization: Learning in Low Resource Modalities via Meta-Alignment](#)
Paul Pu Liang*, Peter Wu*, Liu Ziyin, Louis-Philippe Morency, Ruslan Salakhutdinov
ACM Multimedia 2021 (oral) [\[code\]](#)
25. [StylePTB: A Compositional Benchmark for Fine-grained Controllable Text Style Transfer](#)
Yiwei Lyu*, Paul Pu Liang*, Hai Pham*, Eduard Hovy, Barnabás Póczos, Ruslan Salakhutdinov, Louis-Philippe Morency
NAACL 2021 [\[code\]](#)
24. [Anchor & Transform: Learning Sparse Embeddings for Large Vocabularies](#)
Paul Pu Liang, Manzil Zaheer, Yuan Wang, Amr Ahmed
ICLR 2021 [\[code\]](#)
23. [MOSEAS: A Multimodal Language Dataset for Spanish, Portuguese, German and French](#)
Amir Zadeh, Yansheng Cao, Simon Hessner, Paul Pu Liang, Soujanya Poria, Louis-Philippe Morency
EMNLP 2020
22. [Diverse and Admissible Trajectory Prediction through Multimodal Context Understanding](#)
Seong Hyeon Park, Gyubok Lee, Manoj Bhat, Jimin Seo, Minseok Kang, Jon Francis, Ashwin Jadhav, Paul Liang, L.-P. Morency
ECCV 2020, CVPR 2020 Argoverse competition (honorable mention award) [\[code\]](#)
21. [Towards Debiasing Sentence Representations](#)
Paul Pu Liang, Irene Li, Emily Zheng, Yao Chong Lim, Ruslan Salakhutdinov, Louis-Philippe Morency
ACL 2020 [\[code\]](#)
20. [On Emergent Communication in Competitive Multi-Agent Teams](#)
Paul Pu Liang, Jeffrey Chen, Ruslan Salakhutdinov, Louis-Philippe Morency, Satwik Kottur
AAMAS 2020 (oral) [\[code\]](#)
19. [Empirical and Theoretical Studies of Multimodal Co-learning](#)
Amir Zadeh, Paul Pu Liang, Louis-Philippe Morency
Elsevier Information Fusion 2020
18. [Think Locally, Act Globally: Federated Learning with Local and Global Representations](#)
Paul Pu Liang*, Terrance Liu*, Liu Ziyin, Nick Allen, Randy Auerbach, David Brent, Ruslan Salakhutdinov, Louis-Philippe Morency
NeurIPS 2019 Workshop on Federated Learning (oral, distinguished student paper award) [\[code\]](#)
17. [Deep Gamblers: Learning to Abstain with Portfolio Theory](#)
Liu Ziyin, Zhikang Wang, Paul Pu Liang, Ruslan Salakhutdinov, Louis-Philippe Morency, Masahito Ueda
NeurIPS 2019 [\[code\]](#)
16. [Learning Representations from Imperfect Time Series Data via Tensor Rank Regularization](#)
Paul Pu Liang*, Zhun Liu*, Yao-Hung Hubert Tsai, Qibin Zhao, Ruslan Salakhutdinov, Louis-Philippe Morency
ACL 2019
15. [Multimodal Transformer for Unaligned Multimodal Language Sequences](#)
Yao-Hung Hubert Tsai, Shaojie Bai, Paul Pu Liang, Zico Kolter, Louis-Philippe Morency, Ruslan Salakhutdinov
ACL 2019 [\[code\]](#)
14. [Social-IQ: A Question Answering Benchmark for Artificial Social Intelligence](#)
Amir Zadeh, Michael Chan, Paul Pu Liang, Edmund Tong, Louis-Philippe Morency
CVPR 2019 (oral) [\[code\]](#)
13. [Strong and Simple Baselines for Multimodal Utterance Embeddings](#)
Paul Pu Liang*, Yao Chong Lim*, Yao-Hung Hubert Tsai, Ruslan Salakhutdinov, Louis-Philippe Morency
NAACL 2019 (oral) [\[code\]](#)
12. [Learning Factorized Multimodal Representations](#)
Yao-Hung Hubert Tsai*, Paul Pu Liang*, Amir Zadeh, Louis-Philippe Morency, Ruslan Salakhutdinov
ICLR 2019 [\[code\]](#)
11. [Found in Translation: Learning Robust Joint Representations by Cyclic Translations Between Modalities](#)
Hai Pham*, Paul Pu Liang*, Thomas Manzini, Louis-Philippe Morency, Barnabás Póczos
AAAI 2019 [\[code\]](#)
10. [Words can Shift: Dynamically Adjusting Word Representations Using Nonverbal Behaviors](#)
Yansen Wang, Ying Shen, Zhun Liu, Paul Pu Liang, Amir Zadeh, Louis-Philippe Morency
AAAI 2019 [\[code\]](#)
9. [Computational Modeling of Human Multimodal Language: The MOSEI Dataset and Interpretable Dynamic Fusion](#)
Paul Pu Liang, Ruslan Salakhutdinov, Louis-Philippe Morency
Master's Thesis, CMU Machine Learning Data Analysis Project 2018 (first runner-up award) [\[code\]](#)

