

Harshvardhan Chandrasekar

(+1)7633529820 | harshvardhan.chandira@gmail.com | linkedin.com/in/harshvardhan | github.com/perseusdg

EDUCATION

BITS Pilani, K.K. Birla Goa Campus

Bachelor of Engineering in Mechanical Engineering

Goa, India

Aug. 2018 – May 2022

University of Minnesota Twin Cities

M.S. in Robotics (3.81/4.0 GPA)

Minneapolis, USA

Sep. 2023 – Present

EXPERIENCE

Graduate Student Researcher

Zhu Tian Chen's Lab

Feb. 2024 – Present

Minneapolis, MN

- Developed collaborative AR systems (e.g., SpatialMiro) with LLM-based brainstorming tools and studied group performance.
- Built an interactive Gaussian Splatting AR framework achieving 90 FPS on Apple Vision Pro for real-time rendering.
- Developing machine learning models to generate contact-point deformations for Gaussian splats.
- Engineered a low-latency GPU-CPU-network pipeline using FlatBuffers for real-time training visualization of Gaussian splats.
- Explored egocentric Gaussian Splatting to merge static scenes with dynamic objects for personalized AR experiences.
- Investigated AR-based juggling instruction for muscle memory.
- Trained and developed an uncertainty head for flow-matching vision-language-action models like pi0.
- Pursuing a UIST 2026 publication to explore the future of apps, with a focus on MR and robotics accounting for the evolving spectrum of robot autonomy, to envision a future towards Sutherland's version of the ultimate display.

Research Engineer - Computer Vision

HiPeRT S.R.L

Sep. 2022 – May 2023

Modena, Italy

- Enhanced RTLS for automated forklift operations using multi-device SLAM and network mapping.
- Engineered a custom neural network library on Hailo-8 and Jetson AGX, enabling a YOLOv7 model (AP50: 72.5%, 60+ FPS).
- Engineered a WiFi-based warehouse localization system (1.5 m accuracy), saving €2000 by pre-mapping existing networks under varying loads.
- Collaborated with clients at Technology Innovation Institute (TII UAE) to conduct user surveys on long-range depth estimation networks and evaluate their performance on Jetson NX platforms.
- Documented technology evaluation findings and presented comparative performance metrics to both technical and product management teams.

Research Intern (UG Thesis: Real-time semantic SLAM for embedded systems)

HiPeRT Lab, UNIMORE

May 2021 – Aug. 2022

Modena, Italy

- Integrated tkDNN with ORB-SLAM3 to achieve 24 FPS on AGX Xavier.
- Minimized semantic drift to 2 frames compared to baseline RDS-SLAM (10 frames with Mask R-CNN, 5 with SegNet).
- Evaluated on KITTI sequences showing significant reductions in ATE (11.65%), APE (77.1%), and ARE (74.2%).
- Optimized GPU-network communication and CUDA kernels, cutting execution time by 3.2%.

INTERESTS

- Interactive neural rendering (NeRF, Gaussian Splatting) for AR/VR environments, human-computer interaction for spatial computing, real-time 3D scene reconstruction with interactive capabilities, human-robot interactions using mixed reality, and enhancing teleoperation for everyday robot use.

PROJECTS

Guitar Chord Recognition | JAX, MediaPipe

Oct. 2025 – Present

- Developed a custom neural network library in JAX as a part of CSCI 5561 (Computer Vision) as a learning platform for CNN and graph neural network implementations.
- Implemented real-time chord recognition using hand pose from MediaPipe.
- Developed a normalization method to abstract 21 hand points into 20 hand points that can be considered as a 4x5 image with 3 channels for use in baseline CNN and ResNet implementations.
- Preliminary tests indicate ResNet performs the best at around 88 percent accuracy, but the estimation is highly dependent on the viewing angle and the location of the chord on the fretboard.
- For the next steps, the project intends to mix in visual features by encoding the image through DINOv2 and ensembling it with the existing model.
- Open-source code can be found at: <https://github.com/chrismelville29/GuitarChords>

Learning Muscle Memory through AR | Unity, Quest 3

Sep. 2024 – Jan. 2025

- This course project serves as a user study to understand if muscle memory can be gained relatively quickly when learned in a user-controlled and modified environment in AR.

