JESSE THOMASON

jessethomason.com
thomason.jesse@gmail.com

I am an Assistant Professor at the University of Southern California since Fall 2021. My research focuses on language grounding and natural language processing applications for robotics, including how dialog with humans can facilitate both robot task execution and learning.

RESEARCH and INDUSTRY POSITIONS

- University of Southern California, Assistant Professor 2021-. . .
- Amazon Alexa AI, Visiting Academic 2020-2021
- University of Washington, Research Associate (Postdoc) with Luke Zettlemoyer 2018-20

EDUCATION

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

FELLOWSHIPS and AWARDS

- Charles Lee Powell Faculty Research Award (USC) 2022
- National Science Foundation Graduate Research Fellowship (NSF GRFP) 2015

PUBLICATIONS

Papers and Preprints

- TwoStep: Multi-agent Task Planning using Classical Planners and Large Language Models
  Ishika Singh, David Traum, and **Jesse Thomason**.
  arXiv, 2024.
- ViSaRL: Visual Reinforcement Learning Guided by Human Saliency
  Anthony Liang, **Jesse Thomason**, and Erdem Biyik.
  arXiv, 2024.
- 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 (NAACL), 2024.
- Do Localization Methods Actually Localize Memorized Data in LLMs?
  Ting-Yun Chang, **Jesse Thomason**, and Robin Jia.
  North American Chapter of the Association for Computational Linguistics (NAACL), 2024.
- Efficient End-to-End Visual Document Understanding with Rationale Distillation
  Wang Zhu, Alekh Agarwal, Mandar Joshi, Robin Jia, **Jesse Thomason**, and Kristina Toutanova.
  North American Chapter of the Association for Computational Linguistics (NAACL), 2024.
- 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) @ ICLR, 2024.
- 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.
  arXiv, 2024.

- The COLOSSEUM: A Benchmark for Evaluating Generalization for Robotic Manipulation
  Wilbert Pumacay, Ishika Singh, Jiafei Duan, Ranjay Krishna, **Jesse Thomason**, and Dieter Fox.
  arXiv, 2024.
- Chain-of-Questions Training with Latent Answers for Robust Multistep Question Answering
  Wang Zhu, **Jesse Thomason**, and Robin Jia.
  Empirical Methods in Natural Language Processing (EMNLP), 2023.
- 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 (EMNLP Findings), 2023.
- Exploring Strategies for Efficient Real-World VLN Evaluation
  Abrar Anwar, Rohan Gupta, Elle Szabo, and **Jesse Thomason**.
  Workshop on Language and Robot Learning (LangRob) @ CoRL, 2023.
- 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 Sehyr, and Naomi Caselli.
  Conference on Computers and Accessibility (ASSETS), 2023.
- Exploring Strategies for Modeling Sign Language Phonology
  Lee Kezar, Riley Carlin, Tejas Srinivasan, Zed Sevcikova Sehyr, Naomi Caselli, and **Jesse Thomason**.
  European Symposium on Artificial Neural Networks (ESANN), 2023.
- RREx-BoT: Remote Referring Expressions with a Bag of Tricks
  Gunnar Sigurdsson, **Jesse Thomason**, Gaurav Sukhatme, and Robinson Piramuthu.
  Intelligent Robots and Systems (IROS), 2023.
- 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 (AURO), 2023.
- I2I: Initializing Adapters with Improvised Knowledge
  Tejas Srinivasan, Furong Jia, Mohammad Rostami, and **Jesse Thomason**.
  Conference on Lifelong Learning Agents (CoLLAs), 2023.
- 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 (INTERSPEECH), 2023.
- Iterative Vision-and-Language Navigation
  Jacob Krantz, Shurjo Banerjee, Wang Zhu, Jason J. Corso, Peter Anderson, Stefan Lee, and **Jesse Thomason**.
  Computer Vision and Pattern Recognition (CVPR), 2023.
- Does VLN Pretraining Work with Nonsensical or Irrelevant Instructions?
  Wang Zhu, Ishika Singh, Yuan Huang, Robin Jia, and **Jesse Thomason**.
  Workshop on Open-Domain Reasoning Under Multi-Modal Settings (ODRUM) @ CVPR, 2023.
- 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) @ CVPR, 2023.
- 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 (ICRA), 2023.

