

Sudeep Pillai, Ph.D.

+1 (650) 267-3765 | sudeep.pillai@gmail.com | [website](#) | [gscholar](#) | [github](#)

ABOUT	Technology expert (CS PhD) and engineering lead with strong ML/DL, CV, SLAM and probabilistic robot perception expertise, interested in developing software, algorithms and infrastructure technologies to support highly autonomous vehicles (AVs).
PROFESSIONAL EXPERIENCE	<p>Toyota Research Institute (TRI), Los Altos, CA Jun 2019 - Mar 2021 <i>Manager, Machine Learning Engineering</i> Managed and led the ML Engineering team responsible for the design, development and lifecycle of all ML/DL-driven technologies deployed to the TRI autonomous vehicle (AV) fleet.</p> <ul style="list-style-type: none">• Led the design, development and lifecycle of all deep-learning perception models deployed to the TRI AV fleet• Spearheaded and executed on long-term strategy for large-scale semi-supervised learning leveraging worldwide fleet data• Directed the development of bleeding-edge and high-performance ML infrastructure to support large-scale, distributed deep-learning on production-scale (~100TB) AV datasets• Supervised the production dataset pipeline for all ML models and increased throughput by 10x by unifying processes including ETL, curation, labeling and testing under a common framework• Managed and oversaw the overall production ML modeling budget including labeling services, large-scale cloud compute, professional contracting and experimental POCs <p><i>Research Scientist & Technology Lead, Machine Learning</i> Oct 2017 - May 2019 Responsible for the design and development of advanced machine learning techniques for Toyota's next-generation autonomous vehicles.</p> <ul style="list-style-type: none">• Spearheaded the development of self-supervised learning strategy within Toyota's autonomous driving fleet• Supported several research projects that led to the development and publication of multiple state-of-the-art monocular depth estimation methods at top-tier peer-reviewed conferences <p>PhaseSpace Motion Capture, San Leandro, CA 2009 – 2011 <i>Computer Vision Engineer</i> Software developer responsible for the development of cutting-edge computer vision and motion-capture technology stack. <i>Key Technologies:</i> Real-time sub-pixel fiducial tracker; Vision-based 6-DOF Pose Estimation; Visual-inertial 3D tracking; Real-time Facial Expression Synthesizer; Multi-camera Calibration (upto 32 cameras); Multi-view Geometry, OpenCV, SIMD/CUDA.</p>
EDUCATION	<p>Massachusetts Institute of Technology, Cambridge, MA 2014 – 2017 Doctor of Philosophy (Ph.D), Electrical Engineering and Computer Science (EECS) <i>Thesis: SLAM-aware, Self-Supervised Perception in Mobile Robots</i> Advisor: John J. Leonard</p> <p>Massachusetts Institute of Technology, Cambridge, MA 2011 – 2014 Master of Science (S.M), Electrical Engineering and Computer Science (EECS)</p> <p>University of Michigan, Ann Arbor, MI 2005 – 2008 Bachelor of Science in Engineering (B.S.E), Mechanical Engineering</p>
PROFESSIONAL INTERNSHIPS	<p>Mitsubishi Electric Research Laboratories, Cambridge, MA May-Aug 2014 <i>Research Intern - Spatial Analysis Team</i></p> <p>Segway Inc., Manchester, NH May-Aug 2008 <i>R & D Engineer Intern - Product Development and Future Technologies</i></p>
PATENTS	<p>Auto-labeling of driving logs using analysis-by-synthesis and unsupervised domain adaptation 2020</p> <p>Systems and methods for depth estimation using monocular images 2020</p> <p>High-Speed and Tunable Scene Reconstruction Systems and Methods Using Stereo Imagery 2019</p>

