Akshay Rangesh

3782 1st Ave, Unit 203, San Diego, CA, 92103 • (424) 644-7701

arangesh@ucsd.edu • Website • Google Scholar • GitHub • LinkedIn

Summary

I am a Ph.D. candidate with expertise in applied computer vision and machine learning - specifically in the context of autonomous driving and driver safety. I enjoy building systems for the real world, and have experience in creating multi-sensory automotive testbeds, the capture and curation of public datasets, and the development and deployment of unique real-time models and algorithms.

Education

- **University of California, San Diego (UCSD)**
  Ph.D. in Electrical and Computer Engineering (advised by Prof. Mohan M. Trivedi) — Jul 2016 – present
  M.S. in Electrical and Computer Engineering, GPA 3.84/4.00 — Sep 2014 – Jun 2016
- **National Institute of Technology, Silchar, India (NITS)**
  B.Tech. in Electronics and Communication Engineering, GPA 8.72/10.00 — Jul 2010 – Jun 2014

Publications (selected)

- **Object Detection**
  - LaneAF: Robust Multi-Lane Detection with Affinity Fields [paper,videos,code]
    Proposed a new clustering approach based on affinity fields for lane detection and segmentation that achieves state-of-the-art performance on the CULane dataset.
  - Ground Plane Polling for 6DoF Pose Estimation of Objects on the Road [paper,videos,code]
    Developed an efficient polling scheme using ground planes that allows any 2D object detector become a monocular 3D object detector by predicting few additional keypoints of interest. This resulted in a real-time detector that achieved the best orientation accuracy score on KITTI at the time of publication.
- **Multi-Object Tracking**
  - TrackMPNN: A Message Passing Graph Neural Architecture for Multi-Object Tracking [paper,code]
    Introduced a framework based on dynamic, undirected, bipartite graphs, and created a message-passing neural network that operates on it to produce tracks. This lightweight approach uses only 2D box locations and still manages to beat all competing approaches using RRC detections in the KITTI MOT challenge.
  - No Blind Spots: Full-Surround Multi-Object Tracking for Autonomous Vehicles using Cameras & LiDARs [paper,videos]
    Extended the popular MDP framework for tracking and made it amenable to multiple sensors and different sensor modalities. The approach was validated using a highly sensorized testbed comprised of a full-surround sensor suite.
- **Trajectory Prediction**
  - How would surround vehicles move? A Unified Framework for Maneuver Classification and Motion Prediction [paper,videos]
    Created a unified model that produces multiple plausible trajectories for each agent conditioned on discrete maneuver classes. The most likely maneuver for each agent is then chosen through an energy optimization formulation that minimizes conflicts between neighboring agents.
- **Driver Behavior Analysis**
  - Gaze Preserving CycleGANs for Eyeglass Removal & Persistent Gaze Estimation [paper,videos,code+data]
    Extended our previous work on driver gaze estimation to better handle challenges faced in the real world, namely - illumination and eyeglasses. We conduct experiments to understand the effects of harsh lighting on IR sensors, and propose gaze preserving CycleGANs for eyeglass removal while retaining the direction of gaze.
  - HandyNet: A One-stop Solution to Detect, Segment, Localize & Analyze Driver Hands [paper,videos,code+data]
    Introduced a semi-automatic labelling scheme for driver hand instances using chroma-keying. We then used this dataset to train models to detect, segment, localize and classify driver hands.
  - Forced Spatial Attention for Driver Foot Activity Classification [paper,videos,code+data]
    Proposed a spatial attention loss to forcefully introduce spatial awareness into a foot activity classification model.

- **Tools & Testbeds**
  - 3D BAT: A Semi-Automatic, Web-based 3D Annotation Toolbox for Full-Surround, Multi-Modal Data Streams [paper,videos,code]
    Created an open-source web-based annotation tool for labelling datasets with multiple sensors and modalities. Care was taken to reduce human effort and leverage geometry to automatically transfer labels between different sensor views.
  - LISA-T Testbed [paper,videos]
    Tesla Model S platform comprised of 11 GoPro Hero 4 Blacks, 1 Kinect v2, 1 Velodyne VLP-16 HiRes LiDAR, 1 SwiftNav Duro Inertial, and 2 VL6180 IR sensors. This testbed has produced 100s of hours of high-quality data that is calibrated and timestamped for research purposes. We are planning on a public release of this dataset in the near future.
  - LISA-A Testbed [paper,videos,data]
    Toyota Avalon platform comprised of 8 PointGrey Flea3 RGB cameras, 6 iBeo LiDARS, 4 Delphi SRR2 Radars, and a Mobileye ADAS. Fully calibrated sensory suite used to develop fusion-based perception algorithms.

Patents

- **Surround Vehicle Tracking and Motion Prediction (US20210056713)**
  Akshay Rangesh, Mohan M. Trivedi, Nachiket Deo
- **Neural Network for Object Detection and Tracking (currently filed)**
  Akshay Rangesh, Vahid R. Ramezani, Benjamin Englard, Siddhesh Suhas Mhatre, Meseret R. Gebre, Pranav Maheshwari

Skills

- **Programming Languages:** C++, Python, MATLAB, Shell
- **Deep Learning Frameworks:** Caffe, TensorFlow, Keras, PyTorch, TensorRT
- **Web Design:** HTML, CSS (basics), JavaScript (basics)
- **Other Skills:** Robot Operating System (ROS), Git/SVN, Arduino prototyping

Work Experience

- **Data Fusion Intern, Luminar Technologies**
  Palo Alto, CA — Summer, 2019
- **Researcher, Laboratory for Intelligent & Safe Automobiles (LISA)**
  UC San Diego, CA — Jul 2015 – Present
- **Research Assistant, Indian Institute of Technology**
  Guwahati, India — Summer, 2013

Teaching Experience

- Autonomous Driving Systems (Spring 2020): 100% students recommend
- Special Topics in Robotics and Control Systems (Spring 2018): 100% students recommend
- Introduction to Intelligent Systems (Winter 2016 & Winter 2017): 100% students recommend
- Digital Image Processing (Fall 2015): 97.1% students recommend

Courses (selected)

Computer Vision (I & II) ∘ Artificial Intelligence ∘ Parameter Estimation (I & II)
Statistical Learning ∘ Optimization on Manifolds ∘ Digital Signal Processing
Data Mining and Predictive Analytics ∘ Design and Analysis of Algorithms