

JESSE THOMASON

jessethomason.com
thomason@gatech.edu

<http://glamor.rocks/>
Google Scholar Profile

RESEARCH and INDUSTRY POSITIONS

- Georgia Institute of Technology, Associate Professor 2026–present
- University of Southern California, Assistant Professor 2021–2026
- Amazon Alexa AI, Visiting Academic 2020–2021
- University of Washington, Postdoctoral Scholar 2018–20

EDUCATION

- PhD Computer Science, University of Texas at Austin 2013–18
- BS Computer Science; Mathematics, University of Pittsburgh 2009–13

TEACHING

University Courses: developed 4 (marked with *), taught 5.

- 5.* *CSCI 699: History of Language and Computing*: A role-playing reading seminar for PhD, MS, and undergraduate students that covers primary and secondary texts on the history of thought and paradigms surrounding language and computing. Spring 2025
Spring 2024
- 4.* *CSCI 444: Natural Language Processing*: This undergraduate course introduces core NLP concepts as well as the modern ML techniques behind large language models. I developed the course originally as CSCI 499: NLP for Interactive AI. Fall 2024
3. *CSCI 566: Deep Learning and its Applications*: This Master's course covers the core ML background for deep learning, and specific successful and widely used architectures in computer vision, natural language processing, and robotics. Spring 2023
- 2.* *CSCI 499: Natural Language Processing for Interactive AI*: This undergraduate course introduces core NLP concepts and builds towards AI agents that operate with text input and output as chatbots or embodied agents in simulated or real environments. Fall 2022
- 1.* *CSCI 699: Grounding Natural Language*: This PhD topics course covers multimodal machine learning problems and architectures where natural language serves as an interface between people and AI systems and ML tools. Spring 2022

RESEARCH SUPERVISION

Ph.D Student Alumni (7)

Abrar Anwar

Fall 2021–Spring 2026

Thesis: Language-Guided Robot Learning and Evaluation

Next: Robotics @ NVIDIA

Awards: Horatio Alger Scholar (2023),
Viterbi Research Fellow (2021)

Tejas Srinivasan

Fall 2021–Spring 2026

Thesis: Facilitating Reliable Human-AI Collaboration Under Uncertainty

Next: Microsoft 365

Awards: USC-Capital One CREDIF Fellowship (2025),
USC-Amazon Machine Learning Fellowship (2024),
Viterbi Undergraduate Research Mentoring Award (2023)

Ting-Yun Chang (Co-advised by Robin Jia)

Fall 2021–Spring 2026

Thesis: Dissecting LLMs: From Component Attribution to Actionable Insights

Next: Microsoft 365

Ishika Singh Fall 2021–Spring 2026
Thesis: Adapting Large Pre-Trained Models for Generalized Robot Reasoning and Control
Next: Robotics @ NVIDIA
Awards: Viterbi Research Fellow (2021)

Wang Bill Zhu (Co-advised by Robin Jia) Spring 2022–Spring 2026
Thesis: Domain-Adaptive Programming: Expanding the Boundary of What LLMs can Solve
Next: Postdoc with Serina Chang’s lab at UC Berkeley

Lee Kezar Summer 2022–Summer 2025
Thesis: Phonological Inductive Biases for Computationally Modeling American Sign Language Video
Next: Postdoc with Lorna Quandt’s Action and Brain Lab at Gallaudet University

Jesse Zhang (Co-advised by Erdem Biyik) Fall 2023–Spring 2025
Thesis: Scaling Robot Adaptation with Large Model Guidance
Next: Postdoc with Abhishek Gupta’s WEIRD Lab at the University of Washington

Current Ph.D. Students (8)

Leticia Pinto-Alva (Co-advised by Gale Lucas) Fall 2021–present
Milestones: Qualification Exam (Summer 2025)
Awards: Viterbi Outstanding Mentor Award (2023),
Viterbi Research Fellow (2021)

Yuliang Cai (Co-advised by Mohammad Rostami) Spring 2025–present
Milestones: Thesis Proposal (Spring 2026),
Qualification Exam (Spring 2025)

Zizhao Hu (Co-advised by Mohammad Rostami) Spring 2025–present
Milestones: Qualification Exam (Spring 2025)

Harshitha Rajaprakash Fall 2025–present

Dhananjay Ashok (Co-advised by Jon May) Fall 2025–present
Milestones: Qualification Exam (Spring 2026)
Awards: Viterbi Research Fellow (2024)

Wei Yang Fall 2025–present

Jie Cai Fall 2025–present

Xinyan Yu Spring 2026–present

Former Ph.D. Students (1)

Anthony Liang (Fall 2021–Summer 2024; → *Advised by Erdem Biyik and Stephen Tu; graduated Spring 2026*)

Master’s Student Collaborators (5 current; 13 total)

Aditeya Kamlesh Prajapati (2026–), Mousumi Das (2025–), Asmita Mohanty (2025–), Yusen Luo (2024–), Ryan Lindeborg (2024–), Mann Patel (2024–25; → *Capital One*), Jiahui Zhang (2024–25; → *UT Dallas CS PhD Program*), Yiming Tang (2023–25; → *Baidu NLP Team*), Chengqi Xu (2023), Chenfeng Huang (2023; → *UCLA Statistics PhD Program*), Athang Gupte (2023–24; → *Rockstar Games*), Wilbert Pumacay (2023–24; → *Allen Institute for AI Predoctoral Program*), Yuan Huang (2022–23)

Undergraduate Student Researchers (3 current; 38 total)

Ryan Nene (2025–), Emily Wang (2024–), Richard Peng (2024–), Trevor Asbery (2025; → *Google*), Mike Gee (2025; → *Yale MSCS Program*), Zian Zeng (2024; SURE Fellow), Vedant Raval (2024), Christina Wang (2024–25), Matthew Salaway (2024–25; → *Boeing*), Ryan Wang (2024–25; → *UC Berkeley CS PhD Program*), Lorena Yan (2024–25; → *NYU CS PhD Program*), Qiutong Yi (2024–26), David Bai (2024–25; CURVE Fellow), Keyu He (2023–25; → *CMU CS Masters Program*), Jaiv Doshi (2023), Yuxi Qian (2023), Zain Merchant (2023–25; CURVE Fellow; → *Google*), Cicily Chung (2023; CURVE Fellow), Nidhi Munikote (2023–24; CURVE Fellow), Miaosen Chai (2023–26; CURVE Fellow; → *UChicago CS PhD Program*), Riley Ashford (2023–24; SURE Fellow, CURVE Fellow), Emmanuel Ezirim (2023–24; VSI Fellow, CURVE Fellow), Gwen Bradforth