SELECTED PUBLICATIONS	<p>Lee et al. "End-to-end birds-eye-view flow estimation for autonomous driving", <i>IROS '20</i>.</p> <p>Tang et al. "Neural Outlier Rejection for Self-Supervised Keypoint Learning", <i>ICLR '20</i>.</p> <p>Ambrus et al. "Two Stream Networks for Self-Supervised Ego-motion Estimation", <i>CoRL '20</i>.</p> <p>Guizilini et al. "Robust Semi-Supervised MonoDepth Estimation with Reprojected Dist.", <i>CoRL '20</i>.</p> <p>Guizilini et al. "3D Packing for Self-Supervised Monocular Depth Estimation", <i>CVPR '20</i>.</p> <p>Pillai et al. "SuperDepth: Self-Supervised, Super-Resolved Monocular Depth Estimation", <i>ICRA '19</i>.</p> <p>Pillai et al. "Self-Supervised Visual Place Recognition Learning in Mobile Robots", <i>IROS '17 (Wkshp)</i>.</p> <p>Pillai et al. "Towards Visual Ego-motion Learning in Robots", <i>IROS '17</i>.</p> <p>Moll et al. "Exploring big volume sensor data with Vroom", <i>VLDB '17 (Demo Track)</i>.</p> <p>Fourie et al. "Centralized Graph Databases for Mobile Robotics", <i>ICRA '17</i>.</p> <p>Pillai et al. "High-Performance and Tunable Stereo Reconstruction", <i>ICRA '16</i>.</p> <p>Pillai et al. "Monocular SLAM Supported Object Recognition", <i>RSS '15</i>.</p> <p>Ramalingam et al. "Line-Sweep: Cross-Ratio for Wide-Baseline Matching & Recons.", <i>CVPR '15</i>.</p>																				
INVITED TALKS	<table border="0"> <tr> <td>TWIMLCon 2021, Virtual</td> <td>Feb 2021</td> </tr> <tr> <td>ODSC West 2019, San Francisco, CA</td> <td>Nov 2019</td> </tr> <tr> <td>NVIDIA GPU Technology Conference (GTC), San Francisco, CA</td> <td>Mar 2019</td> </tr> <tr> <td>AutoAI, San Francisco, CA</td> <td>Mar 2019</td> </tr> <tr> <td>RE·WORK Deep Learning Summit 2018, Boston, MA</td> <td>May 2018</td> </tr> <tr> <td>NVIDIA Research, Santa Clara, CA</td> <td>Mar 2017</td> </tr> <tr> <td>Microsoft Analog Research and Development, Seattle, WA</td> <td>Feb 2017</td> </tr> <tr> <td>MIT CSAIL Advisory Board Meeting, Cambridge, MA</td> <td>Apr 2016</td> </tr> <tr> <td>Boston Imaging and Vision Meetup, Cambridge, MA</td> <td>Jan 2016</td> </tr> <tr> <td>Association for the Advancement of Artificial Intelligence (AAAI '15), Austin, TX</td> <td>Jan 2015</td> </tr> </table>	TWIMLCon 2021, Virtual	Feb 2021	ODSC West 2019, San Francisco, CA	Nov 2019	NVIDIA GPU Technology Conference (GTC), San Francisco, CA	Mar 2019	AutoAI, San Francisco, CA	Mar 2019	RE·WORK Deep Learning Summit 2018, Boston, MA	May 2018	NVIDIA Research, Santa Clara, CA	Mar 2017	Microsoft Analog Research and Development, Seattle, WA	Feb 2017	MIT CSAIL Advisory Board Meeting, Cambridge, MA	Apr 2016	Boston Imaging and Vision Meetup, Cambridge, MA	Jan 2016	Association for the Advancement of Artificial Intelligence (AAAI '15), Austin, TX	Jan 2015
TWIMLCon 2021, Virtual	Feb 2021																				
ODSC West 2019, San Francisco, CA	Nov 2019																				
NVIDIA GPU Technology Conference (GTC), San Francisco, CA	Mar 2019																				
AutoAI, San Francisco, CA	Mar 2019																				
RE·WORK Deep Learning Summit 2018, Boston, MA	May 2018																				
NVIDIA Research, Santa Clara, CA	Mar 2017																				
Microsoft Analog Research and Development, Seattle, WA	Feb 2017																				
MIT CSAIL Advisory Board Meeting, Cambridge, MA	Apr 2016																				
Boston Imaging and Vision Meetup, Cambridge, MA	Jan 2016																				
Association for the Advancement of Artificial Intelligence (AAAI '15), Austin, TX	Jan 2015																				
ACADEMIC DUTIES	<p>Tutorial / Workshop Co-organizer:</p> <ul style="list-style-type: none"> • Workshop on Deep Learning for Visual SLAM, CVPR 2018 2018 • Learning for Mapping Workshop, IEEE IROS 2017 2017 • Geometric and Semantic 3D Reconstruction, CVPR 2017 2017 <p>Conference and Journal Reviewer:</p> <ul style="list-style-type: none"> • Conference on Robot Learning (CoRL) 2019-2020 • IEEE International Conference on Robotics and Automation (ICRA) 2014 - Present • IEEE International Conference on Intelligent Robots and Systems (IROS) 2015 - Present • IEEE Robotics and Automation Letters (RA-L) 2016 - Present • IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017 - Present • Robotics Science and Systems (RSS) 2016 • Autonomous Robots (AURO) 2016 - 2018 • International Conference on Computer Vision (ICCV) 2017 																				
PROGRAMMING	<p>Programming Languages: Python, C/C++, C++17, CUDA, Javascript, WebGL/Three.js, SIMD/SSE, Bash, Emacs Lisp, SQL.</p> <p>Software / Libraries: PyTorch, TensorFlow, Docker, Horovod, Ceres Solver, OpenCV, Eigen, PCL, LCM/ROS/Gazebo, ISAM/GTSAM, Apache Spark, CGAL, OpenGL, Boost, Boost.Python, Pybind11, Conda build, OpenMP, GDB/PDB, Valgrind, Linux</p> <p>Python: Keras, Scikit-Learn, NumPy, Boost-Python, Cython, Pandas, SciPy, PyTables, NetworkX, graph-tool, Numba, Anaconda</p>																				

Keywords: Machine Learning (ML), Deep Learning (DL), Self-supervised Learning, Semantic Scene Understanding, Active Learning, Computer Vision, Simultaneous Localization and Mapping (SLAM), Structure-from-Motion (SfM), Self-driving cars, Autonomous Robots, Robot Perception, Object Recognition, Probabilistic Graphical Models, Sensor Fusion.