

Shi (Billy) Chen

M.S. student, Robotic Systems Development, Robotics Institute, Carnegie Mellon University
Co-founder & CTO, Intuition Core Inc. (intuition.dev)

shic@andrew.cmu.edu | shic@intuition.dev | +1 (341) 258-8001 | +86 166-0210-7753
shichen.world | github.com/BillyChern | linkedin.com/in/shibillychen | ORCID 0009-0003-7365-6983

RESEARCH INTERESTS

Robot learning, mobile manipulation, 3D computer vision, and generative modeling. Current work follows two threads: generative models for 3D perception from sparse, imbalanced sensor data; and vision-language-action policies for industrial manipulation, built at Intuition Core and running 24/7 on KUKA arms in production.

EDUCATION

Carnegie Mellon University, School of Computer Science

Pittsburgh, PA

M.S. in Robotic Systems Development (MRSD), Robotics Institute

Aug. 2026 – May 2028 (expected)

- MRSD team project, funded by Nissan: learning manipulation skills from factory video and transferring them to a Unitree G1.

Fudan University

Shanghai, China

B.S. in Computer Science and Technology, College of Computer Science and Artificial Intelligence

Sept. 2022 – July 2026

- GPA 93/100 (major GPA 3.80/4.0); rank 11 of 91.
- Thesis: *Generative Semantic Scene Completion* (生成式语义场景补全), advised by Prof. Weifeng Ge. Outstanding Undergraduate Thesis Award, Fudan University, 2026.
- Coursework: Data Structures; Principles of Computer Systems; Operating Systems; Computer Organization and Architecture; Computer Networks; Algorithm Design and Analysis; Computer Graphics; Machine Learning and Pattern Recognition; Computer Vision; Natural Language Processing.

University of California, Berkeley

Berkeley, CA

Exchange Student

Jan. – Aug. 2025

- GPA 3.91/4.0. Coursework: Programming Languages and Compilers; Robotic Manipulation and Interaction; Advanced Large Language Model Agents; Practical Applications for Artificial Intelligence.

HONORS, AWARDS AND FUNDING

- 2026 **Young Scientist, 75th Lindau Nobel Laureate Meeting**, Lindau, Germany, 28 June – 3 July 2026. One of about 600 young scientists worldwide and one of ten undergraduates from China, selected through the Lindau program of the Sino-German Center for Research Promotion (NSFC–DFG). Listed in the meeting’s Young Scientist directory.
- 2026 **Shanghai Outstanding Graduate** (上海市优秀毕业生), Shanghai Municipal Education Commission.
- 2026 **Outstanding Undergraduate Thesis Award**, Fudan University, for *Generative Semantic Scene Completion*.
- 2024 **National Natural Science Foundation of China, Youth Student Basic Research Project** (undergraduate track), Grant 624B1006, *LiDAR Signal Processing via Multi-Stage Semantic Modeling and Scene-Structure Generation*, principal investigator. One of about 120 undergraduates funded nationwide in 2024 and one of 19 at Fudan.
- 2024 **First with Distinction** (85%), Online Summer Research Programme, Pembroke College, University of Cambridge.
- 2022–25 **DataGrand Scholarship** (达观数据奖学金), a named award within the Fudan University Scholarship for Academic Achievement: First Class 2024–25; Second Class 2023–24; Third Class 2022–23.
- 2023 **Second Prize** (Shanghai Division, undergraduate group), China Undergraduate Mathematical Contest in Modeling (CUMCM).
- 2023 **Third Prize**, Chinese Mathematics Competitions (CMC), undergraduate division.

PUBLICATIONS

Preprints

- [1] **Shi Chen** and Weifeng Ge. **Generative Semantic Scene Completion**. arXiv:2608.26737, August 2026. Under review. [arXiv](#) | [project page](#) | [code](#).

Manuscripts

- [2] **Shi Chen** and Marwa Mahmoud. **Hi-Fi: A High-Quality Synthetic Hand-Over-Face Gestures Dataset with Multimodal Diffusion**. Manuscript, Online Summer Research Programme, Pembroke College, University of Cambridge, 2024.

Thesis

- [3] **Shi Chen**. **Generative Semantic Scene Completion**. B.S. thesis, Fudan University, May 2026. Advisor: Weifeng Ge. Outstanding Undergraduate Thesis Award.

PATENTS

Three Chinese invention patents, all granted by the China National Intellectual Property Administration (CNIPA) in September 2026. First inventor on each; co-inventor Weifeng Ge; patentee Fudan University.

- [P1] **Shi Chen** and Weifeng Ge. **Method and apparatus for three-dimensional semantic scene completion based on embodied intelligence**. Chinese invention patent ZL 2026 1 0956628.1 (CN 122473413 B), filed 30 June 2026, **granted 1 September 2026**.
- [P2] **Shi Chen** and Weifeng Ge. **LiDAR bird’s-eye-view semantic segmentation method and apparatus for embodied intelligence**. Chinese invention patent ZL 2026 1 0968991.5 (CN 122493464 B), filed 1 July 2026, **granted 1 September 2026**.
- [P3] **Shi Chen** and Weifeng Ge. **LiDAR point-cloud data augmentation method and apparatus for embodied intelligence**. Chinese invention patent ZL 2026 1 0956630.9 (CN 122473414 B), filed 30 June 2026, **granted 22 September 2026**.