- Developed Unity functions built on top of Meta Quest building blocks to account for different hand poses for ball grab and release.
- Helped develop a custom physics system to slow down ball speed and gravity to give the user a better understanding of hand placement.
- Helped develop a questionnaire for the user study.

Home Security System based on AVR BLE board | *C, AVR-BLE, Swift*

Mar. 2024 – May 2024

- Developed a real-time embedded security system as the final course project, implementing motion and door magnetic sensors on an AVR-BLE board to monitor and detect potential intrusions within a household.
- Implemented secure data handling by leveraging the on-board secure element chip to sign and encrypt sensitive information related to the security system.
- Designed and developed an iOS application to facilitate user communication and provide real-time push notifications and Bluetooth connectivity for alerting users of any unexpected activity detected by the security system.

Emergency Response Automation: Harnessing RL | *Python, C++, CARLA, PyTorch*

Nov. 2023 – May 2024

- Explored the use of reinforcement learning in autonomous vehicles through CARLA and OpenAI Gym.
- Trained a reinforcement learning model based on PPO to navigate a specific town in CARLA from point A to point B while achieving the fastest time possible.
- Implemented a custom path planning algorithm that called on the local path planning algorithm to generate alternate routes in the event of a road closure or if the need arises for a faster alternative.

Navigating and Manipulating robots with LLMs | *ROS, LLM, Pose-estimation*

Oct. 2023 – Dec. 2023

- Integrated 2D object detection (YOLOv7) with spatial pose estimation to enable a robotic platform (Gazebo simulation - Baxter) to identify objects and determine their positions in 3D space for manipulation based on natural language commands.
- Developed a system for decomposing human commands into robot-understandable formats by breaking them down into a set of predefined actions such as move, pick, and place.
- Successfully integrated whisper.cpp and rosgpt, enabling the translation of human voice commands into broken-down tasks based on predefined prompts.
- Integrated object detection, path planning, and pose estimation components to facilitate the robotic system's ability to identify objects, plan paths, estimate object poses, and perform inverse kinematics for object manipulation tasks.

tkDNN | *Python, C++, CUDA, cuDNN, TensorRT, PyTorch*

May 2021 – Jun. 2023

- Successfully ported tkDNN to Windows based on popular demand in collaboration with HiPeRT.
- Created custom CUDA kernels to make tkDNN compatible with TensorRT 8+ and improved average inference speeds on FP16 for YOLO-based networks by 10%.
- Created additional pipelines to enable inference on transformer-based depth estimation neural networks.
- Ported tkDNN to support inference on edge detectors in combination with Jetson boards.
- The project can be found [here](#).

TECHNICAL SKILLS

Languages: Python, C/C++, MATLAB, Swift, C#, GLSL

Frameworks: React, CUDA, cuDNN, TensorRT, CARLA

Developer Tools: Git, Docker, AWS Bedrock, Google Cloud Platform, GDB, Linux

Libraries: pandas, NumPy, Matplotlib, PyTorch, TensorFlow, FFmpeg, OpenCV, OpenGL, Vulkan, Metal, PCL

Game Engines: Unreal Engine 5, Unity 2022

RELEVANT COURSEWORK

Graduate Coursework: Introduction to intelligent robotics systems, Machine Learning Fundamentals, Reinforcement Learning and Dynamic Programming, Real-Time Embedded Systems, Computer Graphics, DeepRob, Computer Vision, Augmented Reality Systems

MOOCs: CS231n by Stanford, CS234 by Stanford, Flutter Bootcamp on Udemy, Computer Graphics with Modern OpenGL and C++ on Udemy, Accelerated computer science fundamentals on Coursera