

Jingyu Shi

shi537@purdue.edu | +1-571-3153892 | Personal Website | LinkedIn | Google Scholar

Summary & Objective

Ph.D. Candidate in Computer Engineering with a focus on applied **Machine Learning (ML) / Deep Learning (DL)** methodologies to solving **Computer Vision (CV)** challenges in real-world applications, specifically in videos, **Augmented Reality (AR)**, and **Virtual Reality (VR)**, with hands-on experience with **diffusion models**, **Neural Radiance Field (NeRF)**, **action recognition**, **object/hand tracking** etc.

Education

Ph.D. in Computer Engineering <i>Purdue University</i>	Jan. 2021 – expt. Dec. 2025 West Lafayette, IN
M.S. in Computer Engineering <i>Georgia Institute of Technology</i>	Sept. 2019 – Dec. 2020 Atlanta, GA
B.Eng. in Instrument Science and Technology <i>Beihang University</i>	Sept. 2015 – Jun. 2019 Beijing, China

Selected Projects

Occlusion-Aware NeRF Inpainting via Score Distillation Sampling from Diffusion Models 2024

- Designed and implemented collaborative *score distillation sampling* workflow to inpaint NeRF scenes with occlusion awareness and applied a Grid-based denoising pattern to enhance the distillation consistency.
- Collected a challenging dataset for NeRF inpainting with consistently annotated RGB, Depth, and mask images.
- Applied *LoRA to fine-tune Stable Diffusion 2* for more realistic per-scene inpainting.
- Achieved the best performance in consistency and faithfulness on common datasets, compared against SOTA baselines

Diverse Text-to-Image Generation via Bimodal Classifier-Free Guidance 2024

- Implemented *Bimodal Classifier-Free Guidance for diffusion models* to incorporate the (demographic and aesthetic) diversity of the generated images without compromising fidelity.
- Adapted Decoupled Cross-Attention and applied *LoRA to fine-tune bimodal-conditioned diffusion models*
- On human image generation tasks in MSCOCO and OpenImages-v6, achieved the best diversity and image quality performance while maintaining good fidelity and alignment compared with SOTA baselines.

Context-Aware Motion Generation for AR Instructions using Diffusion Models 2023

- Designed and developed CARING-AI, a system that leverages *Motion Diffusion Model* (MDM) to create context-aware humanoid avatar animations for AR instructions.
- Modified MDM to enable conditioning on trajectory data and incorporated a temporal smoothing algorithm to ensure continuous transitions among generated animations.
- Finetuned and evaluated the models on a subset of the HumanML3D dataset, achieving 80% less transition error and 11% less average distance between generated avatars and keypoint conditions.

Consistent Hand-Object Interaction Rendering in AR 2023

- Designed and developed Ubi-TOUCH, a system that enables real-time rendering of hand-object interaction (HOI) in AR to transfer real-world HOIs into virtual rendering with proxy objects.
- Set up a five-camera system to *collect real-world HOI data* with hand-tracking and object pose estimation annotations.
- Fine-tuned and deployed pretrained *hand-tracking* model with 6% reduced MPJPE and *object pose estimation* model with 3% reduced rotation error and 29% reduced translation error for real-time inference in AR.
- Implemented a joint HOI optimization algorithm to render the virtual hands and objects plausibly in AR.

Object-Object Interaction for richer dynamic scene representations in videos 2022

- Collected and annotated a novel video dataset for object-object interactions (OOI), and *benchmarked the dataset* for tasks including OOI classification, Scene-Graph generation, and human-object interaction classification.
- Developed a web-based UI for video data annotation, and recruited and trained annotators for high-quality annotation of the collected video data.
- Using the I3D backbone, *Vision-Transformer*, and *Segment-Anything-Model*, extracted and engineered the features into descriptors of motion relations, location relations, and contacted relations for the OOI-based downstream tasks.

Professional Experience

Research Intern May 2024 – Dec. 2024 *Futurewei Technologies Inc.* Santa Clara, CA

- Conducted research on SOTA generative AI algorithms such as diffusion models for CV and CG tasks
- Spearheaded a research team of three to conduct independent research, contributed to multiple research projects, and delivered **3 CVPR submissions** and **1 patent application**.

Software Engineer Intern Jan. 2018 – Mar. 2018 *Hikvision Digital Technology Co., Ltd.* Hangzhou, China

- Developed and maintained a QT-based front-end software for editing, viewing, and annotating videos in C++

Technical Skills

Programming Languages: **Python**, **C++**, **C#**, Objective C, Shell Scripting

ML/DL Platforms & Libraries: **Pytorch**, TensorFlow, Keras, **Diffusers**, **Transformers**, OpenCV, Scikit-learn

AR/VR & Graphics Development Tools: **Unity3D**, Blender, MeshLab, **Oculus**, **HoloLens**

Selected Publications & Preprints

An Exploratory Study on Multi-modal Generative AI in AR Storytelling

Hyungjun Doh*, **Jingyu Shi***, Rahul Jain*, Heesoo Kim, Karthik Ramani
DIS '25 (Under Review): the ACM Designing Interactive Systems Conference 2025

DesignFromX: Empowering Consumer-Driven Product Design by Facilitating Feature Composition of Existing Products

Runlin Duan, Chenfei Zhu, Yuzhao Chen, Yichen Hu, **Jingyu Shi**, Karthik Ramani
DIS '25 (Under Review): the ACM Designing Interactive Systems Conference 2025

