Jun Gao

jungao@cs.toronto.edu ◇ Homepage ◇ Google Scholar ◇ (+1) 437-985-2877

I am interested in computer vision, computer graphics and machine learning. I develop 3D generative AI models to create realistic, high-quality and diverse 3D content for reconstructing, generating and simulating 3D worlds.

EDUCATION

University of Toronto (UofT) Jan. 2020 - May. 2025
Ph.D. Student in Computer Science
Advisor: Sanja Fidler

University of Toronto (UofT) Sept. 2018 - Jan. 2020
Master of Science in Computer Science
Advisor: Sanja Fidler

Peking University (PKU) Sept. 2014 - Jul. 2018
Bachelor of Science (Summa Cum Laude) in Computer Science
Advisor: Liwei Wang

EMPLOYMENT

University of Michigan Aug. 2025 - Present
Assistant Professor Ann Arbor, U.S.

NVIDIA Toronto AI lab Nov. 2019 - Present
Research Scientist, Manager: Sanja Fidler Toronto, Canada

NVIDIA Toronto AI lab Oct. 2018 - Nov. 2019
Research Intern, Manager: Sanja Fidler Toronto, Canada

Microsoft Research Asia Feb. 2018 - May 2018
Research Intern, Manager: Di He Beijing, China

SELECTED AWARDS AND HONORS

2026 ICML Best Paper Honorable Mention

2025 CVPR Best Paper Award Candidate

2023 SIGGRAPH Asia Best Paper Award

2022 NeurIPS Outstanding Reviewer Award

2020 Platform Computing Graduate Fellowship in Computer Science, University of Toronto

2018 Vector Scholarships in Artificial Intelligence, Vector Institute

2018 SUMMA CUM LAUDE, Peking University

2016 Merit Student Award, Peking University

2015 National Scholarship, China

PROFESSIONAL SERVICE

- Program Committee: ACM SIGGRAPH Asia 2024
- Area Chair: Neural Information Processing Systems (NeurIPS) 2023-2026
- Area Chair: The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2026
- Area Chair: International Conference on 3D Vision (3DV) 2025-2026

PUBLICATIONS

*: Authors with equal contribution. ¶: Students/Interns who I (co-)mentored.

1. OSCAR: Omni-Embodiment Action-Conditioned World Model for Robotics
   Zhuoyuan Wu, **Jun Gao**
   Arxiv, 2026, Project, Paper, Code

2. AFUN: Towards an Affordance Foundation Model for Functionality Understanding
   Zhaoning Wang*, Yi Zhong*, Jiawei Fu, Henrik I. Christensen, **Jun Gao**
   Arxiv, 2026, Project, Paper, Code

3. Data-Forcing Distillation: Restoring Diversity and Fidelity in Few-Step Video Generation
   Siyi Chen, Shaowei Liu, Yixuan Jia, Zian Wang, Huan Ling, Qing Qu, **Jun Gao**
   Arxiv, 2026, Project, Paper, Code

4. Modeling Depth Ambiguity: A Mixture-Density Representation for Flying-Point-Free Depth Estimation
   Siyuan Bian*, Congrong Xu*, **Jun Gao**
   Arxiv, 2026, Project, Paper, Code

5. R^3: 3D Reconstruction via Relative Regression
   Congrong Xu, Huachen Gao, Xingyu Chen, Yuliang Xiu, **Jun Gao**, Anpei Chen
   Arxiv, 2026, Project, Paper, Code

6. Lyra 2.0: Explorable Generative 3D Worlds
   Tianchang Shen*, Sherwin Bahmani*, Kai He, Sangeetha Grama Srinivasan, Tianshi Cao, Jiawei Ren, Ruilong Li, Zian Wang, Nicholas Sharp, Zan Gojcic, Sanja Fidler, Jiahui Huang, Huan Ling, **Jun Gao**, Xuanchi Ren*
   ACM Transactions on Graphics (Proc. **SIGGRAPH Asia**), 2026, Project, Paper, Code

7. γ-World: Generative Multi-Agent World Modeling Beyond Two Players
   Fangfu Liu*, Kai He*, Tianchang Shen, Tianshi Cao, Sanja Fidler, Yueqi Duan, **Jun Gao**, Igor Gilitschenski, Zian Wang†, Xuanchi Ren†
   Arxiv, 2026, Project, Paper, Code

