Fazle Rafsani

+1-602-558-6525 | frafsani@asu.edu | rafsani.github.io | linkedin.com/in/fazle-rafsani | github.com/Rafsani | Google Scholar
Tempe, AZ, USA (open to relocation) · F-1 student, work-authorized via CPT/OPT · Expected graduation Dec 2027

About

PhD candidate in Computer Science (GPA 4.00) working on deep learning, computer vision, and multimodal foundation models, with **7 peer-reviewed publications (3 first-author)** at ICCV, WACV, IEEE TMI (IF 12.4), and Nature Scientific Reports. Build end-to-end ML pipelines, data curation, multi-GPU training, evaluation, and reproducible benchmarking for vision and vision-language models, including DINOv2, CLIP/BioMedCLIP, diffusion models, and V-JEPA. Prior industry experience as an AI/ML software engineer at IQVIA shipping production Python and ML services on healthcare data. Seeking AI/ML engineering and research internships or full-time roles.

Education

- **Arizona State University** *Aug 2023 - Dec 2027 (Expected)*
  PhD in Computer Science *(MS in Computer Science awarded en route, Dec 2026)* Tempe, AZ, USA
  - **GPA: 4.00/4.00**
  - Focus: Deep Learning, Computer Vision, Foundation Models, Self-Supervised Learning, Medical Imaging
- **Bangladesh University of Engineering and Technology (BUET)** *May 2022*
  Bachelor of Science in Computer Science and Engineering Dhaka, Bangladesh

Experience

- **Wu Lab, Arizona State University** *January 2024 - Present*
  Graduate Research Associate Tempe, AZ
  - Build end-to-end deep learning pipelines at the ASU-Mayo Center for Innovative Imaging for 3D segmentation, anomaly detection, and multimodal classification, training on **10k+ 3D volumes** across **9** public and clinical datasets.
  - Adapt **foundation and generative models** (DINOv2, BioMedCLIP, denoising diffusion, V-JEPA) to 3D medical imaging, running **multi-GPU distributed training** on a SLURM cluster with experiment tracking and reproducible evaluation harnesses.
  - Published **5 papers in 18 months** at **ICCV, WACV, IEEE TMI, and Nature Scientific Reports**; released open-source implementations of two methods.
  - Mentor **2** junior graduate researchers on model training, code review, and experiment design.
- **IQVIA** *August 2022 - July 2023*
  Software Engineer, AI/ML Dhaka, Bangladesh
  - Built and shipped Python ML services for the **"Next Best Action"** pharmaceutical recommendation platform, serving **3 client brands** on **10000+** patient and prescriber records.
  - Developed the **recommendation** model and feature pipeline that improved **recommendation acceptance rate by 5%** over the prior heuristic baseline.
  - Owned data preprocessing and API layer (**Python, SQL, REST**) for the recommendation service; the product won the **2023 Stevie Award** for innovation in healthcare technology.
- **Arizona State University** *August 2023 - Present*
  Graduate Teaching Associate Tempe, AZ
  - Instructor of record for four sections of CSE 110 (Principles of Programming), teaching and evaluating **350+ students**.
- **ICT Cell, BUET** *December 2021 - August 2022*
  Software Developer Dhaka, Bangladesh
  - Built a university financial-oversight platform end-to-end using **Django, React.js, and PostgreSQL**, deployed for **100+** administrative users across **6** departments.

Selected Publications

C=Conference, J=Journal, A=Abstract; 3 first-author

[J.1] **Rafsani, F.**; Sheth, D.; Che, Y. et al. (2025). **Leveraging multi-modal foundation model image encoders to enhance brain MRI-based headache classification.** *Scientific Reports* 15, 33256 (2025). [IF 4.0]

[C.1] **Rafsani, F.**; Shah, J.; Chong, C. D.; Schwedt, T. J.; & Wu, T. (2025). **DinoAtten3D: Slice-Level Attention Aggregation of DinoV2 for 3D Brain MRI Anomaly Classification**. In *Proceedings of the IEEE/CVF International Conference on Computer Vision **(ICCV)** Workshop on Anomaly Detection with Foundation Models*.

[J.2] Siddiquee, M. M. R.; **Rafsani, F.**; Wu, T.; Chong, C. D.; Schwedt, T. J.; & Li, B. (2026). **AUCp: Pseudo-AUC for Inference Model Selection with Unlabeled Validation Data in Abnormality Detection**. *IEEE **Transactions on Medical Imaging***. [IF 12.4]

[C.2] Che, Y.; **Rafsani, F.**; Shah, J.; Siddiquee, M. M. R.; & Wu, T. (2025). **AnoFPDM: Anomaly Detection with Forward Process of Diffusion Models for Brain MRI**. In *Proceedings of the Winter Conference on Applications of Computer Vision **(WACV)** Workshops*.

