# Qiao Gu

Ph.D. Student · University of Toronto

q.gu@mail.utoronto.ca | georgegu1997.github.io | @qiaogu1997

## Education

**University of Toronto** — *Toronto ON, Canada*
Ph.D. in Computer Science — *09/2021 - 09/2026 (expected)*
- GPA: 4.0 / 4.0
- Advisor: Prof. Florian Shkurti
- Research Interest: 3D Computer Vision, Robot Learning, Generative Models

**Carnegie Mellon University** — *Pittsburgh PA, U.S.A.*
M.S. in Robotics — *08/2019 - 08/2021*
- GPA: 4.24 / 4.33
- Advisor: Prof. David Held and Prof. Martial Hebert

**Hong Kong University of Science and Technology** — *Hong Kong S.A.R.*
B.Eng. in Computer Science and Electronic Engineering — *09/2015 - 06/2019*
- GPA: 4.15 / 4.30
- Advisors: Prof. Chi-Keung Tang and Prof. Yu-Wing Tai

## Professional Experience

**Mar 2026 – Aug 2026** | **Research Scientist Intern | NVIDIA, Toronto ON., Canada**
With Jiahui Huang and Sanja Fidler
Video Generation | World Modeling | 3D Reconstruction

**June 2025 – Feb 2026** | **Research Scientist Intern | Meta Reality Labs, Redmond WA, U.S.A.**
With Julian Straub
Video Generation with Spatial Memory | Egocentric World Models

**Aug 2023 – Feb 2024** | **Research Scientist Intern | Meta Reality Labs, Redmond WA, U.S.A.**
With Chris Sweeney and Zhaoyang Lv
Open-world Segmentation | 3D Gaussian Splatting | Egocentric Perception

## Selected Projects

**E$^3$C: Video generation with 3D spatial memory and ego–exo pose control** — *Preprint, First Author (Meta Internship)*
- Proposed a controllable **latent video diffusion** model that renders an explicit **SLAM-based 3D environmental memory** to maintain viewpoint-consistent egocentric generation.
- Augmented the 3D memory with per-point **VAE appearance features** and fused them via a context adapter to reduce texture drift and improve object consistency.
- Implemented **ego–exo pose control** (exo 2D skeletons; ego 3D joints + 6DoF wrists) with **persistent pose tokens** for occlusion-robust motion adherence and controllable scene edits.
- **Set best-reported Nymeria results** with **249 FVD** (*-24%* vs best fine-tuned baseline) and **2.4 cm** camera translation error.

**SAFE: Scalable failure estimation for vision-language-action models** — *NeurIPS 2025, First Author*
- Formulated **multitask failure detection** for VLA policies by learning a task-agnostic failure signal directly from internal model representations.
- Trained lightweight temporal predictors (MLP/LSTM) to output calibrated **failure likelihood** scores with minimal runtime overhead and no policy modifications.

- Applied **functional conformal prediction** to calibrate time-varying thresholds, enabling reliable intervention trade-offs across tasks and embodiments.
- Reached **88.4 ROC-AUC** on *unseen* real-robot tasks (OpenVLA/WidowX), improving *+18.4* over the strongest baseline.

**EgoLifter: Open-world 3D segmentation for egocentric perception** *ECCV 2024, First Author (Meta Internship)*

- Built an **open-world 3D instance segmentation** system on top of **3D Gaussian Splatting**, learning per-Gaussian embeddings for object-level decomposition.
- Introduced a self-supervised **transient prediction** module to filter dynamic regions during reconstruction, reducing floaters and sharpening instance features.
- Lifted **SAM** masks into 3D via contrastive learning, enabling promptable 2D/3D segmentation and object-level extraction for scene editing.
- Introduced EgoLifter and an ADT benchmark; achieved **48.8** cross-view **mIoU** (*+27.3* vs Gaussian Grouping) for open-world 3D segmentation from egocentric video.

**ConceptGraphs: Open-vocabulary 3D scene graphs for perception & planning** *ICRA 2024, Co-first Author*

- Implemented an **object-centric 3D mapping** pipeline that constructs **open-vocabulary scene graphs** for downstream perception and planning.
- Engineered robust **localization** and **map update** mechanisms (re-association, persistence, change handling) to maintain consistent object maps over time.
- Prototyped **object captioning** and ran targeted **segmentation/association** studies to improve object proposals and multi-view consistency.

