Sarath Shekkizhar, Ph.D.

shekkizh@usc.edu · shekkizh.com · shekkizh · shekkizh

Education

Ph.D. in Electrical Engineering Aug 2017 - May 2023
University of Southern California, Los Angeles, CA GPA: 3.93
Advisor: Antonio Ortega

M.S. in Computer Science Aug 2017 - May 2022
University of Southern California, Los Angeles, CA GPA: 4.0

M.S. in Electrical Engineering (Computer Vision, Machine Learning) Aug 2012 - Dec 2013
University of Southern California, Los Angeles, CA GPA: 3.86

B.Tech. in Electronics and Communication July 2008 - June 2012
National Institute of Technology, Tiruchirappalli, India GPA: 9.12

Work Experience

Staff Research Scientist, *Salesforce*, San Francisco, CA Oct 2024 - Present
Working on foundational research on (multi) agentic systems and LLM training for improved reasoning and alignment. I continue to spend some of my research time on aspects of voice AI and agentic design for applied research.

Member of Technical Staff, *Tenyx (Acq. by Salesforce)*, Los Altos, CA June 2023 - October 2024
Part of the founding team at Tenyx building Voice AI for customer support. Was primarily focused on research and algorithms for various aspects of voice agents. Key accomplishments include research on continual learning, building TenyxChat series of models, and geometric characterization of LLMs. Was also involved in product development, particularly in endpointing, audio disambiguation, and agent governance.

Research Intern, *Google*, Sunnyvale, CA Sep 2022 - Dec 2022
Worked on understanding the impact of input data used in training graph models and scalable sampling approaches to improve semi-supervised graph learning. Preliminary experiments with proposed graph learning showed 3x increased recall in abuse detection. Host: Mohamed Farghal, Animesh Nandi, Behavior Protections, Counter-Abuse Technology.

Software Engineer 2, *KLA Tencor*, Milpitas, CA Mar 2014 - Oct 2016
Designed and developed tools to classify and visualize defect modulations for Process Window Qualification in wafer fabrication. Also, implemented and co-owned components for analysis and classification using decision trees and random forests.

Publications

[1] S. Shekkizhar, P. Jwalapuram, Z. Ke, V. Patil, J. Wang, S. Bansal, T. H. Nguyen, D. Jiang, D. Liu, R. Zhao, W. Lu, G. Carenini, M. Bansal, S. Savarese, S. Joty. Why Do Agent Teams Fail Across Boundaries? A Failure Taxonomy for Agent-to-Agent Systems with Private States. *arXiv Preprints*, 2026.

[2] Z. Ke, V. Patil, H. Shi, Y. Li, Y. Liu, S. Shekkizhar, A. Koul, J. Wang, X. P. Nguyen, S. Yavuz, M. Bansal, S. Joty. EVOHARNESSBENCH: Can Your Agents Keep Pace with an Evolving Harness?. *arXiv Preprints*, 2026.

[3] R. Cosentino, S. Shekkizhar, A. Earle, S. Savarese. Counterparty Modeling is Not Strategy: The Limits of LLM Negotiators. *arXiv Preprints*, 2026.

