

Yaadhav Raaj

Perception • AI • Robotics • Engineer

yaadhavraaj@gmail.com



raaj.tech



soulslicer



yaadhavraaj



+14126289992



Raaj brings a decade of research and industry experience in modeling, solving and scaling computer vision solutions for Robotics and AR/VR systems

Education

Carnegie Mellon University – (Master of Science in Robotics)

Aug 19 – May 21

- Coursework: Computer Vision, Path Planning, SLAM, Reinforcement Learning
- Research: Human and Car Pose Estimation, Adaptive Lidar / Light Curtains for Self-Driving Vehicles

National University of Singapore – (Bachelor of Engineering in Computer Engineering)

Aug 11 – Feb 16

Developer Skills

Programming Languages:

C++, Python, CUDA, OpenCL, OpenGL, MATLAB, Bash, Javascript

Software Frameworks:

OpenCV, Pytorch, PCL, ROS, QT Framework, GTSAM, Git, CARLA, Mujoco, Unreal Engine, Three.JS, Pytorch Lightning, AWS S3, Tensorboard, ABB Robot Studio, Beckhoff Twincat, 3D Studio MAX, Meshlab, Docker, Kubernetes

Experience

Agility Robotics – (Perception Engineer III - Portland, OR)

Jun 21 – Current

- Lead the development effort in setting up the object detection framework on [Digit](#), our humanoid robot platform
- Developed the end-end system for data collection and annotation, with over a million unique segmentation and keypoint instances
- Developed a custom deep learning model to estimate instances, keypoints and 6DOF pose, leveraging rgb and depth modalities
- Selected and integrated the embedded NVIDIA GPU that goes on our robots, and optimized and deployed the model to run on hw
- Developed *patent-pending* [algorithms](#) to leverage robot arm and environment constraints to do 6DOF multi object tracking
- Our perception workflows are being scaled in *production* at [Amazon](#) and [GXO](#) today, moving over 1000 totes a day
- Developing a general purpose foundation model for object detection, and GenAI tooling for synthetic data generation

Skills: Custom PointNet++ head for Generalized RCNN / Vision Transformer model, Analytical Jacobians implemented for efficient SE3 pose optimization via Factor Graphs, Lie algebra, Kinematically constrained robot pose optimization for optimal viewpoint, Vision conditioned diffusion policies, Custom Libtorch CUDA kernels for RGB / Depth / Surface normals and iOU, ONNX / TensorRT model optimizations and deployment to embedded NVIDIA hardware, Implemented the entire relationship and systems with our data annotation services, Synthetic data generation via Blender Proc / Omniverse

CMU Robotics Institute – (Graduate Student - Pittsburgh, PA)

Feb 18 – May 21

Under: [Dr. Yaser Sheikh \(CMU / Facebook Reality Labs\)](#)

- Core Developer on the 2D Human Pose Detector called [OpenPose](#) with Linux, OSX and Windows support
- This was the main human pose estimation library used by developers through most of 2018 / 2019 with (30k Stars on Github)
- Developed the Recurrent Spatio-Temporal Affinity Fields for real time multi-person tracking accepted at [CVPR 2019](#)
- This was an oral paper at CVPR 2019, and was the top performer on the posetrack challenge in 2018
- Helped develop joint Hand / Body / Face face tracking [ICCV 2019](#) used in Meta as a key technology enabler
- Helped out in development on the [Panoptic Studio Dataset](#)
- Ported Jacobian Computations for the ADAM / SMPL model into Pytorch + CUDA that was used in a Ceres SLAM Optimizer

Under: [Dr. Srinivas Narashiman \(CMU\)](#)

- Wrote a dynamic control and rendering simulator for a special kind of lidar called [Programmable Light Curtains](#)
- Hand built an entire Lidar + Stereo + Light Curtain rig on a vehicle that was then synchronized and calibrated with a custom [dataset](#)
- Developed a deep learning algorithm that estimates a probabilistic depth field from these sensors, accepted at [CVPR 2021](#)

Skills: Network and training architecture updates to OpenPose such as DDP (Distributed Data Parallel) and mixed precision training, Custom CUDA and OpenCL kernels that run on OSX / Windows and Ubuntu for OpenPose, GLSL and CUDA based physics based lidar modeling, rewrote a differentiable CUDA implementation of ADAM model (SMPL equivalent), Multiview Stereo Net for multimodal sensor fusion and modeling of depth uncertainty at a pixel level, Low level LIDAR + Stereo hardware synchronization for driving data collection