(2023–24; CURVE Fellow), Rohan Gupta (2023–25; Provost’s Fellow; →*Synopsys*), Abhinav Gupta (2023–26; Provost’s Fellow, URAP; →*NYU CS MS Program*), Riley Carlin (2023; →*Columbia Statistics PhD Program*), Furong Jia (2022–23; CURVE Fellow; →*Duke CS PhD Program*), Allen Chang (2022–23; →*NSFGRFP; UPenn CS PhD Program*), Aarav Monga (2022–23), Leslie Moreno (2022–26; CURVE Fellow; →*Jarriet*), Elle Szabo (2022–23), Chu Fang (2022–23; URAP), Julie Kim (2022–23; CURVE Fellow), Junu Song (2022; CURVE Fellow), Minh Ngoc Vu (2022; NSF Robotics REU), Kush Bhagat (2022; SURE Fellow), Chidera Iwudyke (2022; SURE Fellow), Tanis Sarbatananda (2022; LACC ASSURE Fellow)

Ph.D. Thesis Defense Committee Member (15) to: Tsung-Shan (Kevin) Yang (2026, USC); Xisen Jin (2025, USC); Zhaoxu Zhang (2025, USC); Bingjie Tang (2025, USC); Kasra Darvish (2025, UMBC); Digbalay Bose (2024, USC); Tom Groechel (2023, USC); Karl Pertsch (2023, USC); Chris Denniston (2023, USC); Lauren Klein (2023, USC); Sara Mohammadinejad (2022, USC); Sebastien Arnold (2022, USC); Aleksei Petrenko (2022, USC); Weiyu Lu (2022, Georgia Tech); Hexiang Hu (2021, USC); Shurjo Banarjee (2021, UMich).

GRANTS, CONTRACTS, and GIFTS

Raised approximately \$1,297,451 since Fall 2021, calculated as the sum of my shares of awarded grants.

Current (3)

14. IARPA Bias Effects and Notable Generative AI Limitations (BENGAL) grant for *Enabling Efficient Unlearning in Pretrained Large Language Models through Information Localization*; PI: Robin Jia; Co-PI: Jesse Thomason; Total \$567,769 (Thomason share \$82,518); Feb 2026–Aug 2027.
13. NVIDIA Accelerating Innovation in Academia gift for *Ensuring Reliable and Safe Pretrained Robot Policy Behavior*; PI: Jesse Thomason; DGX Spark (estimated value \$3,999); Oct 2025–Sep 2026.
12. National Science Foundation (NSF) CISE IIS grant for *HCC: Medium: Language-guided Haptic Texture Modeling and Signal Generation* (Award 2504241); PI: Heather Culbertson; Co-PI: Jesse Thomason; Total \$850,000 (Thomason share \$425,000); Oct 2025–Sep 2029.

Completed (11)

11. USC-Capital One Center for Responsible AI and Decision Making in Finance (CREDIF) grant for *Human-in-the-loop Multi-Agentive AI Systems for Task-oriented Dialogue*; PI: Jesse Thomason; Total \$84,000; Aug 2025–Jul 2026.
10. DARPA ARC Safe and Assured Foundation Robots for Open Environments (SAFRON) grant for *Assuring Pretrained Robot Policy Behavior by Combining Statistical ML and Software Testing Methods* (Award HR0011-25-3-0154); PI: Jesse Thomason; Co-PIs: Souti Chattopadhyay, William G.J. Halfond; Total \$197,457 (Thomason share \$108,223); Jul 2025–Jun 2026.
9. DARPA Friction and Accountability in Conversational Transactions (FACT) AI Exploration grant for *BECAREFUL: Building Embodied Conversational Agent Reliability by Exerting Friction through Uncertain Language* (Award HR00112490376); PI: Dilek Hakkani-Tur; Co-PIs: Gokhan Tur, Malihe Alikhane, Jesse Thomason, Julia Hockenmeier; Total \$700,000 (Thomason share \$215,000); Mar 2024–Aug 2025.
8. USC Undergraduate Research Associates Program (URAP) gift for *Gesture for L2 Lexical Acquisition*; PI: Jesse Thomason; Total \$3,300; Aug 2023–Apr 2024.
7. Laboratory for Analytic Sciences grant for *Multimodal Transformers with Compositional Modules for Continual Learning*; PI: Mohammad Rostami; Co-PI: Jesse Thomason; Total \$13,333; Jan 2023–Apr 2023.
6. NSF Convergence Accelerator Track H grant for *Determining Community Needs for Accessibility Tools that Facilitate Programming Education and Workforce Readiness for Personas with Disabilities* (Award 2236320); PI: Maja Matarić; Co-PIs: Stephen Aguilar, Sook-Lei Liew, Gisele Ragusa, Jesse Thomason; Total \$750,000 (Thomason share \$92,576); Dec 2022–Nov 2024.

5. UPenn’s Alzheimer’s Disease Research Center supported by NIH ‘Penn Artificial Intelligence and Technology Collaboratory for Healthy Aging’ subaward for *An Accessible Machine Learning-Based ADRD Screening Tool for Families and Caregivers* (Award 5-P30-AG-073105-02); PI: Maja Matarić; Co-PI: Jesse Thomason; Thomason share \$8,350; Dec 2022–Nov 2024.
4. USC Undergraduate Research Associates Program (URAP) gift for *Glass is to Shatter as Rubber is to Bounce: Analogies for Natural Language Processing*; PI: Jesse Thomason; Total \$3,300; Aug 2022–Apr 2023.
3. Laboratory for Analytic Sciences grant for *Continual Learning of Few Shot Learners for Natural Language Processing*; PI: Mohammad Rostami; Co-PI: Jesse Thomason; Total \$40,000; May 2022–Dec 2022.
2. Army Research Laboratory: Army Artificial Intelligence Innovation Institute (A2I2) grant for *Communicating with Natural Language Dialogue for Teams of Intelligent Systems and Humans* (Award W911NF-23-2-0010); PI: David Traum; Co-PI: Jesse Thomason; Total \$247,588 (Thomason share \$192,852); Feb 2022–Nov 2025.
1. Amazon AWS Credits for Amazon Vising Academics gift for *Language-Guided Mobile Manipulators*; PI: Jesse Thomason; Total \$25,000; Sep 2021–Aug 2022.