8. [Multimodal Language Analysis with Recurrent Multistage Fusion](#)
Paul Pu Liang, Ziyin Liu, Amir Zadeh, Louis-Philippe Morency
EMNLP 2018 ([oral](#))
7. [Multimodal Local-Global Ranking Fusion for Emotion Recognition](#)
Paul Pu Liang, Amir Zadeh, Louis-Philippe Morency
ICMI 2018
6. [Multimodal Language Analysis in the Wild: CMU-MOSEI Dataset and Interpretable Dynamic Fusion Graph](#)
Amir Zadeh, Paul Pu Liang, Jon Vanbriesen, Soujanya Poria, Edmund Tong, Erik Cambria, Minghai Chen, Louis-Philippe Morency
ACL 2018 ([oral](#)) [[code](#)]
5. [Efficient Low-rank Multimodal Fusion with Modality-Specific Factors](#)
Zhun Liu, Ying Shen, Varun Lakshminarasimhan, Paul Pu Liang, Amir Zadeh, Louis-Philippe Morency
ACL 2018 ([oral](#)) [[code](#)]
4. [An Empirical Evaluation of Sketched SVD and its Application to Leverage Score Ordering](#)
Hui Han Chin, Paul Pu Liang
ACML 2018
3. [Multi-attention Recurrent Network for Human Communication Comprehension](#)
Amir Zadeh, Paul Pu Liang, Soujanya Poria, Prateek Vij, Erik Cambria, Louis-Philippe Morency
AAAI 2018 ([oral](#)) [[code](#)]
2. [Memory Fusion Network for Multi-view Sequential Learning](#)
Amir Zadeh, Paul Pu Liang, Navonil Mazumder, Soujanya Poria, Erik Cambria, Louis-Philippe Morency
AAAI 2018 ([oral](#)) [[code](#)]
1. [Multimodal Sentiment Analysis with Word-level Fusion and Reinforcement Learning](#)
Minghai Chen*, Sen Wang*, Paul Pu Liang*, Tadas Baltrušaitis, Amir Zadeh, Louis-Philippe Morency
ICMI 2017 ([oral](#), [best paper honorable mention](#)) [[code](#)]