Uber ATG – (Software Engineering Intern - Pittsburgh, PA)

Jun 20 – Aug 20

- Built upon the [MultiXNet](#) Architecture to make a Path Relative Metric and equivalent Spatio-Temporal Path Relative Field developed by my team, end-end differentiable, with custom Pytorch + CUDA modules via five point stencil

SpeedCargo – (Robotics Engineer - Singapore)

Jan 16 – Jan 18

- Worked on the Gudel / ABB industrial robot platform, automating the palletization and depalletization of [aviation cargo](#)
- Led the vision development, integrating and calibration a TOF and RGB camera on the end-effector of the robot
- Developed a sub-cm precision object measurement and mesh reconstruction [system](#) over large workspaces
- System [deployed](#) at the Singapore Changi Airport, dimensioning thousands of large cargo items a day

Skills: ROS + Nodelets for modules, Particle filter based optimization for sub-mm pose estimation, Robot and camera calibration over large workspaces with a custom cube calibration rig, Custom OpenCL + CUDA implementation of Octree (Used by Apple), Docker + Kubernetes for customer deployment

Bumblebee Robotics - (Software Developer - Singapore)

Jun 14 – Jun 16

- Worked on the [BBAUV](#) and [BBASV](#) Robot platforms for underwater applications
- Developed [patent-pending algorithms](#) for calibrating and localizing objects underwater, fusing Sonar, Camera, IMU and DVL sensors
- Helped team win 2nd place at [AUVSI RoboSub 2015](#) in San Diego, and 4th place at the [AUVSI RobotX 2016](#) in Hawaii

Publications / Patents

- Generating object references from perception data in unstructured environments and related technology** {Patent Pending} Nov 23
{Yaadhav Raaj, and other authors}
- Generating references for robot-carried objects and related technology** {Patent Pending} Sep 23
{Yaadhav Raaj, and other authors}
- Exploiting and Refining Depth Distributions with Triangulation Light Curtains** (CVPR 2021) Jun 21
{Yaadhav Raaj, Siddharth Ancha, Robert Tamburo, David Held, Srinivasa G. Narasimhan} [[pdf](#)]
- Method and system for object localization, and controlling movement of an autonomous underwater vehicle for object intervention** {WO Patent WO2021045679A1} Mar 21
{Eng Wei Goh, Wan Theng Ruth Chew, Yaadhav Raaj} [[link](#)]
- Active Perception using Light Curtains for Autonomous Driving** (ECCV 2020) Aug 20
{Siddharth Ancha, Yaadhav Raaj, Peiyun Hu, Srinivasa G. Narasimhan, David Held} [[pdf](#)] [[video](#)]
- Single-network whole-body pose estimation** (ICCV 2019) Sep 19
{ICCV 2019 - Poster}
{Gines Hidalgo, Yaadhav Raaj, Haroon Idrees, Donglai Xiang, Hanbyul Joo, Tomas Simon, Yaser Sheikh} [[pdf](#)] [[video](#)]
- Efficient Online Multi-Person 2D Pose Tracking with Recurrent Spatio-Temporal Affinity Fields** (CVPR 2019) Jun 19
{Yaadhav Raaj, Haroon Idrees, Gines Hidalgo, Yaser Sheikh} [[pdf](#)] [[video](#)]
- Adapting the Search Subspace of a Particle Filter using Geometric Constraints** May 17
{Nikhil Somani, Yaadhav Raaj, Suraj Nair, and Alois Knoll} [[pdf](#)] [[video](#)]
- Precise Measurement of Cargo Boxes for Gantry Robot Palletization in Large Scale Workspaces using Low-Cost RGB-D Sensors** (ACCV 2016) Nov 16
{Yaadhav Raaj, Suraj Nair, and Alois Knoll} [[pdf](#)] [[video](#)]
- 3D Object Localization using Forward Looking Sonar and Optical Camera via Particle Filter Calibration and Fusion** (IEEE Oceans 2016) Sep 16
{Yaadhav Raaj, Alex John, Tan Soon Jin} [[pdf](#)] [[video](#)]
- Design And Implementation Of Bumblebee ASV 1.0** [[pdf](#)] [[video](#)] Jun 16
Design and Implementation Of Bumblebee AUV 3.0 [[pdf](#)] [[video](#)]
{Robonation Journal} {One of several authors}

Others

Languages:

English, Tamil, and basic Mandarin

Interests:

Outdoor Rock climbing, Mountaineering