## Publications

** Indicates equal contribution.*

**E$^3$C: Video Generation with 3D Environmental Memory and Ego-Exo Human Pose Control** Preprint
***Qiao Gu**, Lingni Ma, Adam W Harley, Richard Newcombe, Florian Shkurti, Julian Straub*

**SAFE: Scalable Failure Estimation for Vision-Language-Action Models** NeurIPS 2025
***Qiao Gu**, Yuanliang Ju, Shengxiang Sun, Igor Gilitschenski, Haruki Nishimura, Masha Itkina, Florian Shkurti*
Paper Website Code

**MomaGraph: State-Aware Unified Scene Graphs with Vision-Language Models for Embodied Task Planning** ICLR 2026 (Oral)
*Yuanchen Ju, Yongyuan Liang, Yen-Jen Wang, Gireesh Nandiraju, Yuanliang Ju, Seungjae Lee, **Qiao Gu**, Elvis Hsieh, Furong Huang, Koushil Sreenath*
Paper Website

**SurgSemGS: Semantic and Contrastive Gaussian Splatting for Dynamic Surgical Scene Understanding** Preprint
*Liyuan Liu, **Qiao Gu**, Daniel Elson, Baoru Huang*

**SICNav-Diffusion: Safe and Interactive Crowd Navigation with Diffusion Trajectory Predictions** RA-L & ICRA 2026
*Sepehr Samavi, Anthony Lem, Fumiaki Sato, Sirui Chen, **Qiao Gu**, Keijiro Yano, Angela P Schoellig, Florian Shkurti*
Paper

**EgoLifter: Open-world 3D Segmentation for Egocentric Perception** ECCV 2024
***Qiao Gu**, Zhaoyang Lv, Duncan Frost, Simon Green, Julian Straub, Chris Sweeney*
Paper Website Code

**Aria Everyday Activities Dataset** arXiv 2024
*Zhaoyang Lv, Nicholas Charron, Pierre Moulon, Alexander Gamino, Cheng Peng, Chris Sweeney, Edward Miller, Huixuan Tang, Jeff Meissner, Jing Dong, Kiran Somasundaram, Luis Pesqueira, Mark Schwesinger, Omkar Parkhi,* ***Qiao Gu****, Renzo De Nardi, Shangyi Cheng, Steve Saarinen, Vijay Baiyya, Yuyang Zou, Richard Newcombe, Jakob Julian Engel, Xiaqing Pan, Carl Ren*
Paper Website

**ConceptGraphs: Open-vocabulary 3D Scenegraphs for Perception and Planning** ICRA 2024
***Qiao Gu****, Alihusein Kuwajerwala*, Sacha Morin*, Krishna Murthy Jatavallabhula*, Bipasha Sen, Aditya Agarwal, Corban Rivera, William Paul, Kirsty Ellis, Rama Chellappa, Chuang Gan, Celso Miguel de Melo, Joshua B. Tenenbaum, Antonio Torralba, Florian Shkurti, Liam Paull*
Paper Website Code

**ConceptFusion: Open-set Multimodal 3D Mapping** RSS 2023
*Krishna Murthy Jatavallabhula, Alihusein Kuwajerwala*, ***Qiao Gu****, Mohd Omama*, Tao Chen, Shuang Li, Ganesh Iyer, Soroush Saryazdi, Nikhil Keetha, Ayush Tewari, Joshua B. Tenenbaum, Celso Miguel de Melo, Madhava Krishna, Liam Paull, Florian Shkurti, Antonio Torralba*
Paper Website Code

**Preserving Linear Separability in Continual Learning by Backward Feature Projection** CVPR 2023
***Qiao Gu****, Dongsub Shim, Florian Shkurti*
Paper Code

**OSSID: Online Self-supervised Instance Detection by (and for) Pose Estimation** RA-L & ICRA 2022
***Qiao Gu****, Brian Okorn, David Held*
Paper Website Code

**ZePHyR: Zero-shot Pose Hypothesis Rating** ICRA 2021
*Brian Okorn*, **Qiao Gu***, Martial Hebert, David Held*
Paper Website Code

**Deep Video Matting via Spatio-Temporal Alignment and Aggregation** CVPR 2021
*Yanan Sun, Guanzhi Wang*, **Qiao Gu***, Chi-Keung Tang, Yu-Wing Tai*
Paper

**iQUANT: Interactive Quantitative Investment Using Sparse Regression Factors** EuroVis 2021
*Xuanwu Yue,* ***Qiao Gu****, Deyun Wang, Huamin Qu, Yong Wang*
Paper

