Yutian Chen

yutianch@andrew.cmu.edu | www.yutianchen.blog | github.com/MarkChenYutian

Education

Carnegie Mellon University, MS. Robotics — Sep 2025 - Now
Carnegie Mellon University, BSc. Computer Science — Sep 2021 - May 2025

Experience

Research Assistant, Prof. Sebastian Scherer, Carnegie Mellon University — Sep 2022 - Now

- Co-proposed the *MAC-VO* idea and developed the full system from scratch, including the core stereo VO design and implementation, data and training infrastructure, and experimental evaluation.
- Built and deployed real-time perception stacks on NVIDIA Jetson Orin and Thor, designing a multi-process, GPU-accelerated pipeline with ROS and TensorRT, making *MAC-VO* and *UFM* meets compute limit for UAV.
- Developed inertial odometry and uncertainty modules (*AirIO*, *AirIMU*), implementing differentiable Euler/RK4 IMU pre-integration, train and evaluation pipelines across platforms, and integrating them into the SLAM stack.

Research Intern, Field AI — Jun 2025 - Aug 2025

- Led the research of *Co-Me*, a confidence-guided token-merging method that accelerates VGGT and related visual geometric transformers by up to 11.3× without retraining, enabling real-time 3D perception for embodied AI.
- Built the end-to-end training, evaluation, and profiling pipeline for *Co-Me*, wrote custom CUDA kernels and verified the acceleration on both datacenter GPUs and Jetson-class edge compute.

Research Intern, Prof. Chuang Gan, MIT-IBM Watson AI Lab — Apr 2024 - Jan 2025

- Key contributor to *Virtual Community* project, developed a city-scale 3D scene generation pipeline from satellite, aerial, and street-view imagery, including reconstruction, meshing, and refinement for urban environment.
- Integrated this pipeline with the *Genesis* simulation engine to support large-scale physics simulation, enabling scalable evaluation of visuomotor agents in real-world grounded scenes.

Research Assistant, Prof. Bhiksha Raj, Carnegie Mellon University — Mar 2023-Sep 2023

- Designed and implemented an LLM-generated content detector based on token-prediction signals, including model architectures, training pipeline, and large-scale evaluation across diverse domains.
- Curated and released the *OpenLLMText* dataset (human texts with LLM rewrites), defining data collection, filtering, and quality-control procedures and packaging tools for public use.

Publications

MAC-VO: Metrics-aware Covariance for Learning-based Stereo Visual Odometry — ICRA 2025
Best Conference & Perception Paper
Yuheng Qiu*, ***Yutian Chen****, Zihao Zhang, Wenshan Wang, Sebastian Scherer
https://arxiv.org/abs/2409.09479

Co-Me: Confidence-Guided Token Merging for Visual Geometric Transformers — CVPR 2026
Yutian Chen, Yuheng Qiu, Ruogu Li, Jay Patrikar, Sebastian Scherer
https://arxiv.org/abs/2511.14751

UFM: A Simple Path towards Unified Dense Correspondence with Flow — NeurIPS 2025
Yuchen Zhang, Nikhil Keetha, Chenwei Lyu, Bhuvan Jhamb, ***Yutian Chen***,
Yuheng Qiu, Jay Karhade, Shreyas Jha, Yaoyu Hu, Deva Ramanan,
Sebastian Scherer, Wenshan Wang
https://arxiv.org/abs/2506.09278

AirIO: Learning Inertial Odometry with Enhanced IMU Feature Observability — IEEE RA-L
Yuheng Qiu*, Can Xu*, ***Yutian Chen***, Shibo Zhao, Junyi Geng, Sebastian Scherer

https://arxiv.org/abs/2501.15659

AirIMU: Learning Uncertainty Propagation for Inertial Odometry arXiv Preprint
Yuheng Qiu, Chen Wang, Can Xu, ***Yutian Chen***, Xunfei Zhou, Youjie Xia, Sebastian Scherer
https://arxiv.org/abs/2310.04874

Virtual Community: An Open World for Humans, Robots, and Society ICLR 2026
Qinhong Zhou*, Hongxin Zhang*, Xiangye Lin*, Zheyuan Zhang*, ***Yutian Chen***, Wenjun Liu, Zunzhe Zhang, Sunli Chen, Lixing Fang, Qiushi Lyu, Xinyu Sun, Jincheng Yang, Zeyuan Wang, Bao Chi Dang, Zhehuan Chen, Daksha Ladia, Jiageng Liu, Chuang Gan
https://arxiv.org/abs/2508.14893

Token Prediction as Implicit Classification to Identify LLM-Generated Text EMNLP 2023
Yutian Chen*, Hao Kang*, Vivian Zhai, Liangze Li, Rita Singh, Bhiksha Raj
https://aclanthology.org/2023.emnlp-main.810/

PyPose v0.6: The Imperative Programming Interface for Robotics IROS Workshop, 2023
Zitong Zhan, Xiangfu Li, . . . , ***Yutian Chen***, . . . , Jiajun Wu, . . . , Chen Wang
https://arxiv.org/abs/2309.13035

Myocardial Segmentation of Cardiac MRI Sequences With Temporal Consistency for Coronary Artery Disease Diagnosis Frontier Cardiovascular Medicine, 2022
Yutian Chen, Wen Xie, Jiawei Zhang, Hailong Qiu, Dewen Zeng, Yiyu Shi, Haiyun Yuan, Jian Zhuang, Yanchun Zhang, Yuhao Dong, Meiping Huang, Xiaowei Xu
https://www.frontiersin.org/articles/10.3389/fcvm.2022.804442/full

Teaching

Teaching Assistant, Intro to Machine Learning (PhD), Carnegie Mellon Sep 2023 - Jan 2024
- Led 80-minute recitation lectures, prepared course material and assignments, held office hours.

Teaching Assistant, Principle of Imperative Programming, Carnegie Mellon Sep 2022 - Jan 2023
- Led lab session, developed a virtual machine for C-like program heap visualization on browser environments.

Projects

Perceptual Local-manipulation for Humanoid Oct 2025 - Jan 2026
- Extending the *HDMI: Learning Interactive Humanoid Whole-Body Control from Human Videos* by removing privileged state inputs and replacing them with ego-centric perception with uncertainty from *MAC-VO*.
- Designing a visual state-estimation front-end that infers humanoid pose, contact state, and scene geometry directly from stereo images, feeding these estimates into the whole-body control policy.

NASA Lunar Autonomy Challenge (Team Lead, 3rd Place) Sep 2024 - May 2025
- Managed the full development lifecycle for the NASA Lunar Autonomy Challenge, coordinating efforts across perception, planning, and control teams. Orchestrated the system integration and testing of autonomy stack.
- Led the development of visual-inertial SLAM system with less than 0.1cm/s drift in extraterrestrial environment.

Latch-free Parallel B+ Tree Mar 2024 - Apr 2024
- Designed and implemented an efficient latch-free B+ tree data structure using C++ 17. Utilized SIMD, multiprocessing and CAS operation on x86-64 architecture to achieve 20M read/write QPS on laptop CPU.