

FENIL DENISH BARDOLIYA

📞 +1 (602) 574-0317 📩 fenilbardoliya@gmail.com 💬 linkedin.com/in/fenil-bardoliya 🐾 github.com/fenil-b 🌐 fenil-b.github.io

Summary

Engineer and Researcher working on Multimodal Machine Learning, with experience in Vision-Language Models, Structured Reasoning and Evaluation, Model Architecture, and Post-Training of Large Models across research and applied settings.

Education

Master of Science in Computer Science

Arizona State University, Tempe, AZ, USA

Aug 2023 – May 2025

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

Birla Institute of Technology and Science Pilani

Aug 2019 – Jul 2023

Publications

- *The Perceptual Observatory Characterizing Robustness and Grounding in MLLMs* (WACV 2026)
- *Map&Make: Schema Guided Text to Table Generation* (ACL 2025)
- *SPORTSQL: An Interactive System for Real-Time Sports Reasoning and Visualization* (AAACL-IJCNLP 2025 System Demo)
- *Evaluating Multimodal Large Language Models across Distribution Shifts and Augmentations* (CVPR Workshops 2024)

Experience

Complex Data Reasoning & Analysis Lab (CoRAL)

Aug 2024 – Present

NLP Researcher

Tempe, AZ, USA

- Designed an agentic framework for structured generation, enabling high-fidelity text-to-table and schema-aware reasoning, and evaluated frontier LLMs using 5+ structural and semantic metrics.
- Implemented distributed LoRA + DPO fine-tuning on A100/H100/H200 clusters, achieving 20–40% gains on GSM8K and MATH across multiple open-weight LLMs.
- Built a large-scale multimodal robustness evaluation framework, analyzing perceptual, spatial, grounding, and fairness failures under 30+ in-and-out distribution perturbations across 62K+ samples.
- Co-designed SPORTSQL, an interactive NL-to-SQL and visualization system for real-time EPL analytics; contributed 1,793 benchmark queries, achieving up to 80% Exact Match and ~94% LLM-as-Judge accuracy.
- Studied parameter-efficient post-training and distillation trade-offs to balance performance, stability, and memory footprint.
- Designed annotation guidelines and quality-control protocols for human and LLM-assisted evaluations.

Distributed Robotic Exploration and Mapping Systems (DREAMS) Lab

March 2024 – Jul 2024

Computer Vision Research Aide

Tempe, AZ, USA

- Conducted 3D reconstruction and scene modeling of geological rock formations using Gaussian Splatting and NeRF-based pipelines, enabling high-fidelity spatial visualization and analysis.
- Recreated Apollo lunar landing sites for immersive simulation in Dreamscape, focusing on geometric accuracy, photorealism, and viewpoint consistency across reconstructed scenes.

Samsung Semiconductor India Research

Jan 2023 – Jul 2023

Assistant Engineer

Bengaluru, India

- Built automation pipelines for network log capture and analysis; developed an LSTM-based anomaly detection system for SIP/IMS traces over VoLTE.
- Achieved ~80% classification accuracy across 25+ failure modes, including previously unseen error patterns.

Projects

Vermillion: Text-to-Video (T2V) Model

Aug 2024 – Dec 2024

- Contributed to Vermillion, ASU's in-house text-to-video (T2V) generation model, focusing on architectural improvements to the VAE bottleneck for enhanced spatiotemporal fidelity.
- Investigated latent compression, reconstruction quality, and temporal coherence trade-offs, leading to improved video sharpness and stability across generated sequences.

Exploring Unlearning in State Space Models

Aug 2024 – Dec 2024

- Compared machine unlearning methods for SSMs vs Transformers, analyzing privacy, forgetting, and adaptation across multiple datasets.
- Found SSMs exhibit 15–20% lower catastrophic forgetting but 30% slower unlearning than Transformers.

Dark Motions: Low-Light Video Enhancement

Aug 2023 – Dec 2023

- Developed a low-light video enhancement model targeting noise suppression, detail preservation, and temporal consistency under extreme illumination conditions.
- Evaluated enhancement quality using perceptual and signal-based metrics, with emphasis on flicker reduction and cross-frame stability.

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