Teaching

- [MIT Affective Computing and Multimodal Interaction](#) with Rosalind Picard Fall 2025
- [MIT MAS.S60 How to AI \(Almost\) Anything](#) Spring 2025
- [MIT 6.390 Intro to Machine Learning](#) with Shen Shen, Vince Monardo, Priya Donti, Manolis Kellis Spring 2025
- MIT Open Learning Course on Multimodal AI with Dimitris Bertsimas Fall 2024
- MIT Open Learning Course on AI for Human Communication Fall 2024
- [CMU 11-877 Advanced Topics in Multimodal Machine Learning](#), with Daniel Fried Spring 2024
- [CMU 11-777 Multimodal Machine Learning](#) with Louis-Philippe Morency Fall 2023
- [ICML, ICMI, CVPR, NAACL Tutorials on Multimodal ML](#) with Louis-Philippe Morency 2022-2023
- [African Masters of Machine Intelligence course on Multimodal AI](#) Summer 2023
- [CMU 11-866 Artificial Social Intelligence](#) with Louis-Philippe Morency Spring 2023
- [CMU 11-877 Advanced Topics in Multimodal ML](#) with Louis-Philippe Morency Spring 2023
- [CMU 11-777 Multimodal Machine Learning](#) with Louis-Philippe Morency Fall 2022
- [CMU 11-877 Advanced Topics in Multimodal ML](#) with Louis-Philippe Morency, Amir Zadeh Spring 2022

Invited Talks

- **New Advances in Multimodal Reasoning**
M3S Singapore, July 2025
UC Berkeley NLP Group, July 2025
Lawrence Berkeley National Laboratory Deep Learning for Science School, July 2025
Nanyang Technological University, April 2025
Singapore University of Technology and Design, April 2025
University of Cambridge Language Technology Lab Seminar, Mar 2025
MIT Lincoln Labs, Mar 2025
- **Multisensory AI for Health and Wellbeing**
Machine Learning for Health Symposium, Dec 2025
CVPR 2025 Workshop on Multimodality and Medicine, June 2025
Harvard Vision AI Seminar, Mar 2025
Mayo Clinic, Dec 2024
MIT Sloan Global Programs, Oct 2024
Harvard Medical School, Feb 2024
- **Towards High-modality Multimodal Learning**
Chilean Workshop on Multimodal Machine Learning, Dec 2024
Lambda Labs, Dec 2024
MIT Embodied Intelligence Seminar, Oct 2024
University of Glasgow, Sept 2024
MIT NLP Seminar, Sept 2024

ACL 2024 Workshop on Advances in Language and Vision Research, Aug 2024
National University of Singapore, Aug 2024

- **Foundations of Multisensory Artificial Intelligence**

MIT Institute for Data Systems and Society, Oct 2024
MIT Media Lab AI Conference, Sept 2024
MIT EECS, Sept 2024
Valence Labs, July 2024
UC Berkeley Simons Institute for the Theory of Computing, June 2024
University of Maryland CS, April 2024
MIT Media Lab and Schwarzman College of Computing, April 2024
University of Southern California CS, March 2024
Johns Hopkins University CS, March 2024
UIUC School of Information Sciences, March 2024
UT Dallas CS and ECE, March 2024
University of Washington ECE, Feb 2024
UT Austin CS, Feb 2024
Boston University Computing & Data Sciences, Feb 2024
University of Virginia CS and Data Science, Feb 2024
Stanford University Management Science and Engineering, Feb 2024
Georgia Tech CS, Feb 2024
Duke University CS and ECE, Feb 2024
University of Chicago & UC San Diego Rising Stars in Data Science, Nov 2023

- **The Future of Large Language Models: Multimodality and Safety**

American Society for Clinical Pharmacology & Therapeutics Annual Meeting, March 2024
ACM Multimedia 2023 Workshop on Multimodal and Responsible Affective Computing, Oct 2023
Microsoft Research, July 2023
IBM Zurich, March 2023
DeepMind Language Team, Sept 2021

- **Foundations of Multimodal Machine Learning: Principles, Challenges, and Open Questions**

