

Siddharth Choudhary

itzsid.github.io

Email : itzsid@gmail.com

Mobile : +1-404-519-2452

EDUCATION

- **Georgia Institute of Technology** Atlanta, GA
Ph.D. in Computer Science; GPA: 3.84 Aug. 2012 – Aug 2017
- **IIIT Hyderabad** Hyderabad, India
Master of Science in Computer Science; GPA: 4.0 (10.0/10.0) Aug. 2010 – July. 2012
- **IIIT Hyderabad** Hyderabad, India
Bachelor of Technology (Honors) in Computer Science; GPA: 3.31 (8.28/10.0) Aug. 2006 – July. 2010

EXPERIENCE

- **Amazon Lab126** Sunnyvale, CA
Applied Scientist Septemeber 2020-
 - **Amazon Halo:** Working on the CV/ML features for Amazon Halo.
- **Magic Leap** Sunnyvale, CA
Principal Computer Vision Researcher/Engineer October 2017-September 2020
 - **3D Object Recognition in AR Cloud:**
 - * Lead Architect and Developer for **3D Object Recognition feature** deployed in ML-19. The pipeline was designed and built from scratch.
 - * Object Recognition algorithm is scalable in a large number of concurrent users, number of objects or size of map while maintaining a lower memory footprint.
 - * Led the team through research and development phase to deliver Object Recognition for ML19 OTA 1.
 - * Collaborated across other teams like Data, Cloud and Deep Learning teams to build evaluation and visualization tools.
 - * Object Recognition is a key feature for ML19 and one of the most requested feature for ML2 from developer feedback.
 - * Research paper accepted at CVPR 2020 Workshop on AR/VR. Patent filed as well.
 - **Learned Keyframe Selection for SLAM:** Designed and implemented a deep network which learns a frame embedding on top of descriptors to improve retrieval for localization. Improved the localization recall w.r.t Bag of Words algorithm by 10-30% while reducing the memory requirement to store each keyframe by 75%. Optimized the network resulting in similar computational requirement as Bag of Words.
 - **Scalable Infrastructure for SLAM Research:** Designed and implemented a scalable infrastructure to extract data from various stages in the SLAM pipeline along with ground-truth. This enabled scalable training and evaluation of various SLAM related machine learning algorithms.
- **Institute for Robotics and Intelligent Machines, Georgia Tech** Atlanta, GA
Graduate Research Assistant August 2012-August 2017
 - **Distributed Object-based SLAM:** Proposed a distributed algorithm based on Distributed Gauss-Seidel to estimate the 3D trajectories of multiple cooperative robots from relative pose measurements. Our approach has several advantages. It requires minimal information exchange, has an anytime flavor, scales well to large teams, and it has a straightforward implementation. Extended the distributed SLAM framework to include object level semantics resulting in a distributed object level SLAM approach.
 - **Memory-efficient SLAM:** Proposed an approach to scalable SLAM that is exactly sparse. Used Alternating direction method of multipliers to enforce consistency among various subgraphs.
- **Center for Visual Information Technology, IIIT Hyderabad** Hyderabad, India
Research Assistant August 2010-August 2012
 - **Bundle Adjustment on GPU:** Developed a hybrid implementation of sparse bundle adjustment on the GPU using CUDA, with the CPU working in parallel. The algorithm is decomposed into smaller steps, each of which is scheduled on the GPU or the CPU. Our implementation outperforms the CPU implementation significantly, achieving a speedup of 30-40 times over the standard CPU implementation for datasets with upto 500 images on an Nvidia Tesla C2050 GPU.

- **Visibility Probability Structure from SfM Datasets:** We encode the visibility information between and among points and cameras as visibility probabilities. The conditional visibility probability of a set of points on a point (or a set of cameras on a camera) can rank points (or cameras) based on their mutual dependence. We combine the conditional probability with a distance measure to prioritize points for fast guided search for the image localization problem.

PUBLICATIONS

- Multiuser, Scalable 3D Object Detection in the AR Cloud (**CVPR 2020 workshop on AR/VR**)
- Distributed Mapping with Privacy and Communication Constraints: Lightweight Algorithms and Object-based Models. (**IJRR 2017**)
- Data-Efficient Decentralized Visual SLAM (**ICRA 2018**)
- Distributed Trajectory Estimation with Privacy and Communication Constraints: a Two-Stage Distributed Gauss-Seidel Approach (**ICRA 2016**)
- Active planning based extrinsic calibration of exteroceptive sensors in unknown environment (**IROS 2016**)
- Exactly Sparse Memory Efficient SLAM using the Multi-Block Alternating Direction Method of Multipliers (**IROS 2015**)
- Information based Reduced Landmark SLAM (**ICRA 2015**)
- SLAM with Object Discovery, Modeling and Mapping (**IROS 2014**)
- Geometry Directed Browser for Personal Photographs (**ICVGIP 2012**)
- Visibility Probability Structure from SfM Datasets and Applications (**ECCV 2012**)
- Practical Time Bundle Adjustment for 3D Reconstruction on GPU (**CVGPU 2010**)

ACADEMIC DUTIES

- Conference reviewer for IEEE International Conference on Robotics and Automation since 2015.
- Conference reviewer for IEEE/RSJ International Conference on Intelligent Robots and Systems since 2014.
- Journal reviewer for IEEE Transactions of Robotics.
- Journal reviewer for IEEE Robotics and Automations letters.
- Journal reviewer for International Journal for Robotics Research.
- Program Committee member of Robotics Science and Systems conference in 2016.
- Member of the Organizing committee for the tutorial on Perception at Magic Leap at the 2019 Conference on Computer Vision and Pattern Recognition.
- Member of the Organizing committee for the tutorial on Into the MagicVerse at the 2020 Conference on Computer Vision and Pattern Recognition.

PROGRAMMING SKILLS

- **Languages:**C, C++, Python, MATLAB, CUDA
- **Libraries:**OpenCV, GTSAM, PCL, ROS, Eigen, Boost, PyTorch, Numpy, SciPy