

Ahnaf Murshid

murshid.ahnaf@gmail.com | linkedin.com/in/ahnaf-murshid | github.com/amurshid | ahnafmurshid.dev

EDUCATION

University of Houston <i>M.S. in Engineering Data Science and AI</i>	Houston, TX Aug 2026 – Expected May 2028
University of Houston <i>B.S. in Computer Science, Minor in Mathematics</i>	Houston, TX Aug 2024 – Aug 2026
Texas Tech University <i>Computer Science coursework – transferred to UH</i>	Lubbock, TX Aug 2022 – May 2024

EXPERIENCE

University of Houston–Clear Lake <i>Undergraduate Researcher, Robust Encrypted Traffic Classification</i>	Houston, TX Jun 2026 – Aug 2026
- First-authored a paper accepted at IEEE CCSS 2026 on robust transformer-based traffic classification - Diagnosed shortcut learning in a 132M-parameter BERT model: corrupting 16 header bytes cut macro-F1 from 0.98 to 0.35 - Built a training-time regularizer retaining 0.85 vs. 0.53 macro-F1 under full header loss, with 5× lower seed variance - Improved robustness to black-box and white-box adversarial attacks by 10–12 macro-F1 points at under 0.4% training cost - Ran 15 fine-tuning experiments with multi-seed runs, ablations, and bootstrap CIs; caught an unstable metric pre-publication	
University of Houston <i>Research Assistant, SZ3 Scientific Data Compression</i>	Houston, TX Jan 2025 – May 2025
- Achieved a 7× compression ratio at 97% data accuracy across 7+ large-scale weather and seismic datasets - Built automated testing pipelines benchmarking compression algorithms on HPC-scale data with PSNR and L2-norm metrics	
Texas Tech University <i>Student Assistant</i>	Lubbock, TX Jun 2023 – Jul 2024
- Managed academic records for a 500+ student department at 100% data integrity across 2+ semesters - Cut record retrieval time by restructuring and cleaning large datasets in Excel	

PROJECTS

Autonomous Rover – Camera-Only Localization <i>Visual Place Recognition, Deep Learning</i>	<i>PyTorch, DINOv2 ViT, TorchScript, ROS 2</i>
- Built a camera-only localizer – one frame in, map position out – at 0.25 m median error, 90% within 1 m - Fine-tuned a DINOv2 vision transformer on a self-collected 13,875-image day/night dataset - Cut error 62% under unfamiliar lighting (0.71 m → 0.27 m) after tracing four failed runs to a misleading validation split - Deployed to a Raspberry Pi, where one camera frame corrected a 5.66 m position error in 3 seconds - Added natural-language control: an LLM tool-calling agent with Whisper speech-to-text issuing navigation goals	
Lung Cancer CT Scan Classification Pipeline <i>Medical Imaging / Deep Learning</i>	<i>PyTorch, OpenCV, scikit-learn</i>
- Classified 5 lung CT categories at 81% test accuracy with a from-scratch PyTorch CNN trained on 1,535 images Per-class error analysis exposed what accuracy hid: benign, the rarest class, at 0% recall on all 18 test cases - Automated preprocessing for mixed JPG/PNG/RGBA sources: RGBA compositing, 224×224 resizing, stratified 70/15/15 split - Identified class weighting and a pretrained backbone as the fix, with macro-F1 replacing accuracy as the metric	

TECHNICAL SKILLS

Languages: Python, SQL, C++, C, Java, R, JavaScript/TypeScript

ML & AI: PyTorch, scikit-learn, Transformers (BERT, DINOv2 ViT), CNNs, transfer learning & fine-tuning, metric learning, data augmentation, adversarial robustness, model evaluation & experiment design, LLM tool calling

Data & Vision: Pandas, NumPy, OpenCV, Matplotlib, seaborn, statistical analysis (bootstrap confidence intervals)

Tools: Git, Linux, Docker, Kubernetes, AWS, Azure, Flask, ROS 2, TorchScript, Raspberry Pi