Stefan Lee

engr.oregonstate.edu/~leestef
leestef@oregonstate.edu

⚙ Research Positions

Assistant Professor September 2019 to Present
School of Electrical Engineering and Computer Science
at Oregon State University

Past Positions

Visiting Research Scientist Summer 2020
at Facebook AI Research

Research Scientist II August 2017 to August 2019
School of Interactive Computing
at Georgia Institute of Technology

Bradley Postdoctoral Associate August 2016 to August 2017
Machine Learning & Perception Group
at Virginia Tech with Dhruv Batra

Research Assistant May 2012 to August 2016
School of Informatics and Computing,
at Indiana University with David Crandall

Visiting Research Assistant August 2015 to November 2015
Machine Learning & Perception Group
at Virginia Tech with Dhruv Batra

Visiting Research Assistant May 2014 to August 2014
INRIA - WILLOW Project
at L'École Normale Superiéure and UC Berkley with ef Sivic and Alexei A. Efros.

🎓 Education

Ph.D., Computer Science – Indiana University 2016
 Thesis: Data-Driven Computer Vision for Science and the Humanities
 Committee: David Crandall (Chair), Chunfeng Huang, Predrag Radivojac, Michael Ryoo

M.S., Computer Science – Indiana University 2013

B.S., Computer Science – University of West Florida 2011

🎖 Honors & Awards

– **Best Paper Awards**
   International Conference on Learning Representations (ICLR) 2023
   One of 4 best papers (or top 0.08%) out of 4900 submissions to ICLR 2023 which is a top
   venue for research on deep learning.

   Conference on Empirical Methods in Natural Language Processing (EMNLP) 2017
   One of 4 best papers (or top 0.26%) out of 1500 submissions (1466 reviewed, 323 accepted)
   to EMNLP 2017 which is a top venue for research on AI with natural language capabilities.

   CVPR Workshop on Egocentric Vision 2014
   Awarded best papers (or top 7.6%) out of 13 accepted papers to the CVPR Workshop
   on Egocentic Vision, a core workshop for discussion of egocentric (or first-person) vision.

- **Outstanding Reviewer Awards**
  IEEE Conference on Computer Vision and Pattern Recognition (CVPR)   **2017,2019,2020**
    Recognition from areas chairs for quality reviewing (awarded to $\sim 8.5\%$ of reviewers).
  IEEE International Conference on Computer Vision (ICCV)   **2017**
    Recognition from areas chairs for quality reviewing (awarded to $\sim 4.6\%$ of reviewers).
  Neural Information Processing Systems (NeurIPS)   **2017-2018**
    Recognition from areas chairs for quality reviewing (awarded to $\sim 3.6\%$ of reviewers).
  International Conference on Learning Representations (ICLR)   **2018-2019**
  European Conference on Computer Vision (ECCV)   **2020**
- **Notable Area Chair (NeurIPS)**   **2023**

- **Outstanding Teaching Award – College of Engineering, Oregon State University**   **2023**
- **Outstanding Research Scientist – College of Computing, Georgia Tech**   **2019**
- **CoRL Best Presentation Finalist**   **2018**
- **DARPA Riser – Plenary Speaker**   **2018**
  Nominated by DARPA program manager to participate in the DARPA Riser event and was selected from 47 candidates as one of three representative speakers for the plenary session at DARPA's 60th anniversary conference D60.
- **Bradley Postdoctoral Fellowship (Virginia Tech)**   **2016**
- **HANDS Travel Award (CVPR)**   **2016**
- **Doctoral Consortium Travel Award (ICCV)**   **2015**
- **Dissertation Development Award (Indiana University)**   **2015**
- **Heidelberg Laureate Forum Selection (HLF Foundation)**   **2015**

SELECTED PUBLICATIONS

Peer-Reviewed Conference Papers (acceptance rates typically 25-30%)