HONORS and AWARDS

- Best Paper Award—
Special Interest Group on Discourse and Dialogue (SIGDIAL) 2026
- Best Paper Award—
2nd Workshop on Out-of-Distribution Generalization in Robotics @ RSS 2025
- Best Paper Award—
2nd Workshop on Nonverbal Cues for Human-Robot Cooperative Intelligence @ ICRA 2025
- 2nd Place Best Paper Award—
Interactive AI for Human-Centered Robotics Workshop @ Ro-MAN 2024
- Charles Lee Powell Faculty Research Award (USC) 2022
- National Science Foundation Graduate Research Fellowship (NSF GRFP) 2015
- Microelectronics and Computer Development (MCD) Fellowship, University of Texas at Austin 2013
- Computer Science Outstanding Undergraduate Student Award, University of Pittsburgh 2013
- Honors College Chancellor’s Scholarship, University of Pittsburgh 2009

SELECTED MEDIA COVERAGE

- Photo Essay: ‘Human Centered’ The USC School of Advanced Computing in 23 Images, USC Viterbi Magazine, November 2024.
- Scientists Are Putting ChatGPT Brains Inside Robot Bodies. What Could Possibly Go Wrong?, Scientific American, March 2024.

PUBLICATIONS

Refereed Journal Articles (6)

6. Bridging Language and Action: A Survey of Language-Conditioned Robot Manipulation. Hongkuan Zhou, Xiangtong Yao, Oier Mees, Yuan Meng, Ted Xiao, Yonatan Bisk, Jean Oh, Edward Johns, Mohit Shridhar, Dhruv Shah, **Jesse Thomason**, Kai Huang, Joyce Chai, Zhenshan Bing, and Alois Knoll. *International Journal of Robotics Research Special Issue: Foundation Models and Neural-Symbolic AI for Robotics*, 2026. IJRR
5. Better Slow than Sorry: Introducing Positive Friction for Reliable Dialogue Systems. Mert Inan, Anthony Sicilia, Suvodip Dey, Vardhan Dongre, **Tejas Srinivasan**, **Jesse Thomason**, Gokhan Tur, Dilek Hakkani-Tur, and Malihe Alikhani. *Transactions of the Association for Computational Linguistics*, 2025. TACL

4. ProgPrompt: Program generation for situated robot task planning using large language models. [Ishika Singh](#), Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, **Jesse Thomason**, and Animesh Garg. *Autonomous Robots*, 2023. AURO
3. Multimodal embodied attribute learning by robots for object-centric action policies. Xiaohan Zhang, Saeid Amiri, Jivko Sinapov, **Jesse Thomason**, Peter Stone, and Shiqi Zhang. *Autonomous Robots*, 2023. AURO
2. Jointly Improving Parsing and Perception for Natural Language Commands through Human-Robot Dialog. **Jesse Thomason**, Aishwarya Padmakumar, Jivko Sinapov, Nick Walker, Yuqian Jiang, Harel Yedidsion, Justin Hart, Peter Stone, and Raymond J. Mooney. *The Journal of Artificial Intelligence Research* 67, 2020. JAIR
1. BWIBots: A platform for bridging the gap between AI and human-robot interaction research. Piyush Khandelwal, Shiqi Zhang, Jivko Sinapov, Matteo Leonetti, **Jesse Thomason**, Fangkai Yang, Ilaria Gori, Maxwell Svetlik, Priyanka Khante, Vladimir Lifschitz, J. K. Aggarwal, Raymond J. Mooney, and Peter Stone. *The International Journal of Robotics Research*, 2017. IJRR

Refereed Conference Papers (71)

71. Multimodal Behavioral Typicality as a Training-Free Screening Signal for Dementia. [Leticia Pinto-Alva](#), Gale Lucas, Maja Matarić, and **Jesse Thomason**. *International Conference on Multimodal Interaction*, 2026. ICMI
70. TNG-CLIP: Training-Time Negation Data Generation for Negation Awareness of CLIP. [Yuliang Cai](#), **Jesse Thomason**, and Mohammad Rostami. *International Conference on Multimodal Interaction*, 2026. ICMI
69. Mechanistic Finetuning of Vision-Language-Action Models via Few-Shot Demonstrations. Chancharik Mitra, Yusen Luo, Raj Saravanan, Dantong Niu, Anirudh Pai, **Jesse Thomason**, Trevor Darrell, [Abrar Anwar](#), Deva Ramanan, and Roei Herzig. *The European Conference on Computer Vision*, 2026. ECCV
68. DiPS: Dialogue Policy Selection for High-Stakes Persuasion Agents. Tianyi Zhang, Mousumi Das, [Abrar Anwar](#), **Jesse Thomason**, and David Traum. *Special Interest Group on Discourse and Dialogue*, 2026. **Best Paper Award* SIGDIAL
67. Believing without Seeing: Quality Scores for Contextualizing Vision-Language Model Explanations. Keyu He, [Tejas Srinivasan](#), Brihi Joshi, Xiang Ren, **Jesse Thomason**, and Swabha Swayamdipta. *Association for Computational Linguistics*, 2026. ACL
66. Words that make SENSE: Sensorimotor Norms in Learned Lexical Token Representations. Abhinav Gupta, Toben H. Mintz, and **Jesse Thomason**. *Findings of Association for Computational Linguistics*, 2026. ACL Findings
65. Zero-Shot Iterative Formalization and Planning in Partially Observable Environments. Liancheng Gong, [Wang Bill Zhu](#), **Jesse Thomason**, and Li Zhang. *Findings of Association for Computational Linguistics*, 2026. ACL Findings
64. Adjust for Trust: Mitigating Trust-Induced Inappropriate Reliance on AI Assistance. [Tejas Srinivasan](#) and **Jesse Thomason**. *ACM Conference on Intelligent User Interfaces*, 2026. IUI
63. PSALM-V: Automating Symbolic Planning in Interactive Visual Environments with Large Language Models. [Wang Bill Zhu](#), Miaosen Chai, [Ishika Singh](#), Robin Jia, and **Jesse Thomason**. *International Conference on Robotics and Automation*, 2026. ICRA
62. PEEK: Guiding and Minimal Image Representations for Zero-Shot Generalization of Robot Manipulation Policies. [Jesse Zhang](#), Marius Memmel, Kevin Kim, Dieter Fox, **Jesse Thomason**, Fabio Ramos, Erdem Biyik, Abhishek Gupta, and Anqi Li. *International Conference on Robotics and Automation*, 2026. ICRA