- Improving Sign Recognition with Phonology
  Lee Kezar, **Jesse Thomason**, and Zed Sevcikova Sehyr.
  European Chapter of the Association for Computational Linguistics (EACL), 2023.
- 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 (WebSci), 2023.
- 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 (AURO), 2023.
- 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) @ CoRL, 2022.
- 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 (EMNLP), 2022.
- Generalization Differences between End-to-End and Neuro-Symbolic Vision-Language Reasoning Systems
  Wang Zhu, **Jesse Thomason**, and Robin Jia.
  Findings of Empirical Methods in Natural Language Processing (EMNLP Findings), 2022.
- CLiMB: A Continual Learning Benchmark for Vision-and-Language Tasks
  Tejas Srinivasan, Ting-Yun Chang, Leticia Leonor Pinto Alva, Georgios Chochlakis, Mohammad Rostami, and **Jesse Thomason**.
  Neural Information Processing Systems (NeurIPS), 2022.
- VAuLT: Augmenting the Vision-and-Language Transformer with the Propagation of Deep Language Representations
  Georgios Chochlakis, Tejas Srinivasan, **Jesse Thomason**, and Shrikanth Narayanan.
  arXiv, 2022.
- Interactive Learning from Natural Language and Demonstrations using Signal Temporal Logic
  Sara Mohammadinejad, **Jesse Thomason**, and Jyotirmoy V. Deshmukh.
  arXiv, 2022.
- 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 (ACL), 2022.
- 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 (AAAI), 2022.
- 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 @ NeurIPS, 2021.

- Language Grounding with 3D Objects
  Jesse Thomason, Mohit Shridhar, Yonatan Bisk, Chris Paxton, and Luke Zettlemoyer.
  Conference on Robot Learning (CoRL), 2021.
- 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 @ EMNLP, 2021.
- The RobotSlang Benchmark: Dialog-guided Robot Localization and Navigation
  Shurjo Banerjee, **Jesse Thomason**, and Jason J. Corso.
  Conference on Robot Learning (CoRL), 2020.
- 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 (EMNLP), 2020.
- 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 (EMNLP Findings), 2020.
- Interpreting Black Box Models via Hypothesis Testing
  Collin Burns, **Jesse Thomason**, and Wesley Tansey.
  Foundations of Data Science (FODS), 2020.
- 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 (CVPR), 2020.
- 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 (JAIR) 67, 2020.
- Vision-and-Dialog Navigation
  Jesse Thomason, Michael Murray, Maya Cakmak, and Luke Zettlemoyer.
  Conference on Robot Learning (CoRL), 2019.
- Improving Robot Success Detection using Static Object Data
  Rosario Scalise, **Jesse Thomason**, Yonatan Bisk, and Siddhartha Srinivasa.
  Intelligent Robots and Systems (IROS), 2019.
- Augmenting Knowledge through Statistical, Goal-oriented Human-Robot Dialog
  Saeid Amiri, Sujay Bajracharya, Cihangir Goktolga, **Jesse Thomason**, and Shiqi Zhang.
  Intelligent Robots and Systems (IROS), 2019.
- 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 (NAACL), 2019.
- 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 (ICRA), 2019.
- 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 (ICRA), 2019.

- 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 (AAAI-FSS), 2018.
- 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 (IJCAI), 2018.
- Guiding Exploratory Behaviors for Multi-Modal Grounding of Linguistic Descriptions
  Jesse Thomason, Jivko Sinapov, Raymond J. Mooney, and Peter Stone.
  Conference on Artificial Intelligence (AAAI), 2018.
- Maximum-Variance Total Variation Denoising for Interpretable Spatial Smoothing
  Wesley Tansey, **Jesse Thomason**, and James G. Scott.
  Conference on Artificial Intelligence (AAAI), 2018.
- 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 (CoRL), 2017.
- Improving Black-box Speech Recognition using Semantic Parsing
  Rodolfo Corona, **Jesse Thomason**, and Raymond J. Mooney.
  International Joint Conference on Natural Language Processing (IJCNLP), 2017.
- Multi-Modal Word Synset Induction
  Jesse Thomason and Raymond J. Mooney.
  International Joint Conference on Artificial Intelligence (IJCAI), 2017.
- Integrated Learning of Dialog Strategies and Semantic Parsing
  Aishwarya Padmakumar, **Jesse Thomason**, and Raymond J. Mooney.
  European Chapter of the Association for Computational Linguistics (EACL), 2017.
- 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 (IJRR), 2017.
- 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 (IJCAI), 2016.
- 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 (IJCAI), 2015.
- 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 (COLING), 2014.
- Prosodic Entrainment and Tutoring Dialogue Success
  Jesse Thomason, Huy Nguyen, and Diane Litman.
  Artificial Intelligence in Education (AIED), 2013.
- 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 (NAACL), 2013.