Guest lectures at CMU 10-707 Deep Learning, Peking University, University of Florida, 2023
CIFAR DLRL Summer School, July 2023
ICLR 2023 Workshop on Multimodal Representation Learning, April 2023
Guest lectures at CMU 10-707 Deep Learning, 05-618 Human-AI Interaction, 17-728 Machine Learning and Sensing, 2022
Harvard Medical School AI for Pathology Lab, Oct 2022
Heidelberg Laureate Forum, Sept 2022
UC Berkeley Speech Group, Sept 2022
Stanford University MedAI Group, Sept 2022
National University of Singapore, Aug 2022
Amazon AI, Aug 2022
Allen Institute of AI & University of Washington, June 2022
Carnegie Mellon University, May 2022

- **Brainish: Formalizing A Multimodal Language for Intelligence and Consciousness**

Peking University, March 2023
Models of Consciousness Conference, Sept 2022
International Joint Conference on Theoretical Computer Science, Aug 2022

- **Towards Real-World Socially Intelligent AI**

Facebook Fellowship Summit, Sept 2021
DeepMind Multimodal Team, Sept 2021
IJCAI 2021 Workshop on Multimodal Analytics, Aug 2021
Agency for Science, Technology, and Research Singapore, June 2021
Adobe Research, Jan 2021
Carnegie Mellon University, Oct 2020

- **Think Locally, Act Globally: Federated Learning with Local and Global Representations**

Agency for Science, Technology, and Research Singapore, June 2021
NeurIPS 2019 Workshop on Federated Learning, Dec 2019

- **Computational Modeling of Human Multimodal Language**

Google Research, July 2019
Riken AIP Tokyo Machine Learning Seminar, Jan 2019
Riken AIP Kyoto Machine Learning Seminar, Dec 2018
CMU Machine Learning Department Data Analysis Project Presentation, Apr 2018

Current Group

- MAS: Chanakya Ekbote, David Dai, Ray Song, Anku Rani
- EECS PhDs: Lucy Zhao, Megan Tjandrasuwita, Jaedong Hwang
- IDSS PhDs: Ao Qu

- EECS MEngs: Alex Dang, Daniel Li, Jada Li, Shreya Ravikumar, Kevin Zhao, Abhay Basireddy, Lucia Wang, Katherine Zeng, Kathy Choi, Josephine Lee, Daniel Guan, Edgar Morfin, Alex Hu, Divya Shyamal, Patrick Liu, Annika Vivekananthan
- UROPs: Justin Zhang, Regan Song, Marshall Taylor, Jeannie She, Isabella Zhu, Nilay Mishra
- Visiting researchers/post-docs: Fangneng Zhan, Kaichen Zhou, Erik Cambria, Keane Ong, Aaron Han, Ziyin Liu, Alistair Pernigo

Former Students

- Lily Chen (MIT BS, MEng → Stanford PhD)
- Devin Murphy (MIT BS, MEng → Univ of Washington PhD)
- Abe Hou (JHU BS → Stanford PhD)
- Guang-Yuan Hao (HKUST MS → Cornell PhD)
- Haofei Yu (CMU MS → UIUC PhD)
- Yun Cheng (CMU BS, MS → Princeton PhD)
- Rulin Shao (CMU MS → Univ of Washington PhD)
- Yuxin Xiao (CMU MS → MIT PhD)
- Dong Won Lee (CMU BS, MS → MIT PhD)
- Yiwei Lyu (CMU BS, MS → Univ of Michigan PhD)
CRA undergrad research award honorable mention, 2021
- Peter Wu (CMU BS, MS → UC Berkeley PhD)
- Terrance Liu (CMU MS → CMU PhD)
- Chengfeng Mao (CMU MS → MIT PhD)
- Irene Li (CMU BS, MS → SoundHound)
CRA undergrad research award honorable mention, 2019
- Xiang Fan (CMU BS → Univ of Washington PhD)
CRA undergrad research award honorable mention, 2022
- Jivat Neet (BITS Pilani BS → MSR → UC Berkeley PhD)
- Rohan Pandey (CMU BS → OpenAI)
Glushko prize for best senior thesis, 2023
- Samuel Yu (CMU BS → Jump Trading)
CRA undergrad research award finalist, 2022
- Ziyin Liu (CMU BS → Univ of Tokyo PhD → MIT Postdoc)