[J.3] Dip, S. A.; Mallick, D.; Shuvo, U. A.; Soummo, S. B.; **Rafsani, F.**; Paul, B. K.; Moumi, N. A.; Ahmed, S.; & Zhang, L. (2025). **LLMAgent4Bio: LLM Agents for Biological Intelligence Across Genomics, Proteomics, Spatial Biology, and Biomedicine**. *Briefings in Bioinformatics*. [IF 7.7]

[J.4] Zaman, Z.; Rahman, S.; **Rafsani, F.** et al. (2023). **DeepVRM: Deep Learning Based Virtual Resource Management for Energy Efficiency.** *Journal of Network and Systems Management.* [IF 4.3]

[A.1] **Rafsani, F.**; Sheth, D.; Che, Y.; Shah, J.; Siddiquee, M. M. R.; Chong, C.; Nikolova, S.; Dumkrieger, G.; Li, B.; Wu, T.; & Schwedt, T. (2025). **Using Large-scale Contrastive Language-Image Pre-training to Maximize Brain MRI-Based Headache Classification.** *American Academy of Neurology Annual Meeting 2025*.

Research Projects

- **Ultrasound V-JEPA Foundation World Model for Thyroid Nodule Detection** *Mar 2025 - Present*
  Self-supervised video representation learning
  - Pretraining a **Video Joint-Embedding Predictive Architecture (V-JEPA)** on **1000 hours** of unlabeled thyroid ultrasound video to learn spatiotemporal representations for downstream nodule detection.
  - Designed a **reinforcement learning-based adaptive masking policy** that replaces fixed random masking during self-supervised pretraining, targeting clinically salient regions.
- **Diffusion Model-Based Anomaly Segmentation Using Brain MRI (AnoFPDM)** *Jan 2024 - Aug 2024*
  Funded by NIH, USA
  - Developed **AnoFPDM**, a fully weakly-supervised diffusion-based anomaly detection framework that removes any dependence on pixel-level labels during tuning.
  - Reached ≈**80% Dice** on **BraTS21** and **ATLAS v2.0**, surpassing state-of-the-art weakly-supervised baselines; published at **WACV 2025 Workshops**.
- **Multi-Modal Foundation Model Based Headache Detection** *Jan 2024 - Feb 2025*
  Funded by U.S. Department of Defense
  - Built a **BioMedCLIP**-based multimodal framework fusing 3D brain MRI with clinical metadata for headache subtype classification.
  - Improved classification performance by **5-10% AUC** per headache subtype over prior published methods; published in **Nature Scientific Reports** and presented at **AAN 2025**.
- **Automated Thyroid Nodule Segmentation from CT for Overdiagnosis Reduction** *Mar 2025 - Present*
  Foundation model benchmarking
  - Benchmarked **MONAI** 3D architectures (SegResNet, Auto3DSeg) against the **MedSAM** foundation model on thyroid CT, establishing the first internal baseline for automated nodule delineation.
  - Quantified the automation-accuracy trade-off: **74% Dice** fully automated vs. **80% Dice** with prompt-guided MedSAM, both exceeding prior baselines by **8-10%**; now closing the gap via automated prompt generation.

Skills

- **Deep Learning & Computer Vision:** PyTorch, Vision Transformers (ViT), DINOv2, CLIP/BioMedCLIP, Diffusion Models, V-JEPA, Self-Supervised Learning, Contrastive Learning, Semantic Segmentation, Anomaly Detection, MONAI, MedSAM, TensorFlow, Keras, scikit-learn, OpenCV
- **LLMs & Generative AI:** Hugging Face Transformers, parameter-efficient fine-tuning (LoRA/PEFT), multimodal vision-language modeling, LLM agent frameworks, prompt design and retrieval-augmented generation (RAG)
- **Training & Systems:** Multi-GPU distributed training (PyTorch DDP), mixed-precision training, gradient checkpointing, CUDA, SLURM / HPC clusters, Docker, Kubernetes, AWS (EC2, S3), Azure, GCP
- **MLOps & Tooling:** Weights & Biases, TensorBoard, Git/GitHub, CI/CD, FastAPI
- **Programming:** Python, C++, C, Java, SQL, JavaScript, Bash
- **Data & Scientific Computing:** NumPy, Pandas, SciPy, Matplotlib, NiBabel, SimpleITK, PostgreSQL, MongoDB
- **Software Engineering (prior):** Django, React, Node.js, REST APIs, full-stack web development

Awards & Service

- **Gerald Farin Memorial Fellowship** *2024-25, 2025-26*
  Arizona State University — awarded for academic excellence and research potential in computer vision
- **SCAI Conference Award** *Aug 2025*
  School of Computing and Augmented Intelligence, ASU — to present at ICCV 2025
- **Graduate College Travel Award** *Apr, Aug 2025*
  Arizona State University — to present at the American Academy of Neurology Annual Meeting 2025
- **Dean's List** *2021*
  Bangladesh University of Engineering and Technology — top academic standing, final year
- **Reviewer** *2024 - Present*
  IEEE Transactions on Medical Imaging, AAAI, ICCV, WACV, Journal of Headache and Pain