1. Eric Slyman, Stefan Lee, Scott Cohen, Kushal Kafle. FairDeDup: Detecting and Mitigating Vision-Language Fairness Disparities in Semantic Dataset Deduplication. *Computer Vision and Pattern Recognition (CVPR), 2024.*

2. Xiangxi Shi, Zhonghua Wu, Stefan Lee. Viewpoint-Aware Visual Grounding in 3D Scenes. *Computer Vision and Pattern Recognition (CVPR), 2024.*

3. Skand Peri, Iain Lee, Chanho Kim, Li Fuxin, Tucker Hermans, Stefan Lee. Point Cloud Models Improve Visual Robustness in Robotic Learners. *International Conference on Robotics and Automation (ICRA), 2024.*

4. Xiangxi Shi, Stefan Lee. Benchmarking Out-of-Distribution Detection in Visual Question Answering. *Winter Conference on Applications of Computer Vision (WACV), 2024.*

5. Christopher Buss, Jasmin Mousavi, Mikhail Tokarev, Arash Termehchy, David Maier, Stefan Lee. Effective Entity Augmentation by Querying External Data Sources. *Vary Large Databases (VLDB), Volume 16, 2023.*

6. Jacob Krantz, Theophile Gervet, Karmesh Yadav, Austin Wang, Chris Paxton, Roozbeh Mottaghi, Dhruv Batra, Jitendra Malik, Stefan Lee, Davendra Singh Chaplot. Navigating to Object Specified by Images. *International Conference on Computer Vision, 2023.*

7. Eric Slyman, Minsuk Kahng, Stefan Lee. VLSlice: Interactive Vision-and-Language Slice Discovery. *International Conference on Computer Vision, 2023.*

8. Erik Wijmans, Manolis Savva, Irfan Essa, Stefan Lee, Ari Morcos, Dhruv Batra. Emergence of Maps in the Memories of Blind Navigation Agents. *International Conference on Learning Representation (ICLR), 2023.* [Best Paper: 4/4900=0.08% of submissions]

9. Jacob Krantz, Shurjo Banerjee, Wang Zhu, Jason Corso, Peter Anderson, Stefan Lee, Jesse Thomason. Iterative Vision-and-Language Navigation. *Computer Vision and Pattern Recognition (CVPR), 2023.*

10. Zijiao Yang, Arjun Majumdar, Stefan Lee. Behavioral Analysis of Vision-and-Language Navigation Agents. *Computer Vision and Pattern Recognition (CVPR), 2023.*

11. Jacob Krantz, Stefan Lee. Sim-2-Sim Transfer for Vision-and-Language Navigation in Continuous Environments. *European Conference on Computer Vision (ECCV), 2022.* [Oral Paper: 157/5804=2.7% of submissions]

12. Jacob Krantz, Aaron Gokaslan, Dhruv Batra, Stefan Lee, Oleksandr Maksymets. Waypoint Models for Instruction-guided Navigation in Continuous Environments. *International Conference on Computer Vision (ICCV), 2021.* [Oral Paper: 210/6236=3.3% of submissions]

13. Abhinav Moudgil, Arjun Majumdar, Harsh Agrawal, Stefan Lee, Dhruv Batra. SOAT: A Scene- and Object-Aware Transformer for Vision-and-Language Navigation. *Neural Information Processing Systems, 2021.*

14. Yilin Yang, Akiko Eriguchi, Alexandre Muzio, Prasad Tadepalli, Stefan Lee, Hany Hassan. Improving Multilingual Translation by Representation and Gradient Regularization. *Empirical Methods in Natural Language Processing (EMNLP), 2021.*

15. Oleksandr Maksymets, Vincent Cartillier, Aaron Gokaslan, Erik Wijmans, Wojciech Galuba, Stefan Lee, Dhruv Batra. THDA: Treasure Hunt Data Augmentation for Semantic Navigation. *International Conference on Computer Vision (ICCV), 2021.*

