

Minghao Liu

San Jose, CA | (814) 384-2833 | mh.liu.cs@gmail.com | minghao.a.liu@accenture.com | [LinkedIn](#) | [Webpage](#) | [Scholar](#)

SUMMARY

Research scientist and manager working at the intersection of **generative AI** and **physics-grounded simulation** — controllable, physics-constrained video and 3D generation, world models, and diffusion-based content creation — with experience across Accenture, Meta Reality Labs, and ByteDance and publications at NeurIPS, ICML, EGSR, SIGGRAPH Asia, ECCV, and CVPR. Also active in agentic systems and LLM reasoning.

EDUCATION

University of California, Santa Cruz

Ph.D. in Computer Science and Engineering

Sep 2019 — June 2025

University of California, Santa Cruz (UCSC)

B.S. in Computer Science and B.A. in Business Management Economics

Sep 2015 — June 2019

INDUSTRY EXPERIENCE

Accenture @ Advanced AI

Research Scientist Manager

Mountain View, CA

June 2025 — Present

- **Leadership:** Lead a cross-functional team of ~8 scientists, engineers, and analysts across research and client engagements; serve as Research Committee Chair for cross-team research projects.
- **Workflow agents:** Led design and implementation of an internal conversational-agent platform for mid-market clients, using structured workflows and validation to support protocol adherence and safety guardrails alongside task completion; delivered a research prototype and client demos.
- **Self-evolving agents:** Advance agent self-evolution through three complementary directions: (1) a **reward-free framework** using frozen-LLM pairwise validators instead of scalar rewards, matching or exceeding full-reward baselines on prompt- and code-evolution tasks (ICLR submission); (2) a **Jev decision-making router** alongside an LLM agent that selects promising branches in a tree of ideas and iterative experiments to accelerate long-horizon development tasks such as MLE-bench; and (3) **reinforcement learning** to improve agents' task-solving capabilities.
- **LLM reasoning:** Proposed **Confidence Dynamic Gain (CDG)**, a training-free inference-time voting method that lifts reasoning-LLM accuracy by ~4.6% across four frontier open-source LLMs at zero fine-tuning cost. (ICML 2026)
- **Video generation:** Built a controllable car-crash scene-generation pipeline (diffusion models with physics-based constraints) for an insurance-claims client; led technical due diligence on video-generation platforms for a [Netflix acquisition](#) (model quality, IP infringement risk, algorithmic novelty, data sourcing/licensing).
- **World model:** Initiated an ongoing internal research exploration, *Agentic 3D World Model via Video Generation Models*, repurposing video-generation models toward agent-controllable 3D world simulation for scenarios such as autonomous driving.
- **Simulation:** Built simulation infrastructure with synthetic data generation and Vision-Language-Action (VLA) integration for agent training.

Meta @ Reality Lab

Research Scientist Intern (two consecutive internships)

Sausalito, CA

June 2023 — Feb 2025

- **Method:** Proposed a biophysics-bound, latent-diffusion-based generative AI for 3D human avatar generation, grounded in direct physical measurements of skin (e.g., melanin) rather than broad annotations. (EGSR 2025)
- **Rendering:** Integrated the generative model with a physics-based differentiable renderer, achieving cinematic-quality output and demonstrating expertise in traditional 3D asset-based rendering and realistic sensor output generation.
- **Infrastructure:** Trained on Meta's DGX (8×A100 across 10 nodes) with parallel experiments and distributed batch inference; combined diffusion with a refiner to produce 2K-resolution albedo, bump map, and 3D mesh.
- **Data strategy:** Combined a mid-scale capture-in-the-wild dataset with a small high-quality light-stage capture; resulting models remain expressive across identities and substantially reduce reliance on expensive traditional capture.

- **Cartoon avatar auto-creation:** Proposed a self-supervised method that generates editable, fine-grained cartoon avatars from input images; used GAN-based stylization to bridge domain gaps and neural imitators to address non-differentiability in the rendering system. (SIGGRAPH Asia 2022)
- **Tag-based human annotation:** Developed a novel tagging-based annotation methodology for human-cartoon labels that improves quality, reduces annotator confusion, and generalizes to other rendering systems at limited cost; validated through user studies with internal annotation teams and MTurk, designed with clean comparison interfaces for fair and consistent evaluation. (NeSy 2025)

PUBLICATIONS

Generative AI & Agents

- Yu Wang, **Minghao Liu**, Jiayun Wang, Jinrui Huang, Ankit Shah, Wei Wei. *Inference Time Optimization with Confidence Dynamics*. (ICML 2026) (Accenture)
- **Minghao Liu**, Yu Wang, Jiayu Wang, Wei Wei. *Reward-Free Evolving Agents via Pairwise Validator*. (ICLR submission) (Accenture)
- Diji Yang*, **Minghao Liu***, Chung-Hsiang Lo, Yi Zhang, James Davis. *GenIR: Generative Visual Feedback for Mental Image Retrieval*. (NeurIPS 2025) (Co-first author)