RESEARCH EXPERIENCE

Generative Semantic Scene Completion

Fudan University, Shanghai

Project Leader and NSFC Principal Investigator | *Supervisor: Prof. Weifeng Ge*

Sept. 2024 – Aug. 2026

- Recast outdoor LiDAR semantic scene completion as a generative problem, using discrete diffusion in three roles: synthesizing training scenes, completing a scene from noise, and refining a frozen model in one step. Funded by NSFC Grant 624B1006.
- Built **PS³**, an offline paired sparse–dense scene-synthesis pipeline (three cascaded coarse-to-fine multinomial diffusions, a Jensen–Shannon distribution filter, a rare-class object bank, and a virtual HDL-64E ray-tracer), expanding SemanticKITTI training data $2.67\times$ (19,130 real $\rightarrow$ 51,169 pairs).
- Designed **S²D²**, a structured-source discrete-diffusion operator in which the frozen base’s argmax replaces the noise prior, so the network learns only the residual to ground truth, with no base retraining, no distillation and no test-time adaptation.
- One step lifts a frozen SCPNet base to **38.8% mIoU** on the SemanticKITTI hidden test, to our knowledge the best *causal, single-sweep, single-sample* result on that leaderboard, **+2.1 pp** over the previous best published score under that restriction (SCPNet, 36.7). Four correction steps with an eight-view D_4 ensemble reach 39.2% mIoU / 59.0 completion IoU.
- The same one-step correction adds **+1.3 mIoU** to 2D bird’s-eye-view semantic segmentation (34.8 $\rightarrow$ 36.1, val seq 08).
- Outputs: arXiv preprint [1]; three granted patents [P1–P3]; released code, checkpoints and the PS³ dataset (see Software and Datasets).

Mesh and Texture Optimization for 3D Generation

University of Chicago (remote)

Individual Summer Research | *Advisor: Prof. Rana Hanocka*

June – Aug. 2025

- Extended the Continuous Remeshing pipeline to produce high-quality meshes from a single reference image, removing the need for pre-existing goal geometry, via two routes: MV-Adapter multi-view generation with DSINE normal estimation; and OpenLRM’s transformer reconstruction model, rendering multi-view normals from the predicted NeRF tri-plane.
- Introduced vertex-color optimization through a new loss term in the original objective.
- Improved mesh quality over marching cubes applied to the same LRM tri-planes; a crowdsourced user study on Appen showed $>70\%$ preference for our output.

Flying in the Wild, No GPS

University of California, Berkeley

Team Member | Advisor: Prof. Shankar Sastry | Teammates incl. Francesco Crivelli

Apr. – June 2025

- Built a custom quadrotor on the Agilicious framework: 3D-printed frames, Intel RealSense D435i, NVIDIA Jetson Orin Nano, T-Motor F7 controller and Velox V2306 motors.
- Enabled real-time state estimation with SVO Pro over dual infrared cameras and IMU for GPS-free navigation, and built a style-transferred Flightmare/Unity simulator to narrow the Sim2Real gap.
- Deployed the Qwen2.5-VL-3B vision-language model on the Jetson Orin Nano, where it consumes real-time YOLOv8 detections to generate navigation directives for waypoint planning and state-goal estimation; cut hardware cost by several hundred dollars against the Agilicious baseline.
- Documented in the technical report *Flying in the Wild, No GPS*.

Improving Hand Generation Quality in Diffusion Models

Pembroke College, University of Cambridge (online)

Online Summer Research Programme | Advisor: Dr. Marwa Mahmoud

July 2024

- Proposed a multimodal diffusion pipeline integrating ControlNet into Stable Diffusion for hand-over-face gesture generation, raising MediaPipe confidence from 0.248 to 0.556 ($2.24\times$ over baseline).
- Created a synthetic dataset of 170,000+ images across 809 gesture types; crowdsourced Appen validation rated 83.5% of sampled gesture types as “good” (Krippendorff’s $\alpha = 0.61$).
- Awarded First with Distinction (85%); wrote manuscript [2].

Reconstructing Streets and Augmenting Autonomous Driving

Shanghai, China

Researcher | Collaborator: Xingrui Wang (Ph.D. student), Johns Hopkins University

Jan. – July 2024

- Developed a diffusion model in PyTorch generating 3D bounding boxes and BEV perception data on nuScenes; used inpainting to scale scene diversity into a same-sized synthetic dataset.
- Fed synthetic input to MagicDrive for temporally consistent driving video, fine-tuning Stable Diffusion with ControlNet conditioned on 3D boxes and BEV results; reported 2–4% gains on downstream BEV perception and 3D detection with BEVFusion.
- Applied NeRF-based Nerfacto for controllable, editable 3D street-scene rendering.

