

Dipesh Kumar

(646) 316-6937 | dk5288@nyu.edu | linkedin.com/in/dipeshk27 | github.com/DipeshK47 | dipeshk47.github.io | New York City, NY

EDUCATION

New York University

Master of Science in Computer Science, GPA: 3.89/4.0

Sept 2025 - May 2027

New York, NY

Coursework: Computer Vision, Machine Learning, Foundations of Data Science, Neuroinformatics

SRM Institute of Science and Technology

Bachelor of Technology in Computer Science & Engineering, GPA: 8.41/10

Aug 2021 - May 2025

Chennai, India

TECHNICAL SKILLS

Languages: Python, C++, SQL, JavaScript, TypeScript, Bash

ML & Deep Learning: PyTorch, TensorFlow, Scikit-learn, Hugging Face, Large Language Models, LangChain, LangGraph

Computer Vision: Instance Segmentation, Object detection, Image Registration, Image Fusion, OpenCV, CUDA

Data & Backend: NumPy, Pandas, SciPy, Matplotlib, PostgreSQL, Data Pipelines, FastAPI, REST APIs, Streamlit, React, Next.js

Cloud, Infrastructure and Tools: Git, AWS, GCP, CI/CD, Docker, Linux, HPC, Jupyter, Codex, Claude Code

EXPERIENCE

Founding Engineer

April 2026 - Present

EEVSEC Private Limited

eevsec.com

- Designed and built CyberSuraksha as sole engineer end to end, a school cyber-safety and computational thinking platform for Classes 3-10, with 3 schools in onboarding. Architected it on **Next.js**, **TypeScript**, and **Cloudflare Workers** over a multi-tenant **Cloudflare D1** database, with authenticated student, teacher, and admin dashboards, using **React** for gamified interactive modules and synthesized narration for animated lessons, plus a retrieval-grounded **AI assistant** that answers questions from course material with cited sources and reports class progress to teachers.
- Designed and shipped eevsec.com, the company's **20-page** public marketing site, covering product, pricing, and audience pages with responsive dark-mode layouts and search-optimized metadata. Drove **500+ waitlist signups** ahead of launch.

Research Assistant

May 2026 - Present

Neuroinformatics Lab, New York University

New York, NY

- Developed a **cross-modal registration workflow** to align live two-photon GCaMP and fixed confocal RNAscope volumes, cutting median nearest-neighbor distance **32%** across 2,382 in-vivo and 29,168 ex-vivo centroids. Scaled inference to **125,940 cells across 82 volumes and 284 image pairs**.
- Improved the lab's soma-print matcher with ICP, Umeyama, Hungarian refinement, and **Fused Gromov-Wasserstein optimal transport**, reaching **0.920 and 0.918 rank-1 accuracy** on 873 and 691-cell benchmarks and lifting a 40-cell cross-modal benchmark from 0.675 to **0.950**.
- Built a **VIA-to-Cellpose pipeline** converting polygon annotations into z-linked 3D instances, turning **5,036 slice instances into 767 global cells**. Fine-tuned Cellpose-SAM to raise held-out mean AP50 **45% across six public datasets**, and fixed a channel bug causing **20x over-detection**.

AI Engineering Intern

Oct 2024 - Dec 2024

Dalhousie University

Halifax, Canada

- Built a **three-class resume-to-job relevance classifier** in PyTorch using **DistilBERT embeddings** and 1D convolutional layers, training on a corpus of **1.6M+ job postings and 1,962 resumes** auto-labeled by TF-IDF skill similarity. Reached **94% held-out accuracy**, validating labels against a manual holdout and correcting class imbalance with synonym augmentation, resampling, and class-weighted loss.

Computer Vision Intern

Dec 2023 - Jan 2024

Centre for Airborne Systems, DRDO

Bengaluru, India

- Developed an **Oriented R-CNN detector** for arbitrarily oriented small ground objects in aerial imagery, achieving **92.8% proposal recall** with Midpoint Offset Representation and rotated RoI Align, then accelerated it to **15.1 FPS** on 1024x1024 images via structured pruning and CUDA parallelism.

PROJECTS

TubeMind: YouTube Research and Knowledge System

[GitHub](#)

- Built TubeMind to **accelerate YouTube research**, gathering evidence from many videos into a topic-scoped workspace with saved, timestamped answers. Built per-board **Fast GraphRAG** knowledge bases using **1,536-dimensional OpenAI embeddings** and HNSW vector search over FastHTML, HTMX, and SQLite, with Google OAuth, transcript fallbacks, retrieval-quality checks, and rate limiting, verified across **10 pytest modules**.

GARUDA: Indian Market Scanner and Quant Research Platform

[GitHub](#)

- Built a **multi-asset Indian market scanner** running user-defined strategies over **Upstox REST and WebSocket feeds** to fire **real-time alerts**, with a browser dashboard and an **event-driven backtester** covering ORB, Supertrend, and Bollinger with transaction costs, slippage, and risk exits, verified by **52 unit tests**.

ACHIEVEMENTS AND PUBLICATION

- Winner**, Kaggle Cell-Type Classification Competition, with top-10 finishes in several other ML competitions.
- Winner**, OSSOME Hack, GitHub Community Hackathon, for trAIner, an AI health platform.
- Published** "Evaluating Image Fusion Techniques for Improved Low-Light Surveillance" in IJERCSE, June 2024. [\[Paper\]](#)