16. Aayam Shrestha, Stefan Lee, Prasad Tadepalli, Alan Fern. DeepAveragers: Offline Reinforcement Learning by Solving Derived Non-Parametric MDPs. *International Conference on Learning Representations (ICLR), 2021.* [Spotlight Paper: 114/2997=3.8% of submissions]

17. Vincent Cartillier, Zhile Ren, Neha Jain, Stefan Lee, Irfan Essa, Dhruv Batra. Semantic MapNet: Building Allocentric SemanticMaps and Representations from Egocentric Views. *AAAI Conference on Artificial Intelligence (AAAI), 2021.*

18. Saurabh Satish Desai, Stefan Lee. Auxiliary Tasks for Efficient Learning of Point-Goal Navigation. *Winter Conference on Applications of Computer Vision (WACV), 2021.*

19. Peter Anderson, Ayush Shrivastava, Joanne Truong, Arjun Majumdar, Devi Parikh, Dhruv Batra, Stefan Lee. Sim-to-Real Transfer for Vision-and-Language Navigation. *Conference on Robot Learning (CoRL), 2020.*

20. Samyak Datta, Oleksandr Maksymets, Judy Hoffman, Stefan Lee, Dhruv Batra, Devi Parikh. Integrating Egocentric Localization for More Realistic Point-Goal Navigation Agents. *Conference on Robot Learning (CoRL), 2020.*

21. Simon Stepputtis, Joe Campbell, Mariano Phiellipp, Stefan Lee, Chitta Baral, Heni Ben Amor. Language-Conditioned Imitation Learning for Robot Manipulation Tasks. *Neural Information Processing Systems (NeurIPS), 2020.* [Spotlight Paper: 280/9467=3.0% of submissions]

22. Michael Cogswell, Jiasen Lu, Rishabh Jain, Stefan Lee, Devi Parikh, Dhruv Batra. Dialog without Dialog Data: Learning Visual Dialog Agents from VQA Data. *Neural Information Processing Systems (NeurIPS), 2020.*

23. Meera Hahn, Jacob Krantz, Dhruv Batra, Devi Parikh, James M. Rehg, Stefan Lee, Peter Anderson. Where Are You? Localization from Embodied Dialog. *Empirical Methods in Natural Language Processing (EMNLP), 2020.*

24. Yilin Yang, Longyue Wang, Shuming Shi, Prasad Tadepalli, Stefan Lee, Zhaopeng Tu. On the Sub-Layer Functionalities of Transformer Decoder. *Findings of Empirical Methods in Natural Language Processing (EMNLP), 2020.*

25. Jacob Krantz, Erik Wijmans, Arjun Majundar, Dhruv Batra, Stefan Lee. Beyond the Nav-Graph: Vision and Language Navigation in Continuous Environments. *European Conference on Computer Vision (ECCV), 2020.*

26. Arjun Majumdar, Ayush Shrivastava, Stefan Lee, Peter Anderson, Devi Parikh, Dhruv Batra. Improving Vision-and-Language Navigation with Image-Text Pairs from the Web. *European Conference on Computer Vision (ECCV), 2020.* [Spotlight Paper: 161/5025=3.2% of submissions]

27. Abhishek Kadian, Joanne Truong, Aaron Gokaslan, Alexander Clegg, Erik Wijmans, Stefan Lee, Manolis Savva, Sonia Chernova, Dhruv Batra. Are We Making Real Progress in Simulated Environments? Measuring the Sim2Real Gap in Embodied Visual Navigation. *Intelligent Robots and Systems (IROS) / Robotics and Automation Letters (RA-L), 2020.*

28. Jiasen Lu*, Vedanuj Goswami*, Marcus Rohrbach, Devi Parikh, Stefan Lee. 12-in-1: Multi-Task Vision and Language Representation Learning. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020.*

