Fenil Denish Bardoliya

1-602-574-0317 fenildbardoliya@gmail.com linkedin.com/in/fenil-bardoliya github.com/fenil-b fenil-b.github.io

Summary

Engineer–researcher building robust multimodal AI systems in the domain of vision-language modeling, agentic reasoning, and applied perception. My work spans multimodal evaluation and grounding, responsible and reliable AI, distributed post-training and unlearning of large models, and video understanding, with a focus on turning foundation models into scalable, trustworthy systems for real-world deployment.

Education

Master of Science in Computer Science **Aug 2023 – May 2025**

Arizona State University, Tempe, AZ, USA

Coursework: Digital Video Processing, Natural Language Processing, Planning and Learning Methods in AI, Biocomputing, Information Assurance and Security, Image Analytics & Informatics, Data Mining, Statistical Machine Learning, Frontier Topics in GenAI

Bachelor of Engineering in Computer Science **Aug 2019 – Jul 2023**

Birla Institute of Technology and Science Pilani

Coursework: Probability and Statistics, Object Oriented Programming, Database Systems, Data Structures and Algorithms, Operating Systems, Design and Analysis of Algorithms, Network Programming, Software Development for Portable Devices, Reinforcement Learning, Machine Learning, Image Processing, Computer Vision

Publications

[1] Mike Zhou, Fenil Bardoliya, Vivek Gupta, and Dan Roth. "Program-of-Thought Reveals LLM Abstraction Ceilings". In: *Findings of the Association for Computational Linguistics: EACL 2026*. Association for Computational Linguistics, 2026, pp. 4911–4919.

[2] Tejas Anvekar, Fenil Bardoliya, Pavan K. Turaga, Chitta Baral, and Vivek Gupta. "The Perceptual Observatory Characterizing Robustness and Grounding in MLLMs". In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. 2026, pp. 1653–1663.

[3] Sebastian Martinez, Naman Ahuja, Fenil Bardoliya, Suparno Roy Chowdhury, Chris Bryan, and Vivek Gupta. "SPORTSQL: An Interactive System for Real-Time Sports Reasoning and Visualization". In: *Proceedings of The 14th International Joint Conference on Natural Language Processing and The 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics: System Demonstrations*. Association for Computational Linguistics, 2025, pp. 94–101.

[4] Naman Ahuja, Fenil Bardoliya, Chitta Baral, and Vivek Gupta. "Map&Make: Schema Guided Text to Table Generation". In: *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics, 2025, pp. 30249–30262.

[5] Aayush Atul Verma, Amir Saeidi, Shamanthak Hegde, Ajay Therala, Fenil Denish Bardoliya, Nagaraju Machavarapu, Shri Ajay Kumar Ravindhiran, Srija Malyala, Agneet Chatterjee, Yezhou Yang, and Chitta Baral. "Evaluating Multimodal Large Language Models Across Distribution Shifts and Augmentations". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*. 2024, pp. 5314–5324.

[6] Kip V Hodges, Fenil Bardoliya, Jnaneshwar Das, Lauren Gold, Jack Hael, Evan Kohl, Robert LiKamWa, Tracey Lott, and Harrison H Schmitt. "Improving Planetary Field Geology Research and Training Using High-Spatial-Resolution, Immersive Virtual Environments". In: *AGU Fall Meeting Abstracts*. Vol. 2024. 2138. 2024, IN11C–2138.

Experience

Complex Data Reasoning & Analysis Lab (CoRAL) **Aug 2024 – Present**

NLP Researcher *Tempe, AZ, USA*

- Built a scalable evaluation framework for MLLMs that isolates perceptual, spatial, grounding, and fairness failures under 30+ perturbations across 62K+ samples (scalable to 1M+), supporting deployment-facing reliability analysis for foundational VLMs
- Designed Map&Make, a schema-guided and agentic text-to-table pipeline for high-fidelity structured extraction from long-form narratives; evaluated frontier LLMs with 5+ structural and semantic metrics.
- Implemented distributed post-training with LoRA, DPO, and GRPO on A100, H100, and H200 clusters, improving GSM8K and MATH by 20-40% across multiple open-weight LLMs while balancing compute efficiency and stability.
- Co-developed SPORTSQL, an interactive NL-to-SQL and visualization system over live English Premier League data; contributed 1,793 benchmark queries and reached up to 80% exact match and 94% LLM-as-judge accuracy.
- Built annotation and quality-control protocols for human and LLM-assisted evaluation, improving reproducibility, failure analysis, and benchmarking rigor.