8. The Dynamic Prior: Understanding 3D Structures for Casual Dynamic Videos
   Zhuoyuan Wu, Xurui Yang, Jiahui Huang, Yue Wang, **Jun Gao**
   European Conference on Computer Vision (**ECCV**), 2026, Paper

9. MoRight: Motion Control Done Right
   Shaowei Liu, Xuanchi Ren, Tianchang Shen, Huan Ling, Saurabh Gupta, Shenlong Wang, Sanja Fidler, **Jun Gao**
   European Conference on Computer Vision (**ECCV**), 2026, Project, Paper

10. MOTIVE: Motion Attribution for Video Generation
    Xindi Wu, Despoina Paschalidou, **Jun Gao**, Antonio Torralba, Laura Leal-Taixe, Olga Russakovsky, Sanja Fidler, Jonathan Lorraine
    International Conference on Machine Learning (**ICML**), 2026, **(Oral, Best Paper Honorable Mention)**, Project, Paper

11. ChronoEdit: Towards Temporal Reasoning for Image Editing and World Simulation
    Jay Zhangjie Wu*, Xuanchi Ren*, Tianchang Shen, Tianshi Cao, Kai He, Yifan Lu, Ruiyuan Gao, Enze Xie, Shiyi Lan, Jose M. Alvarez, **Jun Gao**, Sanja Fidler, Zian Wang, Huan Ling*
    International Conference on Learning Representations (**ICLR**), 2026, Project, Paper, Code

12. Lyra: Generative 3D Scene Reconstruction via Video Diffusion Model Self-Distillation
    Sherwin Bahmani Tianchang Shen, Jiawei Ren, Jiahui Huang, Yifeng Jiang, Haithem Turki, Andrea Tagliasacchi, David B. Lindell, Zan Gojcic, Sanja Fidler, Huan Ling, **Jun Gao***, Xuanchi Ren*
    International Conference on Learning Representations (**ICLR**), 2026, Project, Paper, Code

13. Cosmos-Drive-Dreams: Scalable Synthetic Driving Data Generation with World Foundation Models
    Xuanchi Ren*,¶, Yifan Lu*, Tianshi Cao*, Ruiyuan Gao, Shengyu Huang, Amirmojtaba Sabour, Tianchang Shen, Tobias Pfaff, Jay Zhangjie Wu, Runjian Chen, Seung Wook Kim, **Jun Gao**, Laura Leal-Taixe, Mike Chen, Sanja Fidler*, Huan Ling*
    Arxiv, 2025, Project, Paper, Code

14. TeraSim-World: Worldwide Safety-Critical Data Synthesis for End-to-End Autonomous Driving
    Jiawei Wang*, Haowei Sun*, Xintao Yan, Shuo Feng, **Jun Gao**, Henry X. Liu
    Arxiv, 2025, Project, Paper, Code

15. PartField: Learning 3D Feature Fields for Part Segmentation and Beyond
Minghua Liu*,¶, Mikaela Angelina Uy*, Donglai Xiang, Hao Su, Sanja Fidler, Nicholas Sharp, **Jun Gao**
International Conference on Computer Vision (**ICCV**), 2025, Project, Paper, Code

16. InfiniCube: Unbounded and Controllable Dynamic 3D Driving Scene Generation with World-Guided Video Models
Yifan Lu*, Xuanchi Ren*,¶, Jiawei Yang, Tianchang Shen, Zhangjie Wu, **Jun Gao**, Yue Wang, Siheng Chen, Mike Chen, Sanja Fidler, Jiahui Huang
International Conference on Computer Vision (**ICCV**), 2025, Project, Paper

17. GEN3C: 3D-Informed World-Consistent Video Generation with Precise Camera Control
Xuanchi Ren*,¶, Tianchang Shen*,¶, Jiahui Huang, Huan Ling, Yifan Lu, Merlin Nimier-David, Thomas Müller, Alexander Keller, Sanja Fidler, **Jun Gao**
Computer Vision and Pattern Recognition (**CVPR**), 2025, **(Highlight)**, Project, Paper, Code

