dheera venkatraman

ai/ml applied scientist and engineer | robotics | computer vision | dheera@dheera.net

objective

Applied AI research in robotics, sustainability, humanity, education, creative tools, lab automation, coding, or the future of computing, and building products for a post-AI world. 15+ years of experience in ML including computer vision, generative AI, and leading technical teams. US citizen, world traveller and strong cross-cultural communicator. Bay-area based, also interested in Seattle / Boston / LA.

skills

Human Languages	English, Chinese/Mandarin, basic level German
Languages/Frameworks	Python, PyTorch, C, C++, ROS, ROS2, JavaScript, Node.js, HTML, CSS, SQL
AI/ML	Diffusion models, LLMs, transformer-based architectures, multi-modal models, NeRFs, Gaussian splats, monocular depth, object detection, segmentation, denoising, classical ML, pre-processing pipelines
Software	Robotic perception, path planning, system design, databases, API design, building web front-ends
Hardware	Building prototype robots with off-the-shelf hardware components and actuators, microcontrollers, basic DIY EE skills, designing simple 3D-printed parts
Science	Quantum mechanics, optics and imaging, signals and systems

selected recent experience

Amazon | Principal Applied Scientist, Amazon AGI | Santa Clara, CA, USA | 2022-2024
Generative AI research in text-to-video, text-to-image, and text-to-3D synthesis. Implemented multiple architectural changes and data processing pipelines; led technical direction among a team of researchers. Worked extensively with NeRFs, Gaussian splats, and diffusion models. Prior to generative AI, worked on cloud mapping and localization algorithms for robotic fleets.

Freedom Robotics | Robotics Software Engineer | San Francisco, CA, USA | 2020-2022
Architected and implemented the on-prem part of an enterprise robot monitoring and management solution and deployed with a top 10 auto manufacturer. Brought the solution from proof-of-concept, which brought the startup it's largest contract to date, and followed thorugh to implementation and post-deployment improvements. Developed integrations for multiple brands of industrial AGVs.

Robby Technologies (Y Combinator S16) | Co-Founder and CTO | Palo Alto, CA, USA | 2016-2020
Co-founded and scaled the company from hand-building the first few robots to leading a team of engineers to design, manufacture, and deploy a fleet of autonomous sidewalk delivery vehicles. Implemented a full perception stack including object detection, segmentation, and tracking. Built data collection and training pipelines. Designed and implemented a novel, robust semantic localization algorithm.

MIT RLE Optical and Quantum Communications Group | Research Assistant | Cambridge, MA, USA | 2007-2014
Experimentally implemented a single-photon LIDAR that was published in Science. Constructed an phase-conjugate optical coherence tomography experiment that de-bunked claims that certain features realized by quantum OCT were non-classical.

MIT-China Innovation and Entrepreneurship Forum (MIT-CHIEF) | Founding Team Member and Technology Director | 2011-2013
Responsible for technical conference infrastructure and logistics for the first 2 years, including a database-backed website, registration system, on-site kiosk displays, bilingual marketing and conference materials, video production, photography, and graphic design.

education

Massachusetts Institute of Technology (MIT)
Ph.D. Electrical Engineering (2015) | M.Eng. EE | S.B. Physics | S.B. EECS

a few of my side projects ...

AI for science: Using agentic AI to do (a) algorithm discovery in a container (b) common repetitive physical tasks in science research
ROSshow/ROSboard: Visualize ROS topics and sensor data (including PointClouds!) using ASCII art and streamed to a web browser
Numerous contributions to the ROS ecosystem including C++ drivers for popular sensors and motor drivers
DIY camera-based ADAS / lane keep system implemented on an actual car from scratch (NVIDIA AGX Xavier-based) (not comma.ai!)
Astrophotography denoising with neural networks to reduce the length of imaging sessions by 10X
BotParty: Built 10 WebRTC-based telepresence robots for <$120 in parts each, for telepresence gatherings during the COVID pandemic
Digital 4x5 back Gigapixel camera that scans and auto-stitches the image plane of a 4x5 view camera with an IMX477 sensor
Luxo: Designed a 3D-printed, servo-activated robotic Pixar lamp that is capable of locomotion by jumping