61. HAND Me the Data: Fast Robot Adaptation via Hand Path Retrieval. Matthew Hong, Anthony Liang, Kevin Kim, Harshitha Rajapraakash, Jesse Thomason, Erdem Biyik, and Jesse Zhang. *International Conference on Robotics and Automation*, 2026. ICRA
60. VeriMoA: A Mixture-of-Agents Framework for Spec-to-HDL Generation. Heng Ping, Arijit Bhattacharjee, Peiyu Zhang, Shixuan Li, Wei Yang, Anzhe Cheng, Xiaole Zhang, Jesse Thomason, Ali Jannesari, Nesreen Ahmed, and Paul Bogdan. *Machine Learning and Systems*, 2026. MLSys
59. Learning to Deliberate: Meta-policy Collaboration for Agentic LLMs with Multi-agent Reinforcement Learning. Wei Yang and Jesse Thomason. *Conference on Artificial Intelligence*, 2026. AAAI
58. Why Do Some Inputs Break Low-Bit LLM Quantization?. Ting-Yun Chang, Muru Zhang, Jesse Thomason, and Robin Jia. *Empirical Methods in Natural Language Processing*, 2025. EMNLP
57. Large Language Models Do Multi-Label Classification Differently. Marcus Ma, Georgios Chochlakis, Niyantha Maruthu Pandiyan, Jesse Thomason, and Shrikanth Narayanan. *Empirical Methods in Natural Language Processing*, 2025. EMNLP
56. Can VLMs Recall Factual Associations From Visual References?. Dhananjay Ashok, Ashutosh Chaubey, Hirona J. Arai, Jonathan May, and Jesse Thomason. *Findings of Empirical Methods in Natural Language Processing*, 2025. EMNLP Findings
55. Multimodal Synthetic Data Training and Model Collapse: Insights from VLMs and Diffusion Models. Zizhao Hu, Mohammad Rostami, and Jesse Thomason. *International Conference on Multimodal Interaction*, 2025. ICMI
54. ReWiND: Language-Guided Rewards Teach Robot Policies without New Demonstrations. Jiahui Zhang, Yusen Luo, Abrar Anwar, Sumedh Anand Sontakke, Joseph J Lim, Jesse Thomason, Erdem Biyik, and Jesse Zhang. *Conference on Robot Learning*, 2025. CoRL
53. Efficient Evaluation of Multi-Task Robot Policies With Active Experiment Selection. Abrar Anwar, Rohan Gupta, Zain Merchant, Sayan Ghosh, Willie Neiswanger, and Jesse Thomason. *Conference on Robot Learning*, 2025. CoRL
52. The American Sign Language Knowledge Graph: Infusing ASL Models with Linguistic Knowledge. Lee Kezar, Nidhi Munikote, Zian Zeng, Zed Sevcikova Sehyr, Naomi Caselli, and Jesse Thomason. *Findings of North American Chapter of the Association for Computational Linguistics*, 2025. NAACL Findings
51. Language Models can Infer Action Semantics for Classical Planners from Environment Feedback. Wang Bill Zhu, Ishika Singh, Robin Jia, and Jesse Thomason. *North American Chapter of the Association for Computational Linguistics*, 2025. NAACL
50. Promoting Cognitive Health in Elder Care with Large Language Model-Powered Socially Assistive Robots. Maria R. Lima, Amy O’Connell, Feiyang Zhou, Alethea Nagahara, Avni Hulyalkar, Anura Deshpande, Jesse Thomason, Ravi Vaidyanathan, and Maja Matarić. *Conference on Human Factors in Computing Systems*, 2025. CHI
49. Systematic Translation from Natural Language Robot Task Descriptions to STL. Sara Mohammadinejad, Sheryl Paul, Yuan Xia, Vidisha Kudalkar, Jesse Thomason, and Jyotirmoy V. Deshmukh. *Bridging the Gap Between AI and Reality*, 2024. AISoLA
48. When Parts are Greater Than Sums: Individual LLM Components Can Outperform Full Models. Ting-Yun Chang, Jesse Thomason, and Robin Jia. *Empirical Methods in Natural Language Processing*, 2024. EMNLP
47. Contrast Sets for Evaluating Language-Guided Robot Policies. Abrar Anwar, Rohan Gupta, and Jesse Thomason. *Conference on Robot Learning*, 2024. CoRL
46. ViSaRL: Visual Reinforcement Learning Guided by Human Saliency. Anthony Liang, Jesse Thomason, and Erdem Biyik. *Intelligent Robots and Systems*, 2024. IROS