Vision Graphics / Digital Human

- **Minghao Liu**, Stéphane Grabli, Sébastien Speierer, Nikolaos Sarafianos, Lukas Bode, Matt Chiang, Christophe Hery, James Davis, Carlos Aliaga. *Controllable Biophysical Human Faces*. (EGSR 2025)
- Shen Sang, Tiancheng Zhi, Guoxian Song, **Minghao Liu**, Chunpong Lai, Jing Liu, Linjie Luo, Xiang Wen, James Davis. *AgileAvatar: Stylized 3D Avatar Creation via Cascaded Domain Bridging*. (SIGGRAPH Asia 2022)
- Jiaheng Wei*, **Minghao Liu***, Jiahao Luo, Qitong Li, James Davis, Yang Liu. *DuelGAN: A Duel Between Two Discriminators Stabilizes the GAN Training*. (ECCV 2022) (Co-first author)
- Jiahao Luo, Fahim Khan, Issei Mori, Akila de Silva, Eric Sandoval Ruezga, **Minghao Liu**, Alex Pang, James Davis. *How much does input data type impact final face model accuracy?* (CVPR 2022 Oral)
- **Minghao Liu**, Jiahao Luo, Xiaohan Zhang, Yang Liu, James Davis. *Low-light Image Enhancement Using Chain-consistent Adversarial Networks*. (ICPR 2022)
- Ankit Parag Shah, **Minghao Liu**, Jinrui Huang, Bhavya Pranav Tandra, Yu Wang, Sakshi Agarwal, Paul Wu, Aishik Konwer, Rhett Rozga, Selli P. Kondamuri, Alper Halbutogullari, Qin Zhang, Wei Wei. *Physical AI: The Next Frontier in AI and Robotics to Build Truly Autonomous Machines*. (Preprint, 2026) (Accenture)

Human-AI

- **Minghao Liu***, Jiaheng Wei*, Yang Liu, James Davis. *Do humans and machines have the same eyes? Human-machine perceptual differences on image classification*. (AAAI 2025 Oral)
- **Minghao Liu**, Zeyu Cheng, Shen Sang, Jing Liu, James Davis. *Tag-based annotation creates better avatars*. (NeSy 2025)
- **Minghao Liu***, Zonglin Di*, Jiaheng Wei*, Zhongruo Wang, Hengxiang Zhang, Ruixuan Xiao, Haoyu Wang, Jinlong Pang, Hao Chen, Ankit Shah, Hongxin Wei, Xinlei He, Zhaowei Zhao, Haobo Wang, Lei Feng, Jindong Wang, James Davis, Yang Liu. *Automatic Dataset Construction (ADC): Sample Collection, Data Curation, and Beyond*. (Under review) (Across institutions: UCSC, CMU, Amazon, MS, and more)
- Vanshika Vats, Marzia Binta Nizam, **Minghao Liu**, Ziyuan Wang, Richard Ho, Mohnish Sai Prasad, Vincent Titterton, Sai Venkat Malreddy, Riya Aggarwal, Yanwen Xu, Lei Ding, Jay Mehta, Nathan Grinnell, Li Liu, Sijia Zhong, Devanathan Nallur Gandamani, Xinyi Tang, Rohan Ghosalkar, Celeste Shen, Rachel Shen, Nafisa Hussain, Kesav Ravichandran, James Davis. *A Survey on Human-AI Teaming with Large Pre-Trained Models*.

ACADEMIA EXPERIENCE

Research Assistant at VIS Lab

Sep 2019 — June 2025

UC Santa Cruz | PI: James Davis, Alex Pang, and Yang Liu

Santa Cruz, CA

- **Autonomous-vehicle perception:** Contributed to AV perception projects with RGB and LiDAR processing for object detection and scene understanding; sensor fusion and 3D point-cloud processing for autonomous systems.
- **Human-AI for AV detection:** Explored how human feedback combines with AI predictions to improve detection accuracy and handle edge cases in autonomous-driving scenarios.
- **Marine ML collaborations (non-author contributor):** Contributed to two VIS Lab marine-vision projects — *RipScout: Realtime ML-Assisted Rip Current Detection and Automated Data Collection Using UAVs* (Khan et al., IEEE J-STARS 2025) and *RipFinder: Real-Time Rip Current Detection on Mobile Devices* (Khan et al., Frontiers in Marine Science 2025).
- **Lab leadership:** Ran weekly group meetings (agenda-setting, cross-project discussion) and organized a paper-reading and brainstorming Slack channel that cultivated reading habits across the lab.

Mentor for Undergraduate Research

Jan 2023 — Dec 2023

UC Santa Cruz | PI: Dr. James Davis

Santa Cruz, CA

- **LLM math reasoning:** Assessed ChatGPT's mathematical capabilities and limitations (arXiv).
- **Annotation methodology:** Improved human annotation for avatar creation.
- **Outcomes:** Two arXiv reports and one conference submission; drew VC interest.

Reviewer experience for journal or conference:

- AI: NeurIPS 2026, AAAI 2025, ICLR 2025, ICML 2025, WWW 2025, AAAI 2024, ICLR 2024, ICML 2024, AAAI 2023, ICLR 2023, NeurIPS 2023, WWW 2023
- Computer Vision: CVPR 2025, CVPR 2024, ECCV 2024, CVPR 2023, ICCV-CLVL 2023

Teaching Assistant — UC Santa Cruz

Sep 2020 — June 2024

- AI/ML courses: *CSE240*: Artificial Intelligence (Graduate course); *CSE142*: Machine Learning and Data Mining ($\times 2$); *CSE144*: Applied Machine Learning — Deep Learning ($\times 2$)
- Graphics: *CSE160*: Introduction to Computer Graphics ($\times 3$)
- Computer Science / Math: Data Structures and Algorithms; Programming in Python; Applied Discrete Mathematics

AWARDS

- **Dissertation-Year Fellowship** recipient for 2024–25 at the University of California, Santa Cruz (3 per year).
- National Interest Waiver (**NIW**) Approval for green card application, considered as an outstanding researcher by USCIS.