

Ye Yuan

Staff Research Scientist & Tech Lead, NVIDIA Research

ROBOT LEARNING / HUMANOIDS / EMBODIED AI

yyuan2@cs.cmu.edu • ye-yuan.com • Google Scholar • GitHub

Research profile

Research scientist and technical leader with extensive experience in reinforcement learning and robotics. Develops humanoid learning systems that connect reinforcement learning, visuomotor control, human data, and generative models. Research spans foundational algorithms, data collection and generation, large-scale training, and real-robot evaluation. Record includes 47 accepted or published conference/journal papers, nine oral/highlight/spotlight selections, and a SIGGRAPH Best Paper Honorable Mention.

Selected robotics research contributions

Scalable reinforcement learning and whole-body control

Co-first author of *SONIC* (*Science Robotics*, 2026): scaled motion tracking using 100M+ motion frames, with a shared control interface for teleoperation and vision-language-action models. SONIC is widely adopted by robotics companies and researchers. Extensive experience with System 1 and System 0 UMI policy training and whole-body teleoperation.

Visuomotor learning and sim-to-real transfer

Contributed to *VIRAL* (CVPR 2026), scaling visual teacher-student training to 64 GPUs and demonstrating 54 continuous zero-shot loco-manipulation cycles on a physical Unitree G1. Related [visual RL research](#) studies active perception and dexterity in simulated humanoids. First-author work (ECCV 2018; ICCV 2019) connected egocentric vision to physics-based humanoid control through imitation learning and RL.

Learning robot skills from human video and world models

Project lead for *GRAIL* (CoRL 2026 Oral), using video foundation models in fully specified 3D scenes to generate 20K+ humanoid demonstrations. Egocentric RGB policies trained solely on GRAIL-generated data achieved 84% pick-up success on seen objects, 80% on unseen objects, and 90% stair-climbing success on a real Unitree G1.

Generative motion and 3D foundations

Main model architect for *Kimodo*, a controllable motion diffusion model trained on 700 hours of motion capture. First author of *PhysDiff* (ICCV 2023 Oral) and last author of *GEM* (formerly GENMO) (ICCV 2025 Highlight). Co-first author of *GAavatar*; additional 3D generation work includes *AGG* and *AdaHuman*.

Research experience

NVIDIA Research

May 2022 – Present

Staff Research Scientist & Tech Lead, DAIR and LPR

Lead research connecting human data, simulation, and generative models to humanoid control and visual policies. Guide collaborations from problem formulation and model design through training, open-source release, and real-robot evaluation.

Carnegie Mellon University, Robotics Institute

Aug 2017 – May 2022

Ph.D. Research Assistant; advisor: Kris Kitani

Developed reinforcement-learning methods for humanoid control (*Residual Force Control*, NeurIPS 2020), visuomotor control, and automated robot design algorithms (*Transform2Act*, ICLR 2022 Oral). Mentored graduate and undergraduate researchers.

Honors

SIGGRAPH Best Paper Honorable Mention	2023
NVIDIA Graduate Fellowship (1 of 5 selected worldwide)	2021–2022
Qualcomm Innovation Fellowship (\$100K award)	2020–2021
Outstanding Reviewer, ECCV / ICLR	2022 / 2021

Selected publications

* Co-first/equal contribution; † project lead. Complete record: full research CV.

Humanoid control and reinforcement learning

1. SONIC: Supersizing Motion Tracking for Natural Humanoid Whole-Body Control
Z. Luo*, **Y. Yuan***, T. Wang*, et al. | Science Robotics 2026
2. Transform2Act: Learning a Transform-and-Control Policy for Efficient Agent Design
Y. Yuan, Y. Song, Z. Luo, W. Sun, K. Kitani | ICLR 2022; Oral
3. Residual Force Control for Agile Human Behavior Imitation and Extended Motion Synthesis
Y. Yuan, K. Kitani | NeurIPS 2020

Visuomotor learning and sim-to-real

4. VIRAL: Visual Sim-to-Real at Scale for Humanoid Loco-Manipulation
T. He*, Z. Wang*, H. Xue*, ..., **Y. Yuan**, et al. | CVPR 2026
5. Emergent Active Perception and Dexterity of Simulated Humanoids from Visual Reinforcement Learning
Z. Luo, C. Tessler, T. Lin, **Y. Yuan**, et al. | arXiv 2025
6. Ego-Pose Estimation and Forecasting as Real-Time PD Control
Y. Yuan, K. Kitani | ICCV 2019
7. 3D Ego-Pose Estimation via Imitation Learning
Y. Yuan, K. Kitani | ECCV 2018

Human data for robot learning

8. GRAIL: Generating Humanoid Loco-Manipulation from World Models
T. Xie*, H. Zhang*, J. Park*, ..., **Y. Yuan†** | CoRL (accepted) 2026; Oral
9. GLAMR: Global Occlusion-Aware Human Mesh Recovery with Dynamic Cameras
Y. Yuan, U. Iqbal, P. Molchanov, K. Kitani, J. Kautz | CVPR 2022; Oral
10. SimPoE: Simulated Character Control for 3D Human Pose Estimation
Y. Yuan, S.-E. Wei, T. Simon, K. Kitani, J. Saragih | CVPR 2021; Oral

Generative motion and 3D modeling

11. Kimodo: Scaling Controllable Human Motion Generation
D. Rempe*, M. Petrovich*, **Y. Yuan**, et al. | arXiv 2026
12. AdaHuman: Animatable Detailed 3D Human Generation with Compositional Multiview Diffusion
Y. Huang, **Y. Yuan**, X. Li, J. Kautz, U. Iqbal | ICCV 2025
13. GAvatar: Animatable 3D Gaussian Avatars with Implicit Mesh Learning
Y. Yuan*, X. Li*, Y. Huang, S. De Mello, K. Nagano, J. Kautz, U. Iqbal | CVPR 2024; Highlight
14. PhysDiff: Physics-Guided Human Motion Diffusion Model
Y. Yuan, J. Song, U. Iqbal, A. Vahdat, J. Kautz | ICCV 2023; Oral

Education

Carnegie Mellon University , Ph.D. Robotics; advisor: Kris Kitani	2022
Carnegie Mellon University , M.S. Computer Science; advisor: Stelian Coros	2016
Zhejiang University , B.E. Computer Science and Technology	2015

Research service

Area Chair: CVPR; **Reviewer:** NeurIPS, ICML, ICLR, CVPR, ICCV, ECCV, ICRA, SIGGRAPH, JMLR, TMLR, TPAMI, TIP, RA-L.

Research software and systems

Project releases (GitHub stars): SONIC / GR00T-WBC (3,773), GRAIL (575), GEM/GENMO (528), GEM-X (484).

Personal projects: PyTorch-RL (1,286), AgentFormer (317), RFC (169), DLow (125). As of October 17, 2026.

Systems: Python, PyTorch, distributed GPU training, Isaac Gym / Isaac Sim, MuJoCo, Unitree G1.