45. Selective "Selective Prediction": Reducing Unnecessary Abstention in Vision-Language Reasoning. Tejas Srinivasan, Jack Hessel, Tanmay Gupta, Bill Yuchen Lin, Yejin Choi, **Jesse Thomason**, and Khyathi Raghavi Chandu. *Findings of Association for Computational Linguistics*, 2024. ACL Findings
44. The COLOSSEUM: A Benchmark for Evaluating Generalization for Robotic Manipulation. Wilbert Pumacay, **Ishika Singh**, Jiafei Duan, Ranjay Krishna, **Jesse Thomason**, and Dieter Fox. *Robotics: Science and Systems*, 2024. RSS
43. Which One? Leveraging Context Between Objects and Multiple Views for Language Grounding. Chancharik Mitra, **Abrar Anwar**, Rodolfo Corona, Dan Klein, Trevor Darrell, and **Jesse Thomason**. *North American Chapter of the Association for Computational Linguistics*, 2024. NAACL
42. Do Localization Methods Actually Localize Memorized Data in LLMs? A Tale of Two Benchmarks. **Ting-Yun Chang**, **Jesse Thomason**, and Robin Jia. *North American Chapter of the Association for Computational Linguistics*, 2024. NAACL
41. Efficient End-to-End Visual Document Understanding with Rationale Distillation. **Wang Bill Zhu**, Alekh Agarwal, Mandar Joshi, Robin Jia, **Jesse Thomason**, and Kristina Toutanova. *North American Chapter of the Association for Computational Linguistics*, 2024. NAACL
40. Chain-of-Questions Training with Latent Answers for Robust Multistep Question Answering. **Wang Bill Zhu**, **Jesse Thomason**, and Robin Jia. *Empirical Methods in Natural Language Processing*, 2023. EMNLP
39. Task-Attentive Transformer Architecture for Continual Learning of Vision-and-Language Tasks Using Knowledge Distillation. **Yuliang Cai**, **Jesse Thomason**, and Mohammad Rostami. *Findings of Empirical Methods in Natural Language Processing*, 2023. EMNLP Findings
38. The Sem-Lex Benchmark: Modeling ASL Signs and Their Phonemes. **Lee Kezar**, Elana Pontecorvo, Adele Daniels, Connor Baer, Ruth Ferster, Lauren Berger, **Jesse Thomason**, Zed Sevcikova Sehry, and Naomi Caselli. *Conference on Computers and Accessibility*, 2023. ASSETS
37. Exploring Strategies for Modeling Sign Language Phonology. **Lee Kezar**, Riley Carlin, **Tejas Srinivasan**, Zed Sevcikova Sehry, Naomi Caselli, and **Jesse Thomason**. *European Symposium on Artificial Neural Networks*, 2023. ESANN
36. RREx-BoT: Remote Referring Expressions with a Bag of Tricks. Gunnar Sigurdsson, **Jesse Thomason**, Gaurav Sukhatme, and Robinson Piramuthu. *Intelligent Robots and Systems*, 2023. IROS
35. I2I: Initializing Adapters with Improvised Knowledge. **Tejas Srinivasan**, Furong Jia, Mohammad Rostami, and **Jesse Thomason**. *Conference on Lifelong Learning Agents*, 2023. CoLLAs
34. Multimodal Speech Recognition for Language-Guided Embodied Agents. Allen Chang, Xiaoyuan Zhu, Aarav Monga, Seoho Ahn, **Tejas Srinivasan**, and **Jesse Thomason**. *Annual Conference of the International Speech Communication Association*, 2023. INTER-SPEECH
33. Iterative Vision-and-Language Navigation. Jacob Krantz, Shurjo Banerjee, **Wang Bill Zhu**, Jason J. Corso, Peter Anderson, Stefan Lee, and **Jesse Thomason**. *Computer Vision and Pattern Recognition*, 2023. CVPR
32. ProgPrompt: Generating Situated Robot Task Plans using Large Language Models. **Ishika Singh**, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, **Jesse Thomason**, and Animesh Garg. *International Conference on Robotics and Automation*, 2023. ICRA
31. Improving Sign Recognition with Phonology. **Lee Kezar**, **Jesse Thomason**, and Zed Sevcikova Sehry. *European Chapter of the Association for Computational Linguistics*, 2023. EACL

30. Geolocated Social Media Posts are Happier: Understanding the Characteristics of Check-in Posts on Twitter. Julie Jiang, **Jesse Thomason**, Francesco Barbieri, and Emilio Ferrara. *Web Sciences*, 2023. WebSci
29. ALFRED-L: Investigating the Role of Language for Action Learning in Interactive Visual Environments. Arjun Akula, Spandana Gella, Aishwarya Padmakumar, Mahdi Namazifar, Mohit Bansal, **Jesse Thomason**, and Dilek Hakkani-Tur. *Empirical Methods in Natural Language Processing*, 2022. EMNLP
28. Generalization Differences between End-to-End and Neuro-Symbolic Vision-Language Reasoning Systems. Wang Bill Zhu, **Jesse Thomason**, and Robin Jia. *Findings of Empirical Methods in Natural Language Processing*, 2022. EMNLP Findings
27. CLiMB: A Continual Learning Benchmark for Vision-and-Language Tasks. Tejas Srinivasan, Ting-Yun Chang, Leticia Pinto-Alva, Georgios Chochlakis, Mohammad Rostami, and **Jesse Thomason**. *Neural Information Processing Systems*, 2022. NeurIPS
26. Vision-and-Language Navigation: A Survey of Tasks, Methods, and Future Directions. Jing Gu, Eliana Stefani, Qi Wu, **Jesse Thomason**, and Xin Eric Wang. *Association for Computational Linguistics*, 2022. ACL
25. TEACH: Task-driven Embodied Agents that Chat. Aishwarya Padmakumar, **Jesse Thomason**, Ayush Shrivastava, Patrick Lange, Anjali Narayan-Chen, Spandana Gella, Robinson Piramuthu, Gokhan Tur, and Dilek Hakkani-Tur. *Conference on Artificial Intelligence*, 2022. AAAI
24. Language Grounding with 3D Objects. **Jesse Thomason**, Mohit Shridhar, Yonatan Bisk, Chris Paxton, and Luke Zettlemoyer. *Conference on Robot Learning*, 2021. CoRL
23. The RobotSlang Benchmark: Dialog-guided Robot Localization and Navigation. Shurjo Banerjee, **Jesse Thomason**, and Jason J. Corso. *Conference on Robot Learning*, 2020. CoRL
22. Experience Grounds Language. Yonatan Bisk, Ari Holtzman, **Jesse Thomason**, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, Nicolas Pinto, and Joseph Turian. *Empirical Methods in Natural Language Processing*, 2020. EMNLP
21. RMM: A Recursive Mental Model for Dialog Navigation. Homero Roman Roman, Yonatan Bisk, **Jesse Thomason**, Asli Celikyilmaz, and Jianfeng Gao. *Findings of Empirical Methods in Natural Language Processing*, 2020. EMNLP Findings
20. Interpreting Black Box Models via Hypothesis Testing. Collin Burns, **Jesse Thomason**, and Wesley Tansey. *Foundations of Data Science*, 2020. FODS
19. ALFRED: A Benchmark for Interpreting Grounded Instructions for Everyday Tasks. Mohit Shridhar, **Jesse Thomason**, Daniel Gordon, Yonatan Bisk, Winson Han, Roozbeh Mottaghi, Luke Zettlemoyer, and Dieter Fox. *Computer Vision and Pattern Recognition*, 2020. CVPR
18. Vision-and-Dialog Navigation. **Jesse Thomason**, Michael Murray, Maya Cakmak, and Luke Zettlemoyer. *Conference on Robot Learning*, 2019. CoRL
17. Improving Robot Success Detection using Static Object Data. Rosario Scalise, **Jesse Thomason**, Yonatan Bisk, and Siddhartha Srinivasa. *Intelligent Robots and Systems*, 2019. IROS
16. Augmenting Knowledge through Statistical, Goal-oriented Human-Robot Dialog. Saeid Amiri, Sujay Bajracharya, Cihangir Goktolga, **Jesse Thomason**, and Shiqi Zhang. *Intelligent Robots and Systems*, 2019. IROS
15. Shifting the Baseline: Single Modality Performance on Visual Navigation & QA. **Jesse Thomason**, Daniel Gordon, and Yonatan Bisk. *North American Chapter of the Association for Computational Linguistics*, 2019. NAACL