18. Difix3D+: Improving 3D Reconstructions with Single-Step Diffusion Models
Jay Zhangjie Wu*, Yuxuan Zhang*, Haithem Turki, Xuanchi Ren, **Jun Gao**, Mike Zheng Shou, Sanja Fidler, Zan Gojcic†, Huan Ling†
Computer Vision and Pattern Recognition (**CVPR**), 2025, **(Oral, Best Paper Award Candidate)**, Project, Paper, Code

19. DiffusionRenderer: Neural Inverse and Forward Rendering with Video Diffusion Models
Ruofan Liang*, Zan Gojcic, Huan Ling, Jacob Munkberg, Jon Hasselgren, Zhi-Hao Lin, **Jun Gao**, Alexander Keller, Nandita Vijaykumar, Sanja Fidler, Zian Wang*
Computer Vision and Pattern Recognition (**CVPR**), 2025, **(Oral)**, Project, Paper, Code

20. SpaceMesh: A Continuous Representation for Learning Manifold Surface Meshes
Tianchang Shen¶, Zhaoshuo Li, Marc Law, Matan Atzmon, Sanja Fidler, James Lucas, **Jun Gao**, Nicholas Sharp
ACM Transactions on Graphics (Proc. **SIGGRAPH Asia**), 2024, Project, Paper

21. LATTE3D: Large-scale Amortized Text-To-Enhanced 3D Synthesis
Kevin Xie*, Jonathan Lorraine*, Tianshi Cao*, **Jun Gao**, James Lucas, Antonio Torralba, Sanja Fidler, Xiaohui Zeng
European Conference on Computer Vision (**ECCV**), 2024, Project, Paper

22. WildFusion: Learning 3D-Aware Latent Diffusion Models in View Space
Katja Schwarz, Seung Wook Kim, **Jun Gao**, Sanja Fidler, Andreas Geiger, Karsten Kreis
International Conference on Learning Representations (**ICLR**), 2024, Project, Paper

23. Adaptive Shells for Efficient Neural Radiance Field Rendering
Zian Wang*,¶, Tianchang Shen*,¶, Merlin Nimier-David*, Nicholas Sharp, **Jun Gao**, Alexander Keller, Sanja Fidler, Thomas Müller, Zan Gojcic
ACM Transactions on Graphics (Proc. **SIGGRAPH Asia**), 2023, **(Best Paper Award)**, Project, Paper

24. Flexible Isosurface Extraction for Gradient-Based Mesh Optimization
Tianchang Shen¶, Jacob Munkberg, Jon Hasselgren, Kangxue Yin, Zian Wang, Wenzheng Chen, Zan Gojcic, Sanja Fidler, Nicholas Sharp*, **Jun Gao***
ACM Transactions on Graphics (Proc. **SIGGRAPH**), 2023, Project, Paper, Code

25. Progressive Learning of 3D Reconstruction Network from 2D GAN Data
Aysegul Dundar, **Jun Gao**, Andrew Tao, Bryan Catanzaro
IEEE Transactions on Pattern Analysis and Machine Intelligence (**T-PAMI**), 2023, Project, Paper

26. Fine Detailed Texture Learning for 3D Meshes with Generative Models
Aysegul Dundar, **Jun Gao**, Andrew Tao, Bryan Catanzaro
IEEE Transactions on Pattern Analysis and Machine Intelligence (**T-PAMI**), 2023, Project, Paper

27. Neural Fields meet Explicit Geometric Representations for Inverse Rendering of Urban Scenes
Zian Wang¶, Tianchang Shen, **Jun Gao**, Shengyu Huang, Jacob Munkberg, Jon Hasselgren, Zan Gojcic, Wenzheng Chen, Sanja Fidler
Computer Vision and Pattern Recognition (**CVPR**), 2023, Project, Paper, Video

28. Magic3D: High-Resolution Text-to-3D Content Creation
Chen-Hsuan Lin*, **Jun Gao***, Luming Tang*, Towaki Takikawa*, Xiaohui Zeng*, Xun Huang, Karsten Kreis, Sanja Fidler*, Ming-Yu Liu*, Tsung-Yi Lin
Computer Vision and Pattern Recognition (**CVPR**), 2023, **(Highlight)**, Project, Paper

29. GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images
Jun Gao, Tianchang Shen, Zian Wang, Wenzheng Chen, Kangxue Yin, Daiqing Li, Or Litany, Zan Gojcic, Sanja Fidler
Conference on Neural Information Processing Systems (**NeurIPS**), 2022, **(Spotlight)**, Project, Paper, Code

30. Extracting Triangular 3D Models, Materials, and Lighting from Images
Jacob Munkberg, Jon Hasselgren, Tianchang Shen, **Jun Gao**, Wenzheng Chen, Alex Evans, Thomas Müller, Sanja Fidler
Computer Vision and Pattern Recognition (**CVPR**), 2022, **(Oral)**, Project, Paper, Code

31. Improving Semantic Segmentation in Transformers Using Hierarchical Inter-Level Attention
Gary Leung, **Jun Gao**, Xiaohui Zeng, Sanja Fidler
Arxiv, 2022, Paper

32. Deep Marching Tetrahedra: a Hybrid Representation for High-Resolution 3D Shape Synthesis
Tianchang Shen[¶], **Jun Gao**, Kangxue Yin, Ming-Yu Liu, Sanja Fidler
Conference on Neural Information Processing Systems (**NeurIPS**), 2021, Project, Paper, Code, Video

33. DIB-R++: Learning to Disentangle Material from Lighting Using a Deferred Image-based Renderer
Wenzheng Chen, Joey Litalien, **Jun Gao**, Zian Wang, Clement Fuji Tsang, Sameh Khamis, Or Litany, Sanja Fidler
Conference on Neural Information Processing Systems (**NeurIPS**), 2021, Project, Paper

34. 3DStyleNet: Creating 3D Shapes with Geometric and Texture Style Variations
Kangxue Yin, **Jun Gao**, Maria Shugrina, Sameh Khamis, Sanja Fidler
International Conference on Computer Vision (**ICCV**), 2021, **(Oral)**, Project, Paper

35. DatasetGAN: Efficient Labeled Data Factory with Minimal Human Effort
Yuxuan Zhang*[,¶], Huan Ling*, **Jun Gao**, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, Sanja Fidler
Computer Vision and Pattern Recognition (**CVPR**), 2021, **(Oral)**, Project, Paper, Code

36. Image GANs meet Differentiable Rendering for Inverse Graphics and Interpretable 3D Neural Rendering
Yuxuan Zhang*[,¶], Wenzheng Chen*, Huan Ling, **Jun Gao**, Yinan Zhang, Antonio Torralba, Sanja Fidler
International Conference on Learning Representations (**ICLR**), 2021, **(Oral)**, Project, Paper

37. Learning Deformable Tetrahedral Meshes for 3D Reconstruction
Jun Gao, Wenzheng Chen, Tommy Xiang, Alec Jacobson, Morgan McGuire, Sanja Fidler
Conference on Neural Information Processing Systems (**NeurIPS**), 2020, Project, Paper, Code

38. Beyond Fixed Grid: Learning Geometric Image Representation with a Deformable Grid
Jun Gao, Zian Wang, Jinchen Xuan, Sanja Fidler
European Conference on Computer Vision (**ECCV**), 2020, Project, Paper, Code

39. Interactive Annotation of 3D Object Geometry using 2D Scribbles
Frank Shen*[,¶], **Jun Gao***, Amlan Kar, Sanja Fidler
European Conference on Computer Vision (**ECCV**), 2020, Project, Paper, Video

40. ScribbleBox: Interactive Annotation Framework for Video Object Segmentation
Bowen Chen*, Huan Ling*, Xiaohui Zeng, **Jun Gao**, Ziyue Xu, Sanja Fidler
European Conference on Computer Vision (**ECCV**), 2020, Project, Paper

41. Learning to Predict 3D Objects with an Interpolation-based Differentiable Renderer
Wenzheng Chen, **Jun Gao***, Huan Ling*, Edward J. Smith*, Jaakko Lehtinen, Alec Jacobson, Sanja Fidler
Conference on Neural Information Processing Systems (**NeurIPS**), 2019, Project, Paper, Code

42. Fast Interactive Object Annotation with Curve-GCN
Huan Ling*, **Jun Gao***, Amlan Kar, Wenzheng Chen, Sanja Fidler
Computer Vision and Pattern Recognition **(CVPR)**, 2019, Paper, Code, Video

