Mudit Jain

👨‍💻 With expertise in 3D machine learning, SLAM, computer vision, GPU programming, and embedded systems, I excel in creating integrated engineering solutions through comprehensive, cross-disciplinary thinking.

💻 Skills

Programming : C++ [6+ years] , Python [6+ years]
Technical Knowledge Domains : Multimodal Large Scale Deep Learning **[Ray, Kubernetes, Pytorch]**, Classical computer vision **[C++, OpenCV]**, 3D Computer Vision, Machine Learning **[Pytorch, JAX]**, SLAM **[ORB-SLAM, VINS Mono]**, Non-linear optimization **[Eigen, g2o, ceres, GTSAM] - Bundle Adjustment, Camera Calibration, Pose Graph Optimization**, IMU Preintegration, Bayesian Inference, Embedded Systems, SIMD Programming **[CUDA]**, Model Optimization **[TensorRT],** 3D reconstruction [**NeRFs, Gaussian Splatting**], 2D/3D Object Detection

⏳ Experience

2024 - Present
Qualcomm, San Diego - *Senior Deep Learning Engineer - Multimodal AI*

- Develop **state of the art lidar and LiDAR-camera fusion deep learning models using foundation models** for complex urban and highway scenarios scalable to **large scale data regimes(Petabytes)**. Work closely with seasoned senior perception engineers in leading **automated driving systems in OEMS & Robotaxi companies**.

2024 - Present
Google Summer of Code - *Developer & Mentor - 3D Reconstruction*

- Developed a pipeline **to convert unconstrained video sequences** into **efficient Gaussian splats**. Collaborated with OpenCV under the mentorship of OpenCV founder and president, Gary Bradski.

2021 - 2024
Qualcomm, San Diego - *Senior Machine Learning Engineer*

- Served as the **technical lead**, overseeing the development of a **hardware/software co-design solution for visual odometry** in **extended reality (XR/VR/AR) applications**. The proposed HW/SW co-design **reduces power usage** while maintaining latency and accuracy. Designed and validated sparse optimizations for various non-linear estimation problems of **Simultaneous Localization and Mapping (SLAM)** namely B**undle Adjustment, Pose Graph Optimization, Loop Closure & Online Calibration**.

2019 - 2021
Drone Lab - UCSD, San Diego - *Graduate Student Researcher [Funded]*

- Designed, implemented & deployed an Attention-based CNN on **incoming data from 600 cameras** to solve the problem of **wildfire plume detection** for the **ALERTWildFire initiative**
- Developed an **Object Detection Pipeline** to process **4x1080p video streams** on AGX Xavier using the **DeepStream SDK** in collaboration with LLNL. The detection pipeline used custom trained **TensorRT models** as detectors.

2016 - 2019
NVIDIA, Bengaluru - *Embedded System Software Engineer*

- Served as the **lead** on designing, implementing and testing the **software pipeline of I2C Virtualization** as per the IS026262 functional safety standards for ARM based NVIDIA SoCs (Xavier/Parker).
- Took an active role in the **development and bring-up of the Xavier System-on-Chip (SoC)**, ensuring that the components under my purview were successfully integrated and operational on both the Field-Programmable Gate Array (FPGA) and the SoC.