14. Improving Grounded Natural Language Understanding through Human-Robot Dialog. **Jesse Thomason**, Aishwarya Padmakumar, Jivko Sinapov, Nick Walker, Yuqian Jiang, Harel Yedidsion, Justin Hart, Peter Stone, and Raymond J. Mooney. *International Conference on Robotics and Automation*, 2019. ICRA
13. Prospection: Interpretable Plans From Language By Predicting the Future. Chris Paxton, Yonatan Bisk, **Jesse Thomason**, Arunkumar Byravan, and Dieter Fox. *International Conference on Robotics and Automation*, 2019. ICRA
12. Multi-modal Predicate Identification using Dynamically Learned Robot Controllers. Saeid Amiri, Suhua Wei, Shiqi Zhang, Jivko Sinapov, **Jesse Thomason**, and Peter Stone. *International Joint Conference on Artificial Intelligence*, 2018. IJCAI
11. Guiding Exploratory Behaviors for Multi-Modal Grounding of Linguistic Descriptions. **Jesse Thomason**, Jivko Sinapov, Raymond J. Mooney, and Peter Stone. *Conference on Artificial Intelligence*, 2018. AAAI
10. Maximum-Variance Total Variation Denoising for Interpretable Spatial Smoothing. Wesley Tansey, **Jesse Thomason**, and James G. Scott. *Conference on Artificial Intelligence*, 2018. AAAI
9. Opportunistic Active Learning for Grounding Natural Language Descriptions. **Jesse Thomason**, Aishwarya Padmakumar, Jivko Sinapov, Justin Hart, Peter Stone, and Raymond J. Mooney. *Conference on Robot Learning*, 2017. CoRL
8. Improving Black-box Speech Recognition using Semantic Parsing. Rodolfo Corona, **Jesse Thomason**, and Raymond J. Mooney. *International Joint Conference on Natural Language Processing*, 2017. IJCNLP
7. Multi-Modal Word Synset Induction. **Jesse Thomason** and Raymond J. Mooney. *International Joint Conference on Artificial Intelligence*, 2017. IJCAI
6. Integrated Learning of Dialog Strategies and Semantic Parsing. Aishwarya Padmakumar, **Jesse Thomason**, and Raymond J. Mooney. *European Chapter of the Association for Computational Linguistics*, 2017. EACL
5. Learning Multi-Modal Grounded Linguistic Semantics by Playing "I Spy". **Jesse Thomason**, Jivko Sinapov, Maxwell Svetlik, Peter Stone, and Raymond J. Mooney. *International Joint Conference on Artificial Intelligence*, 2016. IJCAI
4. Learning to Interpret Natural Language Commands through Human-Robot Dialog. **Jesse Thomason**, Shiqi Zhang, Raymond J. Mooney, and Peter Stone. *International Joint Conference on Artificial Intelligence*, 2015. IJCAI
3. Integrating Language and Vision to Generate Natural Language Descriptions of Videos in the Wild. **Jesse Thomason**, Subhashini Venugopalan, Sergio Guadarrama, Kate Saenko, and Raymond J. Mooney. *Conference on Computational Linguistics*, 2014. COLING
2. Prosodic Entrainment and Tutoring Dialogue Success. **Jesse Thomason**, Huy Nguyen, and Diane Litman. *Artificial Intelligence in Education*, 2013. AIED
1. Differences in User Responses to a Wizard-of-Oz versus Automated System. **Jesse Thomason** and Diane Litman. *North American Chapter of the Association for Computational Linguistics*, 2013. NAACL

Refereed Workshop Papers (14)

14. Value-Aware Stochastic KV Cache Eviction for Reasoning Models. Ting-Yun Chang, Harvey Yiyun Fu, Deqing Fu, Chenghao Yang, **Jesse Thomason**, and Robin Jia. *Efficient Reasoning Workshop*, 2026. CoLM
13. RobotFleet: An Open-Source Framework for Centralized Multi-Robot Task Planning. Rohan Gupta, Trevor Asbery, Zain Merchant, Abrar Anwar, and **Jesse Thomason**. *Workshop on Scalable and Resilient Multi-Robot Systems*, 2025. RSS

