Jonathan Fu

917-703-6343 | jonathan.fu01@gmail.com | linkedin.com/in/jonathan-fu-89031918b | 3-Minute Portfolio Video Summary

Education

Duke University Aug. 2024 – May 2026
Master of Engineering in Mechanical Engineering, 3.56/4.0 *Durham, NC*

Case Western Reserve University Aug. 2019 – May 2024
Bachelor of Science in Mechanical Engineering *Cleveland, OH*

Experience

Mechanical Engineer – Full-Stack Roboticist Aug. 2026 – Present
Tencent America *Bellevue, WA*

- Developed an **intelligent, expressive humanoid face robot** for a **video-game character**, owning the end-to-end robotics pipeline spanning **mechanical embodiment, facial actuation, motion retargeting, and multimodal intelligence**.
  - Iteratively optimized the robot's **physical embodiment through CAD**, redesigning eyelid mechanisms, facial end-effector geometries, and actuation point locations to improve expressive capabilities, motion fidelity, and mechanical integration.
  - Developed an end-to-end **conversational AI pipeline** integrating **Whisper** for speech recognition, **Qwen-30B** for language generation, and **Kokoro** for speech synthesis, establishing a functional baseline with an end-to-end conversational latency of **5–8 seconds**.
  - Rearchitected the sequential inference pipeline using **Qwen3-Omni**, unifying audio-visual perception and language generation to reduce conversational latency from **5–8 seconds to 1–3 seconds**. Subsequently integrated **Qwen3-TTS** with the model's text output to enable character-specific voice cloning.
  - Developed a **speech-driven facial animation and control pipeline** using **wav2vec-based phoneme extraction**, mapping synthesized speech to temporal viseme sequences and predefined facial poses, generating coordinated multi-servo commands for synchronized physical articulation and audio playback.
  - Engineered a **multimodal LLM harness** integrating **ArcFace-based identity recognition**, automatic identity enrollment, persistent memory, structured memory consolidation, relationship modeling, and emotion generation to enable personalized, context-aware human-robot interactions across sessions.

Robotics R&D Engineer Fall 2025
Richtech Robotics *Las Vegas, NV*

- Developed a **hat-pressing robot showcased at NVIDIA GTC** in a **team of three**, receiving direct commendation from **NVIDIA's CTO**. Presented point-cloud alignment pipeline to a team of Starlink precision-machining engineers at **SpaceX Redmond**, including **Starlink's Director of Manufacturing**.
  - Led a four-member annotation team under a two-week deadline to train YOLO for hat-seam and logo detection. While logos were reliably detected, deformable hat fabric caused YOLO detection inconsistencies, so I engineered a Frangi-based edge-detection fallback that extracted horizontal/vertical seam line of best fit using bounding-box filtering. This geometric method robustly localized the hat for the GTC demo.
  - Reduced **eye-in-hand extrinsics error** from **2 mm → 0.04 mm** using **ICP geometric regression** with an **ArUco board**, seeded by Tsai's calibration for the initial estimate. Validated accuracy under **1.6 mm intrinsics variance**.
- Converted URDFs to full **Isaac Sim USD** models of the **Dex upper humanoid** and **mobile robot**, including articulation and differential-drive controllers; assets are being published in **Isaac Sim's prebuilt library** and used for Isaac Cosmos synthetic data generation.
- Integrated **ROS2 joint-command subscribers** and **MoveIt IK** for real-time motion planning in Isaac Sim.
- Fine-tuned **Pi0.5** and **GR00T VLA** models for **XArm**, using **400 leader-follower episodes** to achieve smooth and reliable Rubik's cube grasping.
- Mapped human body pose to a robot arm by extracting joint pixel locations (u, v) with Google's **MediaPipe**, obtaining depth from an Orbbec **RGB-D** camera, and reconstructing 3D points via camera intrinsics and using the rpy from the elbow-to-wrist link before bridging to **ROS2 MoveIt** for inverse kinematics control.