

# MOHAMED NASR

mderaznasr@gmail.com | +1 (613) 794-1005 | linkedin.com/in/mohderaznasr | github.com/MDerazNasr | scholar.google

## Education

### Georgia Institute of Technology | 3.9 GPA

Jan 2026 – Dec 2027

*M.Sc. Computer Science, Specialization in Computer Perception & Robotics*

**Coursework:** Computer Vision, Deep Learning, Reinforcement Learning, AI for Robotics (localization, SLAM, path planning), Machine Learning

### University of Ottawa | 3.7 GPA

Sep 2020 – May 2025

*Bachelor of Science (B.Sc.) in Software Engineering*

**Coursework:** Data Structures & Algorithms, OS (threads, processes, memory, synchronization), Embedded/Distributed Systems

## Technical Skills

**Languages:** C++, Python, Rust, Go, Bash, SQL, JavaScript, TypeScript

**Graphics & GPU:** OpenGL, Vulkan, GLSL, Qt, GLFW, GLM, Dear ImGui, CUDA, OpenCL, Nsight, VTune, Open3D

**Systems & Tools:** Linux, multithreading, memory management, CMake, Git/GitHub, Docker, Kubernetes, AWS EC2, ZeroMQ

**Compute & Inference:** OpenVINO, oneDNN, CUTLASS, TensorRT, ONNX Runtime, Triton, PyTorch, JAX

**Robotics & Simulation:** Isaac Lab, Isaac Sim, MuJoCo, Robosuite, Robomimic, COLMAP, ROS2, PyBullet

## Experience

### Qualia

Remote

*Research Engineer Intern*

Jun 2026 – Present

- Trained Stage-2 residual RLT policy on Pi0.5 (**JAX/PyTorch/Ray/AWS**) for peg insertion in **Isaac Lab**, lifting success **33%** on fixed-ID evals
- Built **Python** residual policy harness over **ZeroMQ/ROS2** to stream actions into **Isaac Lab**, matching **ManiSkill** within **2%** across sim evals

### Georgia Institute of Technology

Remote

*Graduate Researcher – HAAG Lab*

Jul 2026 – Present

- Building 3D mapping pipelines (**VGGT, COLMAP, Neuralangelo, PyTorch, Open3D**) for pose & geometry recovery from multi-view images
- Optimizing **C++/CUDA** GPU memory and latency with **Nsight** while deploying on **Linux** via **Docker, Kubernetes**, and **Git** CI/CD
- Fusing **point clouds** with SDFs, **3D Gaussian Splatting**, Instant-NGP, and **tiny-cuda-nn** for novel-view synthesis & synthetic training data

### OpenVINO (Intel) | Google Summer of Code

Remote

*ML Systems Engineer*

May 2026 – Aug 2026

- Cut Unitree UnifoLM-VLA inference latency **32%** & peak memory **19%** in **C++** (OpenVINO, **OpenCL/oneDNN** GEMM) for live robot control
- Exported Qwen2.5-VL to **ONNX/OpenVINO** at **PyTorch/vLLM** parity, streamed **ROS2** policies on **Isaac Sim** (**79%** success, **1k+** episodes)

### Shopify

Ottawa, Canada

*Machine Learning Engineer Intern*

Jan 2025 – Apr 2025

- Prototyped **PyTorch NLP** text-2-SQL model w/ schema grounding & constrained decoding, reducing querying time by **80%** for **1M+** entries
- Built ML pipelines (**Python/Spark/SQL/Pandas/Numpy/Airflow**) to predict business KPI metrics, improving prediction accuracy by **7%**
- Served real-time **XGBoost** KPI models on **TorchServe (Redis, Docker, AWS)**, cutting merchant dashboard response time by **35%**

### University of Ottawa

Ottawa, Canada

*Undergraduate Research Assistant*

Sep 2024 – Dec 2024

- Built **Reinforcement Learning** retraining controller w/ EWMA drift detection & policy evaluation, reducing model recomputation by **40%**
- Engineered **JAX/Python/C++** pipelines processing **200K+** solar readings daily, achieving **80%** data coverage via **Redis** caching
- Developed **React/Go/GraphQL** dashboard serving **XGBoost** forecasts & live model accuracy at **sub-200 ms** latency for live monitoring

## Publications

**World Model Lens (arXiv 2026)** | Unified interpretability for cross-architecture analysis of RSSM, JEPA, & Transformer world models

**Affinity Map (Zenodo 2026)** | Published meta-learning study for few-shot protein family classification with prototypical networks

## Projects

**3D Reconstruction Platform** – Python, PyTorch, 3D Gaussian Splatting, Nerfstudio, Open3D, COLMAP, cameras, React

- Mapped video to navigable 3D Gaussian scenes: frame extraction, COLMAP camera-pose recovery, and confidence-aware geometry

**Debris Engine** – C++, GLSL, OpenGL, Qt, Vulkan, GLFW, GLM, Dear ImGui, CMake

- Building robot debris physics sim with spatial-hash collision, contact debug, deterministic replay, and sensor-map export

**Inference Engine** – C++17, CUDA 12.x, CUTLASS 3.x, TensorRT, ONNX Runtime, Triton, pybind11, PyTorch

- Built flow-matching engine for Diffusion Policy and VLA DiT action heads (persistent trajectory kernels) targeting real-time robot control

**Reinforcement Learning F1 Sim** – Python, PyTorch, Stable-Baselines3, Gymnasium, TensorBoard, Kubernetes

- Built physics-based sim with PPO reinforcement learning & imitation learning, surpassing expert policy by 56% in race reward