

Bhashwar Ghosh

San Francisco, CA · bhashwarghosh@gmail.com · [linkedin.com/in/bhashwar-ghosh](https://www.linkedin.com/in/bhashwar-ghosh) · github.com/qmicon

Engineer with four years at NVIDIA on Metropolis, a video analytics platform that runs real time streams from hundreds of cameras and sensors. Built the analytics pipelines, the vision language and 3D perception systems that produce the data, and the retrieval layer that makes it answerable in natural language. Four US patent applications filed. Now working on training robot foundation models, mainly on inference and data infrastructure.

EXPERIENCE

NVIDIA / Metropolis Jul 2021 to Sep 2025

Senior Software Engineer (Mar 2025 to Sep 2025), Software Engineer (Jul 2021 to Mar 2025)

- Built a deep video research agent that answers queries over video, using vision language models and graph RAG for retrieval.
- Built the systems that compute vision analytics from real time streams of hundreds of cameras, audio and other sensors, and emit them in a schema built for retrieval.
- Worked on 3D vision language models, including SpatialRGPT, answering spatial questions about distance between objects, whether an object fits a given space, and object dimensions.
- Combined LLMs, VLMs, 3D VLMs, graph and hybrid RAG, grounding SAM and depth models into agentic analytics systems in LangGraph, and optimised those workflows for latency and performance.
- Built copilot systems that generate post processing pipelines by writing Python on the fly against low level analytics features, and agentic systems that query relational databases through generated SQL.
- Worked on a warehouse management AI agent that was shown in the NVIDIA GTC 2025 keynote.
- Promoted twice through NVIDIA's engineering levels, IC1 to IC3, reaching Senior Software Engineer.

Aelior Labs / Founder, Delaware C corp Jul 2026 to present

- Building robots for physical work in changing environments, instructed in plain language, with one person supervising several remotely.
- Entered the Stanford BEHAVIOR Challenge 2026 with a small team. 100 long horizon household manipulation tasks in Omniverse from RGBD and proprioception, against pi 0.5 and GR00T baselines.
- Ran technical and market discovery across the physical AI supply chain, including robot data vendors collecting at scale, distributors, and lab automation and space robotics researchers.

Founders, Inc. / Builder, Canopy program Apr 2026 to May 2026

Qmicon / Founder, incubated at IIT Madras (Nirmaan) Jul 2019 to Jun 2021

- Built a social app aimed at creation and productivity rather than passive consumption, for students.

PATENTS

Four US patent applications, all assigned to NVIDIA Corporation. Filed, not yet published.

- Object behavior identification in video using artificial intelligence models and historical object interaction data (Appl. 19/544,956, filed Feb 2026. First named inventor.)
- Performing operations using predicted outputs of a neural network (Docket 24-PU-0439US01, 2024. First named inventor.)
- Detecting movement of objects using one or more neural networks (Docket 24-PU-0819US01, 2024.)
- Mixture of physical AI agents (Docket 25-SC-0221US01, 2025.)

EDUCATION

Indian Institute of Technology Madras / B.Tech, Engineering Physics. GPA 9.13 2017 to 2021

Stanford University / Professional Education Program, Artificial Intelligence. Reinforcement Learning (XCS234), 2025
Deep Generative Models, Machine Learning with Graphs, Natural Language Processing. GPA 10.0

TECHNICAL

Languages Python, C++

ML PyTorch, vision language models, 3D perception and depth, open vocabulary segmentation, graph and hybrid RAG, LangGraph agentic systems

Systems Real time multi sensor pipelines, inference and latency optimisation, model serving, TensorRT, performance profiling with perf, tracing and eBPF, video codecs and streaming pipelines, WebRTC, Docker

Robotics Isaac Sim, Omniverse, BEHAVIOR-1K, openpi, LeRobot