29. Erik Wijmans, Abhishek Kadian, Ari Morcos, Stefan Lee, Irfan Essa, Devi Parikh, Manolis Savva, Dhruv Batra. DD-PPO: Learning Near-Perfect PointGoal Navigators from 2.5 Billion Frames. *International Conference on Learning Representation (ICLR), 2020.*

30. Arijit Ray, Karan Sikka, Ajay Divakaran, Stefan Lee, Giedrius Burachas. Sunny and Dark Outside?! Improving Answer Consistency in VQA through Entailed Question Generation. *Empirical Methods in Natural Language Processing (EMNLP), 2019.*

31. Jiasen Lu, Dhruv Batra, Devi Parikh, Stefan Lee. ViLBERT: Pretraining Task-Agnostic Visiolinguistic Representations for Vision-and-Language Tasks. *Neural Information Processing Systems (NeurIPS), 2019.*

32. Peter Anderson*, Ayush Shrivastava*, Devi Parikh, Dhruv Batra, Stefan Lee. Chasing Ghosts: Instruction Following as Bayesian State Tracking. *Neural Information Processing Systems (NeurIPS), 2019.*

33. Harsh Agrawal*, Karan Desai*, Xinlei Chen, Rishabh Jain, Dhruv Batra, Devi Parikh, Stefan Lee, Peter Anderson. nocaps: novel object captioning at scale. *IEEE International Conference on Computer Vision (ICCV), 2019.*

34. Ramprasaath R. Selvaraju, Stefan Lee, Yilin Shen, Hongxia Jin, Dhruv Batra, and Devi Parikh. Taking a HINT: Leveraging Explanations to Make Vision and Language Models More Grounded. *IEEE International Conference on Computer Vision (ICCV), 2019.*

35. Ashwin Kalyan, Peter Anderson, Stefan Lee, Dhruv Batra. Trainable Decoding of Sets of Sequences for Neural Sequence Models. *International Conference on Machine Learning (ICML), 2019.*

36. Yash Goyal, Ziyan Wu, Jan Ernst, Dhruv Batra, Devi Parikh, Stefan Lee. Counterfactual Visual Explanations. *International Conference on Machine Learning (ICML), 2019.*

37. Ramakrishna Vedantam, Karan Desai, Stefan Lee, Marcus Rohrbach, Dhruv Batra, and Devi Parikh. Probabilistic Neural-symbolic Models for Interpretable Visual Question Answering. International Conference on Machine Learning (ICML), 2019.

38. Erik Wijmans*, Samyak Datta*, Oleksandr Maksymets*, Abhishek Das, Georgia Gkioxari, Stefan Lee, Irfan Essa, Devi Parikh, Dhruv Batra. Embodied Question Answering in Photorealistic Environments with Point Cloud Perception. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. **[Oral Paper]**

39. Huda Alamri and Vincent Cartillier and Abhishek Das and Jue Wang and Stefan Lee and Peter Anderson and Irfan Essa and Devi Parikh and Dhruv Batra and Anoop Cherian and Tim K. Marks and Chiori Hori. Audio-Visual Scene-Aware Dialog. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

40. Sainandan Ramakrishnan, Aishwarya Agrawal, Stefan Lee. Overcoming Language Priors in Visual Question Answering with Adversarial Regularization. *Neural Information Processing Systems (NeurIPS)*, 2018.

41. Abhishek Das, Georgia Gkioxari, Stefan Lee, Devi Parikh, Dhruv Batra. Neural Modular Control for Embodied Question Answering. *Conference on Robot Learning (CoRL)*, 2018.

42. Jianwei Yang*, Jiasen Lu*, Stefan Lee, Dhruv Batra, Devi Parikh. Learning to Ask Questions to Learn Visual Recognition. *Conference on Robot Learning (CoRL)*, 2018. **[Oral Paper] [Best Presentation Finalist]**

43. Ramprasaath R. Selvaraju, Prithvijit Chattopadhyay, Mohamed Elhoseiny, Tilak Sharma, Dhruv Batra, Devi Parikh, Stefan Lee. Choose Your Neuron: Incorporating Domain Knowledge through Neuron Importance. *European Conference on Computer Vision (ECCV)*, 2018.