Professional Service

- Area Chair: NeurIPS, ICLR, ACL, EMNLP, NAACL, ACL Rolling Review, IJCAI, AISTATS, CHI
- Session Chair: ICLR, AAAI
- Workflow Chair: ICML 2019
- Co-organizer: Artificial Social Intelligence Workshop and Challenge at ICCV 2023, ECCV 2024
- Co-organizer: Workshop on Machine Learning for Cognitive and Mental Health at AAAI 2024
- Co-organizer: Tutorial on Multimodal Machine Learning at ICML 2023, ICMI 2023, CVPR 2022, NAACL 2022
- Co-organizer: Workshop on Foundation Models in Vision and Language at ICDM 2022
- Co-organizer: Workshop on Multimodal Artificial Intelligence at NAACL 2022, NAACL 2021
- Co-organizer: Workshop on Tensor Networks in Machine Learning at NeurIPS 2020
- Co-organizer: Grand Challenge and Workshop on Human Multimodal Language at ACL 2020, ACL 2018
- Conference Program Committee: NeurIPS, ICML, ICLR, ACL, EMNLP, NAACL, EACL, COLING, IJCNLP, AACL, COLM, ACL Rolling Review, CVPR, ICCV, ECCV, WACV, ACCV, AAAI, IJCAI, AISTATS, CHI, ICMI, FG, ACML, ML4H, CHIL, ACM Multimedia, ACM Multimedia Asia, ICASSP, Interspeech, SIGGRAPH
- Workshop Program Committee: NeurIPS workshop on Meta-Learning, NeurIPS workshop on ML for Health, ICLR workshop on Embodied Multimodal Learning, ICLR workshop on Never-ending RL, ICLR workshop on Enormous Language Models, ACL workshop on Multimodal Language, EMNLP workshop on NLP Open Source Software, EMNLP workshop on NLP Beyond Text, NAACL workshop on Trustworthy NLP, IJCAI workshop on Federated Learning, WWW workshop on NLP Beyond Text, ICRA workshop on Social Intelligence in Humans & Robots, NeurIPS workshop on Human Evaluation of Generative Models, NeurIPS workshop on Self-Supervised Learning, EMNLP workshop on Generation, Evaluation & Metrics
- Journal Reviewer: IEEE Transactions on Pattern Analysis and Machine Intelligence, IEEE Transactions on Affective Computing, IEEE Transactions on Audio, Speech and Language Processing, IEEE Transactions on Multimedia, IEEE Transactions on Cybernetics, IEEE Computational Intelligence Magazine, IEEE Signal Processing Letters, Elsevier Information Fusion, Elsevier Computer Speech and Language, Machine Learning, Transactions on Machine Learning Research, Neural Networks, Journal of Artificial Intelligence Research, Journal of Machine Learning Research, Medical Image Analysis, ACM Transactions on Sensor Networks, Proceedings of the National Academy of Sciences, Harvard Data Science Review
- MIT Committee on the Use of Humans as Experimental Subjects (COUHES) Member
- MIT PhD Admissions Committee
- MIT Faculty Reader for MIT Undergraduate Admissions
- CMU Machine Learning Blog Editorial Board: 2019, Chief Editor: 2020, 2021, 2022, 2023
- CMU AI Undergraduate Research Mentor: 2018, 2019, 2020, 2021
- CMU Graduate Applicant Support Program Organizer: 2020
- CMU Machine Learning Department PhD Admissions Committee: 2018, 2019, 2020, 2021, 2022, 2023
- CMU Machine Learning Department Masters Admissions Committee: 2017, 2018
- CMU Singapore Students Association Co-President: 2015

Skills

- Languages: English (fluent), Chinese (fluent)
- Programming Languages: Python, C, C++, R, MATLAB, Standard ML, Haskell, SystemVerilog
- Software: PyTorch, Tensorflow, Theano, Keras, $\LaTeX$