INDUSTRY AND ENTREPRENEURIAL EXPERIENCE

Intuition Core Inc.

Berkeley & San Francisco, CA | Shenzhen, China

Co-founder & Chief Technology Officer

May 2025 – Present

- Co-founded an AI robotics company building vision-language-action (VLA) models and orchestration middleware for industrial robot arms; lead technical strategy, investor relations and a Shenzhen team of **8 engineers**.
- In partnership with KUKA, put our *intuition-1* VLA model and Intuition Scheduling System (ISS) middleware into production on KUKA arms at Midea’s dishwasher factory (Shunde, China), running in-line quality inspection **24/7 at about 200 rounds per robot per day** on a station that was fully manual before.
- Designed a masked-autoencoder world-model pretraining architecture for self-supervised learning over **1,000 hours** of vision and proprioception data, enabling scalable use of low-quality data through curiosity-based learning.
- Reached **sub-100 ms one-way latency** on transpacific teleoperation links (Kauai–SF, Shanghai–SF) for decentralized data collection; built the teleoperation systems, operator interfaces and control frameworks.
- Secured \$750K in pre-seed funding from Michael Kirsten, a Founders Fund LP investing in a personal capacity.

SOFTWARE AND DATASETS

- **GSSC-S2D2**: reference implementation, configs and released checkpoints for [1]. github.com/BillyChern/GSSC-S2D2; checkpoints on [Hugging Face](#). MIT license. Archived on Zenodo, DOI [10.5281/zenodo.22900923](https://doi.org/10.5281/zenodo.22900923).
- **PS³-SemanticKITTI**: 32,039 synthetic paired sparse–dense LiDAR scenes. [Hugging Face](#) and IEEE DataPort, DOI [10.21227/nqgf-9k39](https://doi.org/10.21227/nqgf-9k39). CC BY-NC-SA 4.0.
- **Base predictions and rare-class object bank** for reproducing every number in [1]. [Hugging Face](#).
- **Stateful-Agent**: cross-platform LLM agent with persistent memory (SQLite + ChromaDB), automated arXiv paper discovery, and a Chrome extension that pre-fills web forms from a research profile. Course project, UC Berkeley (Prof. Dawn Song), 2025. [GitHub](#).

PRESS AND MEDIA COVERAGE

- SciPapermill. *Class Imbalance Conquered: A Deep Dive into Latest AI/ML Breakthroughs*, 30 August 2026. Coverage of [1] and PS³. [link](#).
- College of Computer Science and Artificial Intelligence, Fudan University (WeChat). 走近毕业生 | 陈石：探索、惊奇、热爱，坚持做有意义的事 (*Meet the Graduates | Shi Chen: Exploration, Wonder, Passion, and Doing What Matters*). Feature interview, 6 July 2026. [link](#).
- Sino-German Center for Research Promotion, NSFC. 中德科学中心2026年度林岛项目入选学生颁证仪式暨行前培训会在京举行 (*Certificate Ceremony and Pre-Departure Briefing for the 2026 Lindau Program*), 26 June 2026. Pictured among the selected students. [link](#).
- AC Academic Platform (AC学术平台). 985官宣：19名本科生，获国自然项目！ (*19 Undergraduates Awarded NSFC Projects*), 25 September 2024. Named among Fudan's 19 awardees. [link](#).

OUTREACH AND VOLUNTEER TEACHING

- **Volunteer teacher**, Fudan University rural teaching program, Gaojia Village, Jiangmen Town, Luzhou, Sichuan, July 2023 (one month). Taught science, English and mathematics to primary-school pupils in a formerly nationally designated poverty area.
- **Volunteer tutor**, 2023 – Present. After the trip, began tutoring a local middle-school student one-to-one; she was admitted to a key senior high school in 2025 and the mentoring continues into her second year.

TECHNICAL SKILLS

- **Programming languages**: C, C++, Python, OCaml, Lean, Verilog HDL, Pascal, MATLAB.
- **Frameworks**: PyTorch, PyTorch3D, JAX, OpenCV, NumPy, LangChain, spconv.
- **Robotics and simulation**: ROS/ROS2, MuJoCo, OpenAI Gym, Flightmare.
- **Tools**: Git, Linux, Docker, Blender, Unity, L^AT_EX.

LANGUAGES

- Mandarin Chinese (native); English (fluent). Standardized tests: TOEFL 106 (speaking 25); GRE 336 (quantitative 170).

REFERENCES

Prof. Weifeng Ge

wfge@fudan.edu.cn

Associate Professor, College of Computer Science and Artificial Intelligence, Fudan University.
Thesis supervisor and patent co-inventor.

Dr. Marwa Mahmoud

marwa.mahmoud@cl.cam.ac.uk

Senior Lecturer, School of Computing Science, University of Glasgow; Visiting Fellow, Department of Computer Science and Technology, University of Cambridge.
Advisor, Online Summer Research Programme, Pembroke College, University of Cambridge, 2024.

Prof. Rana Hanocka

ranahanocka@uchicago.edu

Assistant Professor of Computer Science, University of Chicago.
Advisor, summer research, 2025.