44. Jianwei Yang*, Jiasen Lu*, Stefan Lee, Dhruv Batra, Devi Parikh. Graph R-CNN for Scene Graph Generation. *European Conference on Computer Vision (ECCV)*, 2018.

45. Ashwin K Vijayakumar, Stefan Lee, Anitha Kannan, Dhruv Batra. Learn From Your Neighbor: Learning Multi-Modal Distributions from Sparse Annotation. *International Conference on Machine Learning (ICML)*, 2018. **[Oral Paper – Long Talk]**.

46. Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, Dhruv Batra. Embodied Question Answering. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. **[Oral Paper]**.

47. Ashwin Vijayakumar, Michael Cogswell, Ramprasath Selvaraju, Qing Sun, Stefan Lee, David Crandall, and Dhruv Batra. Diverse Beam Search for Improved Description of Complex Scenes. *AAAI Conference on Artificial Intelligence (AAAI)*, 2018.

48. Abhishek Das*, Satwik Kottur*, José M.F. Moura, Stefan Lee, and Dhruv Batra. Learning Cooperative Visual Dialog Agents with Deep Reinforcement Learning. *IEEE International Conference on Computer Vision (ICCV)*, 2017. **[Oral Paper: 45/3220 = 1.4%]**.

49. Satwik Kottur, José M.F. Moura, Stefan Lee, and Dhruv Batra. Natural Language Does Not Emerge 'Naturally' in Multi-Agent Dialog. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017. **[Best Short Paper: 4/1500 = 0.26%]**.

50. Aroma Mahendru*, Viraj Prabhu*, Akrit Mohapatra*, Dhruv Batra, and Stefan Lee. The Promise of Premise: Harnessing Question Premises in Visual Question Answering. *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2017.

51. Viraj Prabhu, Prithvijit Chattopadhyay, Deshraj Yadav, Arjun Chandrasekaran, Abhishek Das, Stefan Lee, Dhruv Batra, and Devi Parikh. Evaluating Visual Dialog Agents via Cooperative Human-AI Games. *Proceedings of the Fifth AAAI Conference on Human Computation and Crowdsourcing (HCOMP)*, 2017.

52. Qing Sun, Stefan Lee, and Dhruv Batra. Bidirectional Beam Search: Forward-Backward Inference in Neural Sequence Models for Fill-in-the-Blank Image Captioning. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

53. Stefan Lee, Senthil Purushwalkam, Michael Cogswell,Viresh Ranjan, David J. Crandall, and Dhruv Batra. Stochastic Multiple Choice Learning for Training Diverse Deep Ensembles. *Neural Information Processing Systems (NeurIPS)*, 2016.

54. Sven Bambach, Stefan Lee, David Crandall, Chen Yu, Lending A Hand: Detecting Hands and Recognizing Activities in Complex Egocentric Interactions. *IEEE International Conference on Computer Vision (ICCV)*, 2015.

55. Stefan Lee, Nicolas Maisonneuve, David Crandall, Josef Sivic, Alexei A. Efros. Linking Past to Present: Discovering Style in Two Centuries of Architecture. *IEEE International Conference on Computational Photography (ICCP)*, 2015.

56. Stefan Lee, Haipeng Zhang, David Crandall. Predicting Geo-informative Attributes in Large-scale Image Collections using Convolutional Neural Networks. *IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2015.

57. Stefan Lee, Sven Bambach, David Crandall, John Franchak, and Chen Yu. This Hand Is My Hand: A Probabilistic Approach to Hand Disambiguation in Egocentric Video. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Workshop on Egocentric Vision,* 2014. [**Best Paper: 1/13 = 7.6%**].

