

Xi Fang

Machine Learning / Research Engineer | Embodied AI & VLA | Robot Learning | 3D Perception

Cupertino, CA • (518) 960-2197 • xifang96@gmail.com • fang-xi.com • LinkedIn • GitHub • Google Scholar

SUMMARY

Machine Learning / Research Engineer with **7+ years of AI/ML research and engineering experience** across academia and industry. Adapts pretrained **vision-language-action** models through **parameter-efficient and reinforcement-learning post-training**, evaluates them in **closed-loop simulation**, and deploys them to real robot hardware; ships production 3D perception and multimodal registration models in a regulated surgical-robotics product. Publications in MedIA, IEEE TMI, MICCAI, and IPCAI.

ROBOT LEARNING & EMBODIED AI

Vision-Language-Action Post-Training (PEFT + RL) & Sim-to-Real Manipulation

May 2026 – Present

- **Reproduced and benchmarked** Isaac GR00T N1.7 post-training recipes (full fine-tuning, LoRA, QLoRA) under DeepSpeed, tuning the embodiment projector and **flow-matching action head** while keeping the VLM backbone frozen; reached **93.5% success on the LIBERO-Long suite**.
- **Extended post-training from imitation to reinforcement learning** on RLinf: implemented critic-free **group-relative advantage (GRPO)** with zero-variance group filtering, and **Flow Policy Optimization** to derive a likelihood ratio for the flow-matching action head.
- **Designed the closed-loop evaluation protocol** — fixed seeds, a per-suite rollout budget, and bootstrap confidence intervals — and re-evaluated released checkpoints under the same harness for like-for-like comparison.
- Collected **100** teleoperated demonstrations and fine-tuned the policy for an **SO-101** arm; deployed with **TensorRT** and reached **72% real-world task success** on pick-and-place and cover-the-cube, with failure-mode analysis of the sim-to-real gap; documented results across LIBERO and an Open-H-Embodiment endoscope dataset in a technical report.

Learned Forward Models for Surgical Planning

May 2021 – Mar 2023

- Developed an attention-based **correspondence model** that predicts soft-tissue deformation from planned bony movement, matching finite-element accuracy while cutting simulation time from **5 hours to 1.2 minutes** (MICCAI 2022, Young Scientist Award).
- **Used the learned forward model inside an optimization loop** to search for bony movements that reach a target facial outcome — one of the first outcome-driven approaches to orthognathic planning (MICCAI 2023).

EXPERIENCE

Machine Learning Engineer — Noah Medical, San Carlos, CA

Dec 2023 – Present

- Led **Real-time Lesion localization** in intra-operative tomography, achieving **<3 mm mean and <2 mm median** lesion error across **7** different C-Arm types; integrated the optimized model into the core production **C++** system for real-time image-guided biopsy.
- Led development and production deployment of an airway segmentation system, from annotation / QA strategy through customized U-Net training and **TensorRT** inference; replaced a third-party API, achieved **>95% overall recall**, and reduced segmentation-related complaints by **90%**.
- Built an AI-assisted fluoroscopy video annotation pipeline using **foundation models for pre-labeling and cross-frame tracking**, with **human-in-the-loop** refinement and QC criteria co-defined with the data team; cut per-video labeling time **83%** (12h → 2h) and established a continuous model-improvement loop.
- **Scoped** forward-looking ML directions with product leadership — learned **world models** for tomographic image quality and **VLA-based robotic autonomy** — and delivered technical assessments that informed the roadmap.

Immersive Media Research Intern — Tencent America, Palo Alto, CA

Feb 2022 – May 2022

- Built 360° scene-understanding pipelines for Tencent's 3D Virtual Tour product, using panoptic segmentation and depth with MSeg, camera-model-based panorama-to-perspective projection, and **unsupervised domain adaptation** to support 3D reconstruction across **100+ classes**.

SKILLS

Robot Learning: Vision-language-action post-training (full fine-tuning, LoRA / QLoRA), imitation learning,
RL post-training (GRPO, Flow Policy Optimization), flow-matching policies, action chunking, closed-loop simulation
evaluation (LIBERO / MuJoCo), sim-to-real transfer
Perception: 3D segmentation, multimodal registration, object detection, tracking, action / phase recognition, semantic
correspondence, 360° scene understanding, self-supervised learning
Systems: Python, C++, PyTorch, TensorRT, DeepSpeed, PEFT, LeRobot, Transformers, Weights & Biases, MLflow

EDUCATION

Ph.D., Biomedical Engineering — Rensselaer Polytechnic Institute Aug 2018 – Nov 2023
Dissertation: *Deep Learning-Based Orthognathic Surgical Planning* | Advisor: Dr. Pingkun Yan
B.S., Computing — Wuhan University Sep 2014 – Jun 2018

SELECTED PUBLICATIONS & HONORS

Patient-Specific Reference Model Estimation for Orthognathic Surgical Planning. Fang X et al. *IPCAI 2024 / IJCARS* special issue. **Early Acceptance; NVIDIA Bench-to-Bedside Award Runner-up.**
Deep Learning-Based Facial Appearance Simulation Driven by Surgically Planned Bony Movement. Fang X et al. *MICCAI 2022*. **Oral; MICCAI Young Scientist Award (Top 0.3%); Best of MICCAI.**
Correspondence Attention for Facial Appearance Simulation. Fang X et al. *Medical Image Analysis*, 2024.
Soft-Tissue Driven Craniomaxillofacial Surgical Planning. Fang X et al. *MICCAI 2023*. **Early Acceptance (Top 13.6%).**
Multi-Organ Segmentation over Partially Labeled Datasets with Multi-Scale Feature Abstraction. Fang X, Yan P. *IEEE Transactions on Medical Imaging*, 2020.
MICCAI Student Travel Award (2022). Full publication list on Google Scholar.