

PROFESSIONAL SUMMARY

MASc graduate from the University of Waterloo with over 4 years of specialized experience in machine learning and computer vision. Proven expertise in real-time multi-object tracking, object detection, and segmentation, with a strong background in deploying end-to-end ML pipelines in autonomous vehicles and industrial automation. Recognized for academic contributions, including publications in top-tier conferences and a pending patent. Committed about advancing technologies in computer vision, machine learning, and robotics.

EDUCATION

MASc in Systems Design Engineering, *University of Waterloo, Canada*

Sep 2022 — Mar 2025

B.Tech in Electronics and Communication, *Guru Gobind Singh Indraprastha University, India*

Aug 2016 — Sep 2020

SKILLS

Programming & Platforms

Python, C++, C, Linux, Git, Docker, SSH, CI/CD

Machine Learning & AI

PyTorch, LibTorch, CUDA, TensorFlow, Keras, ONNX, Scikit-learn

Computer Vision

OpenCV, Pillow, YOLO, Object Detection, Segmentation, Feature Extraction

Data & Visualization

NumPy, Pandas, Matplotlib, Plotly, Seaborn, Weights & Biases

Hardware

VLP-16 LiDAR, FLIR Pointgrey Cameras, Raspberry Pi 3B, Arduino, NVIDIA Jetson

WORK EXPERIENCE

Computer Vision Researcher | VIP Lab, University of Waterloo & ATS Automation

Jan 2023 — Present

- Leading research on Visual Place Recognition (VPR) for aerial imagery; exploring transformer-based, cross-view, and multimodal architectures (CNNs, DINOv2, ViTs) to extract robust features for UAV-based geo-localization.
- Improved real-time multi-object tracking accuracy by 7% on team sports and dance scenarios by designing data-driven motion prediction model integrating Mamba state-space models with self-attention.
- Increased annotation speed by 10x by creating a novel automatic annotation pipeline for multi-object videos using Point Tracking, Segment Anything, and Object Detection (YOLO-v8).
- Containerized multi-object tracking pipeline for industrial deployment, enabling velocity estimation of 100+ assembly parts in high-throughput manufacturing environment, increasing production efficiency.

Computer Vision Research Engineer | LENS Corporation, India

Feb 2022 — May 2022

- Migrated a latent fingerprint extraction and matching pipeline from Python to C++ using Libtorch to enable faster GPU and CPU-based deployment, optimizing performance and reducing latency.
- Programmed custom signal processing functions, including Fourier Transforms and Gabor Filtering, in C++ from scratch to process fingerprint data, overcoming constraints in existing matrix and signal processing C++ libraries.

Perception Research Engineer | Autonomous Vehicle Project, IIIT Delhi, India ([View Project](#))

Oct 2020 — Feb 2022

- Built and deployed a multi-sensor calibration system integrating 2 FLIR Pointgrey cameras and 3 VLP-16 LiDARs using ROS for synchronized data capture and visualization in RViz; filtered sensor data with RANSAC and DBSCAN for robust checkerboard detection.
- Applied linear algebra (SVD), 3D geometry, and machine learning techniques in Python and scikit-learn to estimate and fine-tune inter-sensor transformations, achieving <4° rotational and <10 cm translational error across all modalities.
- Implemented LiDAR-LiDAR alignment with Iterative Closest Point (ICP) to minimize Euclidean distance between point clouds, strengthening overall sensor fusion accuracy.
- Developed and optimized a real-time lane detection pipeline in C++ and TorchScript; projected lane boundaries onto HD maps to identify driveable regions, maintaining 12+ FPS on NVIDIA Jetson Xavier.

Research Intern | Microsoft Research India ([View Project](#))

Feb 2020 — Sep 2020

- Redesigned and scaled the Automated Driver License Testing (ALT) project for Regional Transport Offices (RTO) across India. The scaling led to the deployment of the project in 10+ cities in India with 99% accuracy in automated driving test results.

PROJECTS

Image Enhancement and Object Detection in Rainy Weather Conditions | *Deep Learning, GANs, PyTorch*

Jan 2023 — Apr 2023

- Implemented deep learning models for image enhancement in rainy weather, integrating GAN-based denoising with morphological transformations and optimized loss functions, improving visual clarity by increasing PSNR by 5.8%.

RecipeDB | *Pandas, Data Visualizations* ([View Project](#))

Jun 2019 — Sep 2019

- Curated and integrated recipe data from 7 online sources, performed multi-level statistical analysis across 22 regions, and programmed interactive visualizations using Pandas, Matplotlib, and Plotly to present insights on regional dietary and nutritional patterns.

PUBLICATIONS

- SportMamba: Adaptive Non-Linear Multi-Object Tracking with State Space Models for Team Sports**, IEEE/CVF CVPR Workshop 2025.
- Attention-Mamba for Multi-Object Tracking**, Conference on Robots and Vision (CRV), 2025.
- POPCat: Propagation of Particles for Complex Annotation Tasks**, Conference on Robots and Vision (CRV), 2024.
- Recipedb: A resource for exploring recipes**, Database Journal (Oxford), pp. baaa077, Nov 2020.
- NTIRE 2019 Challenge on Video Super-Resolution: Methods and Results**, Co-author of team paper, IEEE/CVF CVPR Workshop 2019.