**LADN: Local Adversarial Disentangling Network for Facial Makeup and De-Makeup** ICCV 2019
***Qiao Gu****, Guanzhi Wang*, Mang Tik Chiu, Yu-Wing Tai, Chi-Keung Tang*
Paper Website Code

**Characterizing Fluid Response and Sepsis Progression in Emergency Department Patients** EMBC 2019
***Qiao Gu****, Varesh Prasad, Thomas Heldt*
Paper

**PreserVis, a Visual Analytic System for Traffic and Pollution Patterns** VAST 2017
***Qiao Gu****, Hang Yin, Lian Chen, Haotian Li, Chengzhong Liu, Xuanwu Yue, Huamin Qu*
Paper

## Awards, Scholarships, & Grants

2025 **Mary H. Beatty Fellowship** *CAD 10,000*
2025 **University Nomination for Google PhD Fellowship**
2025 **Finalist of Qualcomm Innovation Fellowship**

| | | |
|---|---|---|
| 2025 | **Kwok Sau Po Scholarship** | *CAD 1,750* |
| 2025 | **The Fahiem Bacchus Memorial Graduate Student Scholarship** | *CAD 2,000* |
| 2024 | **NeurIPS 2024 Top Reviewer** | |
| 2024 | **Mary H. Beatty Fellowship** | *CAD 10,000* |
| 2024 | **Ray Reiter Graduate Award** | *CAD 1,000* |
| 2023 | **Acres Productive Technologies Inc./Joseph Yonan Memorial Fellowship** | *CAD 2,000* |
| 2022 | **Ontario Graduate Scholarships** | *CAD 15,000* |
| 2019 | **HKUST Academic Achievement Medal (Top 1%)** | |
| 2017-2019 | **HKSAR Government Scholarships** | *HKD 160,000* |
| 2015-2019 | **Dean's List** | |
| 2018 | **Mr. Armin and Mrs. Lillian Kitchell Undergraduate Research Award** | |
| 2017 | **High Fashion Charitable Foundation Exchange Scholarships** | *HKD 50,000* |
| 2017 | **HKSAR Government Scholarship Fund - Reaching Out Award** | *HKD 10,000* |
| 2017 | **HKSAR Government Scholarship Fund - Talent Development Scholarship** | *HKD 10,000* |
| 2016 | **Simatelex Charitable Foundation Scholarship** | *HKD 20,000* |
| 2016 | **University's Scholarship Scheme for Continuing Undergraduate Students** | *HKD 20,000* |
| 2015 | **University Admission Scholarship** | *HKD 40,000* |

## Invited Talks

03/2024. *ConceptGraphs: Open-Vocabulary 3D Scene Graphs for Perception and Planning*. Boston Dynamic AI institute

02/2024. *Open-World 3d Segmentation with Project Aria*. Meta Reality Labs

10/2023. *ConceptGraphs: Open-Vocabulary 3D Scene Graphs for Perception and Planning*. Meta Reality Labs

## Teaching Experience

| | | |
|---|---|---|
| Fall 2022 | **CSC384: Intro to Artificial Intelligence** | Teaching Assistant |
| Winter 2022 | **CSC384: Intro to Artificial Intelligence** | Teaching Assistant |
| Fall 2021 | **CSC110: Foundations of Computer Science I** | Teaching Assistant |

## Mentorship

| | | |
|---|---|---|
| 2025-present | **Zhengyang Liang** | PhD student, University of Toronto |
| 2024-present | **James Ross** | Master student, University of Toronto |
| 2024-present | **Yuanliang Ju** | PhD student, University of Toronto |
| 2024-2025 | **Owen Sun** | Undergraduate student, University of Toronto |
| 2022-2025 | **Anthony Lem** | Master student, University of Toronto |
| 2023-2024 | **Sirui Chen** | Undergraduate student, University of Toronto |
| 2023 | **Edward Li** | Undergraduate student, University of Toronto |
| 2023 | **Spencer Li** | Undergraduate student, University of Toronto |

## Outreach & Professional Development

### Service and Outreach

| | | |
|---|---|---|
| 2024-2025 | **UofT RI Reading Group** | Organizer |
| 2022 | **Toronto AI in Robotics Seminar** | Organizer |

### Reviewer Service

TRO, TPAMI, IJRR, TMLR, RA-L, NeurIPS, R:SS, CVPR, ICLR, ICCV, ICRA, IROS, ECCV, ICML, WACV

## Professional Memberships

Student member, IEEE