ARNABI DUTTA

LinkedIn | *GitHub* | *arnabidutta000@gmail.com* | *(+91) 9932740791*

Education

Delhi Technological University (DTU), *B.Tech in Computer Science & Engineering* **2024 – 2028**

Dhruva Public School (12th: 80%) | **St. George's School (10th: 98%)** **2020 – 2023**

Technical Skills and Coursework

Languages & Coursework: Python, C, C++, Probability and Statistics, OOPs, Software Engineering, DBMS, Computer Networks, Computer System & Architecture, Machine Learning, Deep Learning, Data Structures and Algorithms
Frameworks: PyTorch, Scikit-learn, Hugging Face, ROS2, mmcv, Docker, Git, LaTeX
Hardware: NVIDIA Jetson Nano/Orin, Raspberry Pi, Arduino, Pixhawk Cube

Publications & Projects

SPACE-HOP: Spacecraft 6-DoF Pose Estimation using Embedding-Predictive Pretraining and Hopf Map *CVPR 2026*

→ *[AI4Space Workshop Accept]* Introduced the first application of **Joint-Embedding Predictive Architectures (JEPA)** to spacecraft domain. Designed a keypoint masking strategy that forces the network to learn deep 3D structures, effectively overcoming severe orbital occlusions.
→ Structured the rotation as a Classification-and-Refinement task across a **uniform Hopf fibration grid**, refined with **continuous Lie algebra offsets** for **sub-degree accuracy**.
→ Built and **deployed** a low-latency inference **optimized for edge hardware**. Successfully bypassed computationally heavy PnP solvers to guarantee real-time orbital deployment, achieving **18.0 ± 3.8 ms inference latency** on an NVIDIA Jetson Orin.

TiPAI-TSPO: Tournament Inpainting for Patch-level Alignment in text-to-image diffusion models with Tournament Sampling Policy Optimization *EMNLP 2026 Under Review*

→ Proposed **TiPAI**, a decoding-time alignment framework for diffusion models using localized timestep-aware auditing and targeted inpainting for policy-compliant generation.
→ Introduced **TSPO** (Tournament Sampling Policy Optimization), a lightweight reinforcement learning framework for selecting optimal inpainting configurations under quality and compute constraints.
→ Developed a guarded tournament decoding pipeline with monotone non-regression guarantees using faithfulness, safety, and seam-quality based acceptance criteria.

Pathological Tremor Detection via Dense Point Tracking and LiDAR Perception *[GitHub]*

→ Engineered a **multi-modal 3D perception stack** integrating **LiDAR and depth cameras**, executing pointcloud alignment, **auto-calibration** (SuperGlue), and human segmentation (PointNet++) within a **ROS/RViz2** environment.
→ Augmented macroscopic 3D tracking with a **dense optical flow pipeline (AllTracker/RAFT)** and a custom statistical metric (Directness Ratio) to isolate flow vector variance, successfully distinguishing micro-scale **involuntary tremors** from voluntary locomotion.

WireFall: Self-Hardening Transformer WAF via Incremental Learning *[GitHub]* *Team Lead - SIH 2025 Top 5 Finalist*

→ Built a firewall that catches zero-day attacks utilizing a **DistilBERT** encoder trained via **Masked Language Modeling (MLM)** for reconstruction loss, coupled with a frozen contrastive [CLS] head to detect zero-day anomalies with **93.78% recall**.
→ Red-Team self-hardening loop was designed where a fine-tuned Llama agent interprets high-loss anomalies to autonomously generate and inject regex signatures into the static engine, instantly patching new exploits.
→ A **continual learning framework** was implemented to prevent catastrophic forgetting, leveraging Jensen-Shannon (JS) Divergence for stable knowledge distillation; derived a gradient-sensitivity metric to dynamically rank and stack LoRA adapters based on optimization relevance.

Experience

Machine Learning Research Lab, DTU , ***Deep Learning Research Intern*** *May 2025 – Present*

→ Working on **Vision Based Pose Estimation for Non-Cooperative Docking/Inspection Satellite** sponsored by **ISRO**.
→ Implemented **PVNet**, conducted **RANSAC** ablation, tested **ViTPose/HigherHRNet**,**Fastpose-ViT** solved **PnP** problem, applied continuous representation of 3d rotation matrices, and created high-fidelity synthetic spacecraft pose dataset (100k images) to enhance model robustness.

Pragya Lab, BITS Goa , ***Research Intern*** *March 2026 – May 2026*

→ Co-developed the TiPAI-TSPO architecture for decoding-time alignment in text-to-image diffusion models using reinforcement learning, under the mentorship of Aman Chadha and Prof. Dr. Amitava Das.

Unmanned Aerial Systems (UAS) DTU , ***Software Technician*** **[Website]** *Aug 2024 – June 2025*

→ Secured **2nd Place globally** (60k USD Phase 1, 150k USD Phase 2) in the **DARPA TRIAGE CHALLENGE** as a **self-funded student team**, deploying a fleet of autonomous **UAVs** and **UGVs** for mass-casualty triage.
→ Engineered an edge-optimized multi-modal perception stack on **Jetson Orin**, fusing foundation models (**YOLO, LLaVA, SAM, DINO, CLIP, ViTPose**) to autonomously detect casualties in smoke and darkness, executing real-time injury grading (ocular, verbal, and motor vitals) within <**30 minutes**.

EHAX (Cybersecurity DTU) , ***Lead of Research & Impact*** **[Website]** *Aug 2024 – Present*

→ Lead AI security projects: **BERT-powered Network Intrusion Detection System(NIDS)** and **LLM-driven vulnerability analysis agent**; performed AI threat modeling at CRAC Learning to investigate LLM exploitation vectors (OWASP Top 10).
→ **Led outreach** to mentor women in cybersecurity, deep learning research, and reverse engineering.

Achievements

→ Secured 2nd prize (60k USD Phase 1 2024, 150k USD Phase 2 2025) as core software team member of **UAS-DTU** in **DARPA Triage Challenge**
→ **National Finalist (Top 5)** as **Team Lead** at **Smart India Hackathon 2025**: Built Transformer WAF. Secured ISRO internship.
→ **National finalist (Top 12)** at **NCIIPC – AICTE Pentathlon 2025**, biggest govt initiated cybersecurity event in India.
→ Winner **VisionXAI Hackathon 2025** for the task of Person Re-Identification (ReID) in CCTV cameras.
→ **Chitra Visharad** in **Fine Arts**: 12 yrs exp. [Certificate]