43. Representation Problem in Training Natural Language Generation Models
Jun Gao*, Di He*, Xu Tan, Tao Qin, Liwei Wang, Tie-Yan Liu
International Conference on Learning Representations (**ICLR**), 2019, Paper, OpenReview

44. Mimicking the in-camera color pipeline for camera-aware object compositing
Jun Gao, Xiao Li, Liwei Wang, Sanja Fidler, Stephen Lin
Arxiv, 2019, Paper

45. Deepspline: Data-Driven Reconstruction of Parametric Curves and Surfaces
Jun Gao, Chengcheng Tang, Vignesh Ganapathi-Subramanian, Jiahui Huang, Hao Su, Leonidas J Guibas
Arxiv, 2019, Paper

46. DeepPrimitive: Image Decomposition by Layered Primitive Detection
Jiahui Huang, **Jun Gao**, V. G.Subramanian, Hao Su, Yin Liu, Chengcheng Tang, Shi-Min Hu, Leonidas J. Guibas
Computational Visual Media (**CVM**), 2018, Paper

47. Learning to Navigate for Fine-grained Classification
Ze Yang, Tiange Luo, Dong Wang, Zhiqiang Hu, **Jun Gao**, Liwei Wang
European Conference on Computer Vision (**ECCV**), 2018, Paper, Code

48. Dropout Training, Data-dependent Regularization and Excess Risks
Wenlong Mou, Yuchen Zhou, **Jun Gao**, Liwei Wang
International Conference on Machine Learning (**ICML**), 2018, Paper

TEACHING EXPERIENCE

Teaching Assistant CSC 420: Introduction to Image Understanding at UofT, website — Winter 2022

Guest Lecturer CSC 420: Introduction to Image Understanding at UofT, website — Winter 2022 & 2023

Guest Lecturer CS 479: Machine Learning for 3D Data at KAIST, website — Fall 2023

Guest Lecturer CSCI 677: Advanced Computer Vision at USC — Fall 2023

MENTORSHIPS

Siyi Chen, Ph.D. student at the University of Michigan, 2026-now

Siyuan Bian, Ph.D. student at the University of Michigan, 2025-now

Zhaoning Wang, Ph.D. student at the University of Michigan, 2025-now

Xuanchi Ren, Ph.D. student at the University of Toronto, 2024-now

Tianchang Shen, Ph.D. student at the University of Toronto, 2019-now

Weiwei Sun, Ph.D. student at the University of British Columbia, 2022

Gary Leung, Master student at the University of Toronto, 2020

Jinchen Xuan, Undergrad student at the Peking University, 2019

Yuxuan Zhang, Undergrad student at the University of Waterloo, 2020

Yinan Zhang, Undergrad student at the University of Waterloo, 2019

INVITED TALKS

- *Towards a 3D-Grounded Generative World Simulator*
  - ICCV Wild3D Workshop — Oct. 2025
- *3D Grounding from and for Video Generative Models*
  - Stanford University — April 2025
- *Generative AI for 3D Content Creation*
  - Stanford University — April 2024
  - Purdue University — April 2024
  - University of Michigan — Mar. 2024
  - University of Waterloo — Mar. 2024

- *Revisit Representation, Data and Algorithm for Scaling Up 3D Content Generation*
  - Cornell University: Noah Snavely's research group — Sept. 2023
  - MIT: Vision and Graphics Seminar — Sept. 2023
  - Max Planck Institute for Intelligent Systems: Michael Black's research group — Aug. 2023
  - Waabi: Raquel Urtasun's research group — Aug. 2023
- *Towards High-Quality 3D Content Creation with a Hybrid Representation*
  - BIRS Workshop on 3D Generative Models — July 2023
- *Machine Learning for 3D Content Generation*
  - CVPR 2023 Workshop: Structural and Compositional Learning on 3D Data, recording — June 2023
  - Tsinghua University: Li Yi's research group — May 2023
- *Towards Generative Modeling of 3D Objects Learned from Images*
  - Johns Hopkins University: Alan Yuille's research group — Feb. 2023
  - University of Toronto: Toronto AI in Robotics Seminar, recording — Nov. 2022
  - University of Oxford: Visual Geometry Group — Sept. 2022
  - Peking University: Baoquan Chen's research group — Sept. 2022
- *Learning Geometric Representation for Computer Vision*
  - GAMES: Graphics And Mixed Environment Symposium — Nov. 2020
  - University of Alberta: Deep Learning Seminar — May 2020