PowVRtool: A Handheld Haptic Device for Realistic Power Tool Feedback in VR-Based Manufacturing Training

Mayank Patel, Asim Unmesh, Ananya Ipsita, Levi Erickson, Priyam Maheshwari, Rahul Jain, **Jingyu Shi**,
Laura H Blumenschein, Karthik Ramani
DIS '25 (Under Review): the ACM Designing Interactive Systems Conference 2025

CARING-AI: Towards Authoring Context-aware Augmented Reality INstruction through Generative Artificial Intelligence

Jingyu Shi*, Rahul Jain*, Seunggeun Chi*, Hyungjun Doh, Hyung-gun Chi, Alexander J. Quinn, Karthik Ramani
CHI '25 (Accepted): the 2025 CHI Conference on Human Factors in Computing Systems

Transparent Barriers: Natural Language Access Control Policies for XR-Enhanced Everyday Objects

Kentaro Taninaka, Rahul Jain, **Jingyu Shi**, Kazunori Takashio, Karthik Ramani
CHI '25 (Accepted): the 2025 CHI Conference on Human Factors in Computing Systems

An HCI-Centric Survey and Taxonomy of Human-Generative-AI Interactions

Jingyu Shi*, Rahul Jain*, Hyungjun Doh, Ryo Suzuki, Karthik Ramani
CSUR (Major Revision): ACM Computing Surveys

Visualizing Causality in Mixed Reality for Manual Task Learning: An Exploratory Study

Rahul Jain*, **Jingyu Shi***, Andrew Benton, Moiz Rasheed, Hyungjun Doh, Subramanian Chidambaram, Karthik Ramani
TVCG (Accepted): IEEE Transactions on Visualization and Computer Graphics

OccludeNeRF: Handling Occlusions in 3D Scene Inpainting with Collaborative Score Distillation in NeRF

Jingyu Shi, Achleshwar Luthra, Jiazhi Li, Xiyun Song, Zongfang Lin, Heather Yu, Liang Peng
CVPR '25 (Under Review): the IEEE/CVF Conference on Computer Vision and Pattern Recognition 2025

Fair, Creative, and Faithful: Diverse Text-to-Image Generation via Bimodal Classifier-Free Guidance

Jiazhi Li, Mi Zhou, Mahyar Khayatkhoei, **Jingyu Shi**, Jiageng Zhu, Hanchen Xie, Xiyun Song, Zongfang Lin, Heather Yu,
Liang Peng, Jieyu Zhao
CVPR '25 (Under Review): the IEEE/CVF Conference on Computer Vision and Pattern Recognition 2025

Generalized TensorRF: Efficient Multi-Scene Radiance Fields and View-Consistent 3D Editing

Achleshwar Luthra, **Jingyu Shi**, Xiyun Song, Zongfang Lin, Heather Yu, Liang Peng
CVPR '25 (Under Review): the IEEE/CVF Conference on Computer Vision and Pattern Recognition 2025

avaTTAR: Table Tennis Stroke Training with Embodied and Detached Visualization in AR

Dizhi Ma, Xiyun Hu, **Jingyu Shi**, Mayank Patel, Rahul Jain, Ziyi Liu, Zhengzhe Zhu, Karthik Ramani
UIST '24: Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology

Parametric: Empowering In-Situ Parametric Modeling in Augment Reality for Personal Fabrication

Runlin Duan, Xiyun Hu, Min Liu, **Jingyu Shi**, Karthik Ramani
ASME 2024 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference

ConceptVis: Generating and Exploring Design Concepts for Early-Stage Ideation Using Large Language Model

Runlin Duan, Nachiketh Karthik, **Jingyu Shi**, Rahul Jain, Maria C. Yang, Karthik Ramani
ASME 2024 International Design Engineering Technical Conferences and Computers and Information in Engineering Conference

Interacting objects: A dataset of object-object interactions for richer dynamic scene representations

Asim Unmesh, Rahul Jain, **Jingyu Shi**, VK Chaithanya Manam, Hyung-Gun Chi, Subramanian Chidambaram,
Alexander Quinn, Karthik Ramani
RAL: IEEE Robotics and Automation Letters (Volume: 9, Issue: 1, January 2024)

Ubi-TOUCH: Ubiquitous Tangible Object Utilization through Consistent Hand-object interaction in Augmented Reality

Rahul Jain*, **Jingyu Shi***, Runlin Duan, Zhengzhe Zhu, Xun Qian, Karthik Ramani
UIST '23: Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology

Understanding Generative AI in Art: An Interview Study with Artists on G-AI from an HCI Perspective

Jingyu Shi*, Rahul Jain, Runlin Duan, Karthik Ramani
arXiv preprint arXiv:2310.13149

Ubi Edge: Authoring Edge-Based Opportunistic Tangible User Interfaces in Augmented Reality

Fengming He, Xiyun Hu, **Jingyu Shi**, Xun Qian, Tianyi Wang, Karthik Ramani
CHI '23: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems

Patents

Authoring edge-based opportunistic tangible user interfaces in augmented reality

Karthik Ramani, Fengming He, Xun Qian, **Jingyu Shi**, Xiyun Hu
US20240312154A1 (Pending)

Visualizing Causality in Mixed Reality for Manual Task Learning

Karthik Ramani, **Jingyu Shi**, Rahul Jain
US20240135831A1 (Pending)