Thesis work

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

TEACHING

CSCI 699: History of Language and Computing: This course is designed for early career PhD students with an interest in understanding the bases and common assumptions in modern natural language processing research. We will study the history of thought and paradigms surrounding language and computing. We will read original texts as well as retrospectives and summary arguments from influential writers and researchers in recent history as well as those predating modern computation. Students will draw connections between historical perspectives and abstractions to modern day technological innovations and assumptions in natural language processing. Students will develop a rich understanding of the historical context of their own work in computing and language, and be better prepared to situate their research contributions in the long context of language processing.

- Spring 2024 [syllabus].

CSCI 566: Deep Learning and its Applications: Recently, deep learning has advanced many AI-related problems: image retrieval, video analysis, natural language processing, self-driving, medical applications, and more. Our goal is to guide students to get familiar with these recent cutting-edge deep learning (DL) advances in computer vision and natural language processing. Through this course, students will gain a basic understanding of DL algorithms, and how to set up and solve problems involving deep learning techniques. The course will include a couple of practical assignments and a final course project. For the final course project, students will be encouraged to pick their own topics, but can also select from a provided list of projects.

- Spring 2023 [website] [syllabus].

CSCI 499: Natural Language Processing for Interactive AI: Natural Language Processing for Interactive AI is an upper division undergraduate course in which students explore how natural language can serve as an interaction medium between users and AI agents. We cover topics in natural language processing, computer vision, and machine learning, as well as the intersection of planning and search-oriented machine learning algorithms with such language understanding techniques and paradigms. The core modules of the course cover text classification, language modeling with LSTMs and word embeddings, attention mechanisms and Transformers, and multimodality and reinforcement learning. Deliverables include paper reviews, a paper presentation, three increasingly complex coding assignments, and a course project expected to be carried out throughout the semester.

- Fall 2022 [syllabus].

CSCI 699: Grounding Natural Language: Grounding Natural Language is a PhD seminar course introducing the broad space of both multimodal language processing, for example language and vision models, and language models for decision making, for example dialogue systems and language-guided robotics. The course explores the ways in which other sensory modalities, especially visual input and embodiment in 3-dimensional space, can influence and guide representation learning for language. Deliverables include hour-long paper presentations, in which students digest research papers and present them in the context of modern NLP, and a course project expected to be carried out throughout the semester.

- Spring 2022 [syllabus].

PROFESSIONAL ACTIVITIES and SERVICE

Invited Talks

- Utah Robotics Center Seminar 2024
- 6th Robot Learning Workshop 2023
- CMU LTI Colloquium 2023
- Workshop on Interactive Learning with Implicit Human Feedback ICLM'23
- USC/ISI NL Seminar 2021
- Stanford NLP Seminar 2020
- Workshop on NLP for Conversational AI ACL-20
- Visually Grounded Interaction and Language (ViGIL) Workshop NeurIPS-19
- Utah Robotics Center Seminar 2019
- Semantic Policy and Action Representations for Autonomous Robots (SPAR) Workshop IROS-19
- Microsoft Research 2019

Selected Invited Lectures

- USC Computational Human-Robot Interaction 2022
- USC CSCI 662 2021
- USC CSCI 699 2021
- UPenn CIS700 2020
- UPitt CS3730 2020
- Princeton COS 598C 2020

Journal Editor

- **Guest Editor**: Robotics and Autonomous Systems Special Issue: Semantic Policy and Action Representations for Autonomous Robots

Conference Organization / Senior Area Chair / Area Chair

- **Senior Area Chair**: EMNLP-24—Multimodality and Language Grounding to Vision, Robotics and Beyond
- **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

- Transactions of the Association for Computational Linguistics (TACL) — 1 article
- RSS 2021 Special Issue — 1 article
- IEEE Robotics and Automation Letters (RA-L) — 2 articles
- IEEE Transactions on Robotics (T-RO) — 1 article
- Computational Linguistics (CL) — 2 articles
- Elsevier Journal of Artificial Intillengence (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

- 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
- International Conference on Robotics and Automation (ICRA) — 21, 23-24
- Computer Vision and Pattern Recognition (CVPR) — 22-23
- NeurIPS Datasets & Benchmarks — 22
- ACL Rolling Review – off cycle — 9 submissions
- Human-Robot Interaction (HRI) — 18, 22
- Conference on Robot Learning (CoRL) — 18-21
- Robotics: Science and Systems (RSS) — 17-18, 21
- European Chapter of the Association for Computational Linguistics (EACL) — 21
- 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
- International Conference on Machine Learning (ICML) — 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