PATENTS

1. Extracting Triangular 3D Models, Materials, and Lighting from Images
   Carl Jacob Munkberg, Jon Niklas Theodor Hasselgren, Tianchang Shen, **Jun Gao**, Wenzheng Chen, Alex John Bauld Evans, Thomas Müller-Höhne, Sanja Fidler
   U.S. Patent Application 17/827,918
2. Transferring Geometric and Texture Styles in 3D Asset Rendering Using Neural Networks
   Kangxue Yin, **Jun Gao**, Masha Shugrina, Sameh Khamis, Sanja Fidler
   U.S. Patent Application 17/467,792
3. Synthesizing High Resolution 3D Shapes from Lower Resolution Representations for Synthetic Data Generation Systems and Applications
   Tianchang Shen, **Jun Gao**, Kangxue Yin, Ming-Yu Liu, Sanja Fidler
   U.S. Patent Application 17/718,172
4. Hybrid Differentiable Rendering for Light Transport Simulation Systems and Applications
   Wenzheng Chen, Joey Litalien, **Jun Gao**, Zian Wang, Clement Tse Tsian Christophe Louis Fuji, Sameh Khamis, Or Litany, Sanja Fidler
   U.S. Patent Application 17/826,611
5. Neural Rendering for Inverse Graphics Generation
   Wenzheng Chen, Yuxuan Zhang, Sanja Fidler, Huan Ling, **Jun Gao**, Antonio Torralba Barriuso
   U.S. Patent Application 17/981,770
6. Generating Labels for Synthetic Images Using One or More Neural Networks
   Yuxuan Zhang, Huan Ling, **Jun Gao**, Wenzheng Chen, Antonio Torralba Barriuso, Sanja Fidler
   U.S. Patent Application 17/020,649
7. Systems and Methods for Polygon Object Annotation and a Method of Training an Object Annotation System
   Sanja Fidler, Amlan Kar, Huan Ling, **Jun Gao**, Wenzheng Chen, David Jesus Acuna Marrero
   U.S. Patent Application 11/556,797

SELECTED PRESS COVERAGE

- **GET3D: A Generative Model of High Quality 3D Textured Shapes Learned from Images**

  Two Minute Papers: NVIDIA's New AI: Generating 3D Models!

  Yahoo: NVIDIA's new AI model quickly generates objects and characters for virtual worlds

  Engadget: NVIDIA's new AI model quickly generates objects and characters for virtual worlds

  VentureBeat: Nvidia CEO Jensen Huang says AI will autopopulate the 3D imagery of the metaverse

  NVIDIA: World-Class: NVIDIA Research Builds AI Model to Populate Virtual Worlds With 3D Objects, Characters

- **Magic3D: High-Resolution Text-to-3D Content Creation**

  Two Minute Papers: NVIDIA's New AI: Wow, 8x Better Text To 3D!

  Forbes: What Nvidia's New Text-To-3D Means For Engineering & Product Design

  Ars Technica: 3D for everyone? Nvidia's Magic3D can generate 3D models from text

- **FlexiCubes: Flexible Isosurface Extraction for Gradient-Based Mesh Optimization**

  VectureBeat: NVIDIA's FlexiCubes uses generative AI to create 3D meshes

  Two Minute Papers: NVIDIA's New AI: 20% Faster Game Graphics!

  NVIDIA: Better 3D Meshes, from Reconstruction to Generative AI

- **NVDiffrec: Extracting Triangular 3D Models, Materials, and Lighting From Images**

  Two Minute Papers: NVIDIA's New AI Grows Objects Out Of Nothing!

  NVIDIA: AI in the Big Easy: NVIDIA Research Lets Content Creators Improvise With 3D Objects

- **DIB-R: Learning to Predict 3D Objects with an Interpolation-based Differentiable Renderer**

  Two Minute Papers: This Neural Network Creates 3D Objects From Your Photos