58. Stefan Lee, Jerome Mitchell, David Crandall, and Geoffery Fox. Estimating Bedrock and Surface Layer Boundaries And Confidence Intervals In Ice Sheet Radar Imagery Using MCMC. *International Conference on Image Processing (ICIP),* 2014.

Journals

59. Drew Penney, Bin Li, Jaroslaw J. Sydir, LizhongChen, Charlie Tai, Stefan Lee, Eoin Walsh, Thomas Long. PROMPT: Learning dynamic resource allocation policies for network applications. *ACM Future Generation Computer Systems, Volume 145. 2023*

60. S. Solanki, S. Lee, A. Jebakumar, J. Lum, M. Hamidi-Haines, C. Denison, M. Sundheim, K. Schauer, P. Stevenson, J. Hintzman, E. Torniainen. Machine learning for predicting microfluidic droplet generation properties. *Computers & Fluids, Volume 247, 2022.*

61. Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav,Stefan Lee, José M. F. Moura, Devi Parikh, and DhruvBatra. Visual Dialog. *Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2018.

Book Chapters

62. David J. Crandall, Yunpeng Li, Stefan Lee, and Daniel P. Huttenlocher. Recognizing Landmarks in Large-Scale Social Image Collections. Large-Scale Visual Geo-Localization. Ed. Amir R. Zamir, Asaad Hakeem, Luc Van Gool, Mubarak Shah, Richard Szeliski. Springer, 2016.

Extended Abstracts / Technical Reports

63. Sam Greydanus, Stefan Lee, Alan Fern. Piecewise-constant Neural ODEs. arXiv, 2021.

64. Amit Raj, Cusuh Ham, Huda Alamri, Vincent Cartillier, Stefan Lee, James Hays. *Compositional Generation of Images. NeurIPS Workshop on Visually-Grounded Interaction and Language,* 2017.

65. Ashwin K Vijayakumar, Michael Cogswell, Ramprasath R. Selvaraju, Qing Sun, Stefan Lee, David Crandall, Dhruv Batra. Diverse Beam Search: Decoding Diverse Solutions from Neural Sequence Models. arXiv, 2016.

66. Sven Bambach, Stefan Lee, David Crandall, John Franchak, Chen Yu. Tracking Hands of Interacting People in Egocentric Video. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Workshop on Observing and Understanding Hands in Action*, 2015.

67. Stefan Lee, Senthil Purushwalkam, Michael Cogswell, David J. Crandall, Dhruv Batra. Why M Heads are Better than One: Training a Diverse Ensemble of Deep Networks. *arXiv,* 2015.

68. Stefan Lee and David Crandall. Learning to Identify Local Floral with Human Feedback. *IEEE Conference on Computer Vision and Pattern Recognition, Workshop on Computer Vision and Human Computation,* 2014.

TEACHING

Oregon State University
CS535 - Deep Learning	**Spring 2022,2023**
CS434 - Machine Learning and Data Science	**Spring 2021, Fall 2022,2023**
CS539 - Natural Language Processing with Deep Learning	**Winter 2020,2022**
CS539 - Embodied AI	**Fall 2019**

Georgia Institute of Technology
CS8903 - Special Problems	**Fall 2017 - Spring 2019**

Virginia Tech
ECE5424 - Introduction to Machine Learning	**Fall 2016**

Indiana University
B659 - Image Processing and Recognition (Assistant Instructor)	**Spring 2015**
I399 - Research Methods for Informatics and Computing (Mentor)	**Fall 2013**

Invited Talks

Telling Embodied Agents What To Do (And The Agents Sometimes Doing It In Reality)
Invited Speaker – Stanford University NLP Seminar Series   **2021**

Learning Transferable Visiolinguistic Representations
Invited Speaker – Oregon State AI Seminar Series   **2021**
Invited Speaker – Arizona State AI Seminar   **2021**
Invited Speaker – Faceook AI Research   **2020**

Exploring the Sim-2-Real Gap in Vision-and-Language Navigation
Invited Speaker – Lantern Workshop @ COLING   **2020**