12. M3PT: A Transformer for Multimodal, Multi-Party Social Signal Prediction with Person-aware Blockwise Attention. Yiming Tang, **Abrar Anwar**, and **Jesse Thomason**. *The 2nd Workshop on Nonverbal Cues for Human-Robot Cooperative Intelligence*, 2025. **Best Paper Award* ICRA
11. TwoStep: Multi-agent Task Planning using Classical Planners and Large Language Models. David Bai, **Ishika Singh**, David Traum, and **Jesse Thomason**. *Workshop on Language and Semantics of Task and Motion Planning*, 2025. ICRA
10. Evaluating Creativity and Deception in Large Language Models: A Simulation Framework for Multi-Agent Balderdash. Parsa Hejabi, Elnaz Rahmati, Alireza S. Ziabari, Preni Golazizian, **Jesse Thomason**, and Morteza Dehghani. *Wordplay: When Language Meets Games*, 2024. ACL
9. Generating Contextually-Relevant Navigation Instructions for Blind and Low Vision People. Zain Merchant, **Abrar Anwar**, Emily Wang, Souti Chattopadhyay, and **Jesse Thomason**. *Interactive AI for Human-Centered Robotics (InterAI) Workshop*, 2024. **Best Paper Award (2nd Place)* Ro-MAN
8. WinoViz: Probing Visual Properties of Objects Under Different States. Woojeong Jin, **Tejas Srinivasan**, **Jesse Thomason**, and Xiang Ren. *Workshop on Secure and Trustworthy Large Language Models (SeT LLM)*, 2024. ICLR
7. Exploring Strategies for Efficient Real-World VLN Evaluation. **Abrar Anwar**, Rohan Gupta, Elle Szabo, and **Jesse Thomason**. *Workshop on Language and Robot Learning (LangRob)*, 2023. CoRL
6. Does VLN Pretraining Work with Nonsensical or Irrelevant Instructions?. **Wang Bill Zhu**, **Ishika Singh**, Yuan Huang, Robin Jia, and **Jesse Thomason**. *Workshop on Open-Domain Reasoning Under Multi-Modal Settings (ODRUM)*, 2023. CVPR
5. Curriculum Learning for Data-Efficient Vision-Language Alignment. **Tejas Srinivasan**, Xiang Ren, and **Jesse Thomason**. *Workshop on Open-Domain Reasoning Under Multi-Modal Settings (ODRUM)*, 2023. CVPR
4. CLIP-Nav: Using CLIP for Zero-Shot Vision-and-Language Navigation. Vishnu Sashank Dorbala, Gunnar Sigurdsson, Robinson Piramuthu, **Jesse Thomason**, and Gaurav Sukhatme. *Workshop on Language and Robot Learning (LangRob)*, 2022. CoRL
3. LUMINOUS: Indoor Scene Generation for Embodied AI Challenges. Yizhou Zhao, Kaixiang Lin, Zhiwei Jia, Qiaozi Gao, Govind Thattai, **Jesse Thomason**, and Gaurav Sukhatme. *Controllable Generative Modeling in Language and Vision (CtrlGen) Workshop*, 2021. NeurIPS
2. Embodied BERT: A Transformer Model for Embodied, Language-guided Visual Task Completion. Alessandro Suglia, Qiaozi Gao, **Jesse Thomason**, Govind Thattai, and Gaurav Sukhatme. *Novel Ideas in Learning-to-Learn through Interaction (NILLI) Workshop*, 2021. EMNLP
1. Interaction and Autonomy in RoboCup@Home and Building-Wide Intelligence. Justin Hart, Harel Yedidsion, Yuqian Jiang, Nick Walker, Rishi Shah, **Jesse Thomason**, Aishwarya Padmakumar, Rolando Fernandez, Jivko Sinapov, Raymond J. Mooney, and Peter Stone. *AI-HRI AAAI Fall Symposium Series*, 2018. AAAI

Thesis Work (2)

2. Continually Improving Grounded Natural Language Understanding through Human-Robot Dialog. **Jesse Thomason**. *Department of Computer Science, The University of Texas at Austin*, 2018. Thesis
1. Continuously Improving Natural Language Understanding for Robotic Systems through Semantic Parsing, Dialog, and Multi-modal Perception. **Jesse Thomason**. *Doctoral Dissertation Proposal*, 2016. Thesis

INVITED TALKS

Invited Summit, Workshop, and Industry Talks

- Large Foundation Model for Interactive Robot Learning Workshop RSS'25
- Workshop on Semantic Reasoning and Goal Understanding in Robotics RSS'25
- Workshop on Language and Robot Learning CoRL'24
- Summit on Responsible Computing, AI, and Society @ Georgia Tech 2024
- Human-Robot Interaction for Wellbeing Workshop RO-MAN'24
- Shonan Meeting on Intelligent Interaction with Autonomous Assistants in the Wild 2024
- Berggruen Institute: Embodiment & Intelligence Salon and Workshop 2024
- 6th Robot Learning Workshop NeurIPS'23
- Workshop on Interactive Learning with Implicit Human Feedback ICLM'23
- Workshop on NLP for Conversational AI ACL'20
- Visually Grounded Interaction and Language (ViGIL) Workshop NeurIPS'19
- Semantic Policy and Action Representations for Autonomous Robots (SPAR) Workshop IROS-19
- Microsoft Research 2019

Selected Invited Colloquium Talks and Seminar Lectures

- CSCI 544: Natural Language Processing; Guest Lecture for instructor Robin Jia Spring 2026
- UT Austin Forum for Artificial Intelligence 2025
- Penn NLP Computational Linguistics lunch 2025
- Utah Robotics Center Seminar 2024
- CMU LTI Colloquium 2023
- UT Austin; CS 395T: Grounded Natural Language Processing 2022
- USC/ISI NL Seminar 2021
- UPenn; CIS700 2020
- Stanford NLP Seminar 2020
- UPitt; CS3730 2020
- Princeton; COS 598C 2020
- Utah Robotics Center Seminar 2019

SERVICE

Journal Editor

- **Action Editor:** Transactions of the Association for Computational Linguistics (TACL) 2025–6
- **Guest Editor:** Robotics and Autonomous Systems Special Issue: Semantic Policy and Action Representations for Autonomous Robots 2021–2

