

TEJAS SRINIVASAN

tejas.srinivasan@usc.edu ♦ [Website](#) ♦ [Google Scholar](#)

RESEARCH INTERESTS

My research interests lie in **Multimodal Learning**, particularly **Vision-and-Language**, and how visual commonsense can ground language in the real world. I am also interested in **Continual Learning** in real-world multimodal settings.

EDUCATION

2021 - Curr.	Ph.D. in Computer Science, University of Southern California Advisor: Prof. Jesse Thomason	GPA: 4.0/4.0
2018 - 2020	M.Sc. in Language Technologies, Carnegie Mellon University Advisors: Prof. Florian Metze, Prof. Louis-Phillippe Morency, Prof. Yonatan Bisk	GPA: 3.87/4.0
2014 - 2018	B.Tech. Mechanical Engineering, Indian Institute of Technology, Bombay	GPA: 9.15/10.0

INDUSTRY EXPERIENCE

Summer 2021	Microsoft Research (MSR)	NLP Research Intern
2020 - 2021	A.I. Foundation	NLP Research Scientist

SELECTED PUBLICATIONS

- 2022 [Curriculum Learning for Data-Efficient Vision-Language Alignment](#)
Tejas Srinivasan, Xiang Ren, Jesse Thomason
arXiv preprint
We **align pre-trained vision and language models** to each other to achieve better cross-modal alignment than CLIP, using $< 1\%$ as much image-text data!
- 2022 [CLiMB: A Continual Learning Benchmark for Vision-and-Language Tasks](#)
Tejas Srinivasan, Ting-Yun Chang, Leticia Pinto Alva, Georgios, Mohammad Rostami, Jesse Thomason
NeurIPS 2022 Datasets and Benchmarks Track
We introduce CLiMB, a benchmark to study **continual learning in multimodal settings**.
- 2022 [Worst of Both Worlds: Biases Compound in Pre-trained Vision-and-Language Models](#)
Tejas Srinivasan, Yonatan Bisk
Workshop on Gender Bias in Natural Language Processing, NAACL 2022
We analyze how **multimodal LMs exhibit gender biases** learned from each modality.
- 2020 [Looking Enhances Listening: Recovering Missing Speech Using Images](#)
Tejas Srinivasan, Ramon Sanabria, Florian Metze
International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2020
We utilize visual context to make ASR models more robust to audio corruption.
- 2019 [Structured Fusion Networks for Dialog](#)
Shikib Mehri, Tejas Srinivasan*, Maxine Eskenazi*
Special Interest Group on Discourse and Dialogue (SIGDIAL) 2019
Best Paper Award
We incorporate structure from traditional dialog systems into neural dialog models.

RESEARCH SKILLS

NLP/DL TRAINING	PyTorch, TensorFlow, Huggingface Transformers, adapter-transformers, Weights&Biases
COMPUTING TOOLS	NumPy, SciPy, Pandas, Git, L ^A T _E X