[4] S Shekkizhar. Beyond the Assistant Turn: User Turn Generation as a Probe of Interaction Awareness in Language Models. *arXiv Preprints*, 2026.
[5] S Shekkizhar, A Earle. Interaction Theater: A case of LLM Agents Interacting at Scale. *arXiv Preprints*, 2026.
[6] S Shekkizhar, R Cosentino, A Earle, S Savarese. Echoing: Identity Failures when LLM Agents Talk to Each Other. *arXiv Preprints*, 2025.
[7] R Cosentino, S Shekkizhar, A Earle. Convergence dynamics of Agent-to-Agent Interactions with Misaligned objectives. *arXiv Preprints*, 2025.
[8] S Shekkizhar, R Cosentino. AGI Is Coming... Right After AI Learns to Play Wordle. *arXiv Preprints*, 2025.
[9] A. Gulati, X. Dong, C. Hurtado, S. Shekkizhar, S. Swayamdipta, A. Ortega. Out-of-Distribution Detection through Soft Clustering with Non-Negative Kernel Regression. *Findings of the Association for Computational Linguistics: EMNLP*, 2024.
[10] R Cosentino, S Shekkizhar. Reasoning in Large Language Models: A Geometric Perspective. *arXiv Preprints*, 2024.
[11] R Balestriero, R Cosentino, S Shekkizhar. Characterizing Large Language Model Geometry Solves Toxicity Detection and Generation. *International Conference on Machine Learning (ICML)*, 2024.
[12] P. Das, S. Shekkizhar, A. Ortega. Towards a geometric understanding of Spatio Temporal Graph Convolution Networks. *IEEE Open Journal of Signal Processing*, 2024.
[13] S. Shekkizhar, N. Bulut, M. Farghal, S. Tavakkol, M. Bateni, A. Nandi. Data Sampling using Locality Sensitive Hashing for Large Scale Graph Learning. *Mining and Learning with Graphs, Knowledge Discovery and Data Mining (KDD)*, 2023.
[14] S. Shekkizhar, A. Ortega. A data-driven graph framework for geometric understanding of deep learning. *Graph Signal Processing Workshop 2023*, 2023.
[15] C. Hurtado, S. Shekkizhar, J. Ruiz-Hidalgo, A. Ortega. Study of Manifold Geometry using Multi-scale Non-Negative Kernel Graphs. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
[16] R. Cosentino, S. Shekkizhar, M. Soltanolkotabi, S. Avestimehr, A. Ortega. The geometry of self-supervised learning models and its impact on Transfer learning. *arXiv Preprints*, 2022.
[17] S. Shekkizhar, A. Ortega. NNK-Means: Data summarization using dictionary learning with non-negative kernel regression. *IEEE 30th European Signal Processing Conference (EUSIPCO)*, 2022.
[18] D. Bonnet, A. Ortega, J.Ruiz-Hidalgo, S.Shekkizhar. Channel redundancy and overlap in convolutional neural networks with Channel-wise NNK graphs. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
[19] D. Bonnet, A. Ortega, J.Ruiz-Hidalgo, S.Shekkizhar. Channel-Wise Early Stopping without a ValidationSet via NNK Polytope Interpolation. *Asia Pacific Signal and Information Processing Association (APSIPA)*, 2021.
[20] S. Shekkizhar, A. Ortega. Model selection and explainability in neural networks using a polytope interpolation framework. *Asilomar Conference on Signals, Systems, and Computers*, 2021.
[21] S. Shekkizhar, A. Ortega. Revisiting local neighborhood methods in machine learning. *IEEE Data Science and Learning Workshop (DSLW)*, 2021.
[22] S. Shekkizhar, A. Ortega. Efficient graph construction for image representation. *IEEE International Conference on Image Processing (ICIP)*, 2020. **Best student paper**
[23] K. Nonaka, S. Shekkizhar, A. Ortega. Graph-based Deep Learning Analysis and Instance Selection. *IEEE International Workshop on Multimedia Signal Processing (MMSP)*, 2020.
[24] S. Shekkizhar, A. Ortega. DeepNNK: Explaining deep models and their generalization using polytope interpolation. *arXiv Preprints*, 2020.
[25] S. Shekkizhar, A. Ortega. Graph Construction from Data by Non-Negative Kernel Regression. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020.
[26] S. Shekkizhar, A. Ortega. Neighborhood and Graph Constructions using Non-Negative Kernel Regres-

sion. *arXiv*, 2019.
[27] S. Deivalakshmi, S. Shekkizhar, P. Palanisamy. Detection and removal of Salt and Pepper noise in images by improved median filter. *IEEE Recent Advances in Intelligent Computational Systems*, 2011.

Patents

[1] S. Shekkizhar, R. Cosentino. Knowledge base for voice large language model applications. US63752613, Filed: January 2025. *Provisional.*
[2] R. Cosentino, S. Shekkizhar. Gradient-free optimization of large language models. US63752618, Filed: January 2025. *Provisional.*
[3] R. Cosentino, D. Kalajdzievski, S. Shekkizhar, A. Earle. Machine learning model compression. US18905761, Filed: October 2024. *Provisional.*
[4] D. Kalajdzievski, R. Cosentino, S. Shekkizhar, A. Earle. Training a target activation sparsity in a neural network. US18802235, Filed: August 2024. *Pending.*
[5] S. Shekkizhar, A. Earle. Domain aware large language model governance. US18745562, Filed: June 2024. *Granted.*
[6] R. Cosentino, S. Shekkizhar, A. Earle, D. Kalajdzievski, J. Weissenberger, I. Arel. Fine-tuning machine learning models while retraining accumulated knowledge. US18496698, Filed: October 2023. *Pending.*
[7] S. Shekkizhar, N. Bulut, M. Farghal, S. Tavakkol, M. Bateni, A. Nandi. Data sampling using Locality Sensitive Hashing for large scale graph learning. US63517869, Filed: August 2023. *Provisional.*
[8] M. Plihal, E. Soltanmohammadi, S. Paramasivam, S. Ravu, A. Jain, S. Shekkizhar, P. Uppaluri. Optimizing training sets used for setting up inspection-related algorithms. US10267748, Filed: April 2019. *Granted.*

Awards & Honors

- IEEE Rising Star in Signal Processing - ICASSP 2023
- IEEE Best Student Paper Award - ICIP 2020
- Ming-Hsieh Ph.D. Scholar Finalist 2022-23

Academic Activities

- Reviewer: IEEE Journals (JSAIT, TSIPN, SPL, TNNLS)
- Reviewer: Conferences (ICASSP, ICLR, NeurIPS, LoG, ICML, COLM)
- Area Chair: NeurIPS