Conference Organization / Senior Area Chair / Area Chair

- **Session Chair:** AAAI-26—Natural Language Processing
- **Area Chair:** CoLM-26, CoLM-25
- **Associate Editor:** ROMAN-26
- **Area Chair:** AAAI-26
- **Area Chair:** EMNLP-25—Language Grounding to Vision, Robotics and Beyond
- **Session Chair:** RSS-25—Vision Language Action (VLA) Models
- **Area Chair:** RSS-25
- **Senior Area Chair:** NAACL-25—Multimodality and Language Grounding to Vision, Robotics and Beyond
- **Senior Area Chair:** EMNLP-24—Multimodality and Language Grounding to Vision, Robotics and Beyond
- **Workshops and Tutorials Co-Chair:** IEEE RO-MAN 2024
- **Area Chair:** CoLM-24
- **Area Chair:** ACL-24—Language Grounding to Vision, Robotics and Beyond
- **Area Chair:** NAACL-24—Language Grounding to Vision, Robotics and Beyond

- **Publication Chair:** ACL-23
- **Organizer:** Workshop on Language and Robot Learning (LangRob) at CoRL-22
- **Senior Area Chair:** EMNLP-22—Speech, Vision, Robotics, Multimodal Grounding
- **Area Chair:** COLING-22—Multimodal and Grounded Language Acquisition, HRI
- **Session Chair:** NAACL-22—6A: Language Grounding to Vision 2
- **Organizer:** Semantic Policy and Action Representations for Autonomous Robots (SPAR) Workshop at IROS-21
- **Area Chair:** AKBC-21—Computer Vision
- **Area Chair:** ACL-IJCNLP-21—Language Grounding to Vision, Robotics and Beyond
- **Area Chair:** NAACL-21—Language Grounding to Vision, Robotics and Beyond
- **Area Chair:** IJCAI-21
- **Organizer:** Embodied Vision, Actions & Language Workshop (EVAL) at ECCV-20
- **Area Chair:** ACL-20—Language Grounding to Vision, Robotics and Beyond
- **Organizer:** First Workshop on Advances in Language and Vision Research (ALVR) at ACL-20
- **Co-Chair:** Combined Workshop on Spatial Language Understanding (SpLU) and Grounded Communication for Robotics (RoboNLP) at NAACL-19
- **Co-Chair:** Special Session on Physically Situated Dialog (RoboDIAL) at SIGDIAL-18
- **Organizer:** Workshop on Communicating with Robots Naturally (CWRN) at RSS-18

Journal Reviewing

- | | |
|--|-------------------|
| • The International Journal of Robotics Research (IJRR) | 1 article; 24 |
| • IEEE Robotics and Automation Letters (RA-L) | 3 articles; 22-23 |
| • Transactions of the Association for Computational Linguistics (TACL) | 1 article; 21 |
| • RSS 2021 Special Issue | 1 article |
| • IEEE Transactions on Robotics (T-RO) | 1 article |
| • Computational Linguistics (CL) | 2 articles |
| • Elsevier Journal of Artificial Intelligence (AIJ) | 1 article |
| • Springer Autonomous Agents and Multi-Agent Systems (AGNT) | 1 article |
| • ACM Transactions on Interactive Intelligent Systems (TiiS) | 1 article |
| • Springer Autonomous Robots (AURO) | 1 article |

Conference Reviewing

- | | <i>Service Years</i> |
|---|----------------------|
| • European Conference on Computer Vision (ECCV) | 26 |
| • Robotics: Science and Systems (RSS) Workshop Proposals | 26 |
| • International Conference on Machine Learning (ICML) | 19, *26 |
| Recognized as “Gold” Technical Reviewer | |
| • Computer Vision and Pattern Recognition (CVPR) | 22-23, 26 |
| • Robotics: Science and Systems (RSS) | 17-18, 21, 26 |
| • European Chapter of the Association for Computational Linguistics (EACL) | 21, 26 |
| • Human-Robot Interaction (HRI) | 18, 22, 26 |
| • International Conference on Robotics and Automation (ICRA) | 21, 23-24, 26 |
| • Winter Conference on Applications of Computer Vision (WACV) | 26 |
| • Conference on Robot Learning (CoRL) | 18-21, 25 |
| • Robotics: Science and Systems (RSS) Pioneers; senior reviewer | 25 |
| • International Conference on Intelligent Robots and Systems (IROS) | 18-19, 21, 24 |
| • Association for Computational Linguistics (ACL) | 18-19, 24 |
| • North American Chapter of the Association for Computational Linguistics (NAACL) | 16, 19, 24 |
| • NeurIPS Datasets & Benchmarks | 22 |
| • ACL Rolling Review – off cycle | 9 submissions |
| • International Conference on Learning Representations (ICLR) | 21 |
| • Empirical Methods in Natural Language Processing (EMNLP) | 17-20 |

- Conference on Computational Natural Language Learning (CoNLL) 19
- Advances in Neural Information Processing Systems (NeurIPS) 17, *19
- *Top 50% of reviewers.
- International Symposium on Robot and Human Interactive Communication (RO-MAN) 19
- Autonomous Agents and Multi-Agent Systems (AAMAS) 15-16, 19
- International Conference on Computational Linguistics (COLING) 18
- International Joint Conference on Artificial Intelligence (IJCAI) 17

Workshop Reviewing

- Dialog System Technology Challenge 9 (DSTC9) AAAI'21
- Human in the Loop Dialogue Systems (HLDS) NeurIPS'20
- Spatial Language Understanding (SpLU) EMNLP'20
- Language in Reinforcement Learning (LaReL) ICML'20
- Language Grounding for Robotics (RoboNLP) ACL'17

Service to the USC CS Department

- Member, PhD Fellowship Committee Spring'25
- Member, CS DEI Committee Fall'24–Spring'25
- Member, Teaching Faculty Promotion Committee Spring'23
- Member, CS DEI Committee Fall'22–Spring'23
- Chair, Area 1 (AI) PhD fellowship committee Fall'22–Spring'23
- Member, Engineering Faculty Council Fall'22
- Member, Area 1 (AI) PhD fellowship committee Fall'21–Spring'22
- Member, Area 4 (Robotics) PhD fellowship committee Fall'21–Spring'22