

Mission

Augment humanity with physical intelligence

TEAM AND NETWORK



Yue Wang




- USC Assistant Professor, director of the Physical Superintelligence Lab, and former [NVIDIA Research](#) senior faculty researcher.
- Research focus: 3D world models, humanoids, and dexterous manipulation.
- Raised **\$3.6M** to support a **35-person team** at USC within only **2.5 years**.
- [MIT PhD](#); Dynamic Graph CNN is an [ICBS Best Paper](#) (only **4 awarded** in 2022) and the most-cited paper in ACM Transactions on Graphics.
- **21K+ citations** and **80+ papers** across AI, computer vision, and robotics venues.
- **8× best-paper awards**, NVIDIA Fellowship (only **5 awarded** in 2020), and Toyota Young Faculty Fellowship; advised students winning **2× NVIDIA**, **4× Qualcomm**, **1× Amazon**, and **1× Capital One** fellowships.



Di Huang



- Co-founded [Google Cloud Vertex AI](#), building production AI in [Fei-Fei Li's](#) organization.
- **Co-lead** of [Waymo's](#) first-generation Foundation Model.
- Scaled the model from **7B to 140B parameters** — the industry's **first robotic foundation model in production**, deployed on **1,000+ robotaxis**.
- Led Waymo's first-generation perception world model.
- Deep expertise across infrastructure, engineering, and operations.



World Engine

by the numbers

Time in operation

10 months

Current staff

30 employees

Enterprise customers

8 frontier labs

BEST IN SILICON VALLEY

Run rate 2026

~\$100M

Human data processed

200K hours

EGO + GLOVE + TELEOP

Best-paper awards

3 at CVPR 2026


2 BEST + 1 RUNNER-UP

12-month data target

12.5M hours

Physical AI Is Hitting a Wall

	LLM agents GPT/Claude/Gemini	Autonomous driving FSD V14	Embodied intelligence Foundation models
DEPLOYMENT SURFACE	1B-scale users	6M-6.5M vehicles	Sparse robot deployments
MODEL SCALE	3T-10T parameters	10B-50B parameters	3B-class models
TRAINING DATA	50-100T tokens	100-1,000T tokens	Only 1-2T tokens




OUR SOLUTION

Turn Every Human into a **Physical API**



EXISTING CUSTOMERS






\$15M to deliver in 2 months

IN THE PIPELINE




Hard Part #1: Making Human Data Speak Humanoid



We solve this by owning the stack end to end — capture hardware, post-processing algorithms, and operations — so every hour of human data is engineered to speak humanoid.

Hard Part #2: Only the **Model** Knows What Data It Needs



An **open model** becomes the reference design the ecosystem trains on — and every developer on it drives demand for **our data**.

PROOF POINTS

The Ingredients Already **Work**

CVPR 2026 3D VLA/LLM • Best Paper

RSS 2026

Official site: psi-lab.ai/Psi0/

40%+

higher overall success
than strong humanoid
baselines

800h


human video + 30h
robot data — only **8%** of
the data other methods
use

What it proves: the right data curriculum beats
brute-force scaling.

DEMO


PSI-0

Autonomous humanoid execution




BENCHMARK

Real-world evaluation: cross-task success vs public baselines



Full-Stack Research for Humanoid Intelligence

DATA LAYER



Humanoid Everyday

CVPR 2026 Embodied AI • Best Paper

ICRA 2026

The industry's highest-spec humanoid interaction dataset for real everyday tasks.

humanoideveryday.github.io

MODEL LAYER



Psi-0

CVPR 2026 3D VLA/LLM • Best Paper

RSS 2026

Humanoid foundation model for long-horizon loco-manipulation — 40%+ above strong baselines.

psi-lab.ai/Psi0

WORLD MODEL LAYER



PhysWorld

CVPR 2026 3D VLA/LLM • Best Paper Runner-Up

ICRA 2026

Physical world models a full year ahead of embodied-AI startups — on Google DeepMind's radar.

jiageng.me/PhysWorld_Web

FD-Loss

State-of-the-art image & video generation loss — directly lifts world-model generation quality.


Fast SAM 3D Body

Accelerating SAM 3D Body for real-time full-body human mesh recovery — built for humanoid control.

PhysBench

Benchmarking and enhancing vision-language models for physical world understanding.


Three Revenue Lines Compound the **Flywheel**



Near term: sell **data**. Mid term: license **infrastructure**. Long term: power **every robotics company and vertical user**.

COMPETITION

Full-Stack, Full-Loop for Humanoid Dexterous Manipulation



Vertically integrated in data, **horizontally integrated** with the model — we own the infrastructure loop.

Building the **Physical API** to Augment Humanity



	6 months	12 months	24 months
Data	1.25M hours	12.5M hours	62.5M hours
Model	Reference-design models for 10 common tasks	Scale to 100 tasks	Continuous learning at deployment time
Infra	Complete the full stack	End-to-end deployment	Model-agnostic full stack
Business	Data / evaluation as a service	1–3 vertical deployments	The default data & infrastructure layer for humanoid intelligence