Distributed Robotic Exploration and Mapping Systems (DREAMS) Lab **March 2024 – Jul 2024**

Computer Vision Research Aide *Tempe, AZ, USA*

- Developed 3D reconstruction and scene-modeling pipelines for geological formations using Gaussian Splatting and NeRF, emphasizing geometry, viewpoint consistency, and photorealism for simulation-oriented visual analysis.

- Recreated Apollo lunar landing sites for immersive simulation environments, strengthening experience in neural rendering, spatial fidelity, and world-model-like scene representation.

Samsung Semiconductor India Research **Jan 2023 – Jul 2023**

Assistant Engineer *Bengaluru, India*

- Built automation pipelines for network log capture and analysis, and developed an LSTM-based anomaly detector for SIP and IMS traces over VoLTE, reaching about 80% classification accuracy across 25+ failure modes.

Projects

Vermillion: Text-to-Video (T2V) Model **Aug 2024 – Dec 2024**

- Improved the VAE bottleneck in ASU's in-house text-to-video model to study latent compression, reconstruction quality, and temporal coherence trade-offs, yielding sharper and more stable generated videos.

Exploring Unlearning in State Space Models **Aug 2024 – Dec 2024**

- Built an evaluation pipeline to study targeted unlearning in State Space Models and Transformers, focusing on privacy, selective forgetting, and retention of useful capabilities after model updates. Implemented Gradient Ascent and Gradient Ascent + Mismatch across OPT-1.3B, Pythia-1.4B, and Mamba-1.4B, and evaluated behavior using Perplexity, BLEU, ROUGE-L, and BLEURT.
- Identified a key systems trade-off: SSMs showed 15-20% lower catastrophic forgetting but 30% slower unlearning, highlighting their promise for controllable and compliant large-model adaptation.

Medical imaging, localization, segmentation, and self-supervision **Jan 2024 – May 2024**

- Built chest X-ray pipelines for classification, age estimation, semantic and instance segmentation, localization, anomaly detection, and clustering on MiniJSRT using ResNet18, VGG-UNet, and k-means; achieved 92% gender classification, Dice 0.83 semantic segmentation, and AP 0.751 localization.
- Extended the work to larger-scale medical-imaging baselines using ConvNeXt, Swin, InternImage, ViT-B, and I-JEPA across MIMIC, CheXpert, ChestX-ray14, VinDr-CXR, and ChestX-Det.

Dark Motions: Low-Light Video Enhancement **Aug 2023 – Dec 2023**

- Proposed a framework for extreme low-light video restoration that jointly addresses noise suppression, perceptual detail retention, and inter-frame stability.

Invited Talks

- Voxel51 Best of WACV 2026

Services

Reviewer - CVPR 2026 Workshop CogVL, ACL 2026 Workshop SURGeLLM, ICML 2026 Workshop AI4Science

Teaching Experience

Student Project Mentor, CSE 576: Natural Language Processing **Spring 2025**

Informal TA, CSE 576: Natural Language Processing **Spring 2025**

Grader, CSE 575: Statistical Machine Learning **Spring 2025**

Thesis Mentor **Fall 2024 – Spring 2025**

Student Project Mentor, CSE 576: Natural Language Processing **Fall 2024**

Skills

Programming Languages: Python, Kotlin, Java, C++, C, SQL, JavaScript, HTML, CSS, XML
Databases/Cloud: Firebase Firestone, Firebase Realtime Database, MySQL, MongoDB, MariaDB, AWS(EC2, S3, SageMaker, Bedrock), GCP, Azure(Azure ML, Data Lake), Kubernetes, TensorBoard, W&B
Libraries/Technologies/Frameworks/Tools: Scikit-learn, TensorFlow, Keras, PyTorch, HuggingFace, OpenCV, Albumentations, Pillow, Matplotlib, Numpy, Pandas, NLTK, Spacy, Networkx, REST API, Git, JUnit, Docker, Wireshark, AutoML
Concepts: Deep Learning, Machine Learning, LLMs, Computer Vision, Generative AI, Image Processing, Software Development, Android App Development, RAG, Prompt Engineering, LLM Agents, Diffusion/Flow models, RL Algorithms