Questions at the Intersection of Emergent Communication and Natural Language
Invited Speaker – Emergent Communication Workshop @ NeurIPS   **2019**

Training Embodied Agents in Semantically and Perceptually Rich Simulations
DARPA 60th Anniversary (D60) - DARPA Riser Plenary   **2018**

Towards Goal-Driven,Visually Grounded Dialog Agents
Carnegie Mellon University - VASC Seminar   **2018**
Facebook AI Research (FAIR) - AI Seminar   **2018**
University of Trento - DISI Seminar   **2017**
Georgia Tech - Deep Learning Course Guest Speaker   **2017**

Training Diverse Deep Ensembles
Diversity Meets Deep Networks - CVPR Tutorial   **2016**

Media Coverage

AI with no visual sensors generates an internal map to finds its way – NewScientist 2023

Facebook proposes 3D navigation task for training autonomous robots – VentureBeat 2020

Facebook, Georgia Tech & OSU ViLBERT Achieves SOTA on Vision-and-Language Tasks – Medium 2019

Facebook helped create an AI scavenger hunt that could lead to the first useful home robots - MIT Technology Review 2018

How A Virtual Scavenger Hunt Could Train Robots To Find Things In Your Home - FastCompany 2018

Facebook is training AI to answer questions like humans do – Digital Journal 2018

Research Scientist, Assistant Professor Represent IC in DARPA Risers Event - ML@GT Blog 2018

What is Graph R-CNN? - ML@GT Blog 2018

Choose Your Neuron: Incorporating Domain Knowledge through Neuron-Importance - ML@GT Blog 2018

Embodied Question Answering - ML@GT Blog 2018

Service

Served as area chair for
 Computer Vision and Pattern Recognition (CVPR), 2021-2022

 Neural Information Processing Systems (NeurIPS) 2019-2021,2023

 International Conference on Machine Learning (ICML) 2020

 AAAI Conference on Artificial Intelligence (AAAI) 2019

Regularly review or serve on the program committee for
Computer Vision and Pattern Recognition (CVPR)

European Conference on Computer Vision (ECCV)

International Conference on Computer Vision (ICCV)

Neural Information Processing Systems (NeurIPS)

International Conference on Learning Representations (ICLR)

Social Network Analysis and Mining (Springer Journal)

Image and Vision Computing (Elsevier Journal)

Workshop / Tutorial Organization
Embodied AI Workshop @ CVRP 2020-2022

Embodied Agents Workshop @ CVRP 2019

Visually-Grounded Interaction and Language (ViGIL) Workshop @ NeurIPS 2017 / 2018 / 2019 & NAACL 2021

Visual Learning and Embodied Agents in Simulation Environments (VLEASE) Workshop @ ECCV 2018

Shortcomings in Vision and Language (SiVL) Workshop @ ECCV 2018 / NAACL 2019

VQA Challenge and Visual Dialog Workshop @ CVPR 2018 - 2019

Diversity Meets Deep Networks Tutorial @ CVPR 2016

STUDENTS SUPERVISED

Current Students

Zijaio Yang, Oregon State University	[PhD]
Xiangxi Shi, Oregon State University	[PhD]
Skand Peri, Oregon State University (co-advised with Fuxin Li)	[PhD]
Eric Slyman, Oregon State University	[PhD]
Abhinav Jain, Oregon State University (co-advised with Cindy Grimm)	[PhD]
Akhil Perincherry, Oregon State University	[PhD]
Joe Nguyen, Oregon State University	[MS]

Alumni

Jacob Krantz, Oregon State University	[PhD, 2023]
Drew Penney, Oregon State University (co-advised with Lizhong Chen)	[PhD, 2023]
Yilin Yang, Oregon State University (co-advised with Prasad Tadepalli)	[PhD, 2022]
Saurabh Desai, Oregon State University	[MS, 2021]
Shridhar Mishra, Oregon State University	[MS, 2021]