SUNGHWAN CHO

shcho1551@gmail.com
linkedin.com/in/chosunghwan
github.com/sunnycho100

EDUCATION

University of Wisconsin–Madison Expected May 2027
B.S., Computer Engineering and Computer Science GPA 3.91/4.00

Relevant Coursework: Computer Architecture, Digital System Design and Synthesis, Digital Circuits, Machine Organization, Microprocessor Systems, Robotics, Database Management Systems, Electronic Circuits

RESEARCH EXPERIENCE

CATS Lab, University of Wisconsin–Madison Sep 2025 – Present
Research Intern, advised by Prof. Xiaopeng Li Madison, WI
Turning existing traffic camera video into trajectories to identify automated and human-driven vehicles

Roadside Camera Calibration

- Developed camera calibration in Python for uncalibrated traffic cameras, benchmarking three single-image intrinsics models and recovering camera pose from lane vanishing points
- Reconstructed the camera site in 3D from 200 Street View images (COLMAP, SuperGlue) to cross-verify the calibration, matching focal length within 1% and ruling out camera angle error as a 1 m GPS offset source

3D Vehicle Detection and Tracking

- Built a monocular 3D vehicle detection pipeline with bird's-eye-view height estimation (BEVHeight) in PyTorch, a 3D multi-object tracker (AB3DMOT), track stitching and an off-road feature mask
- Fine-tuned the detector through self-training on pseudo-labels from stabilized vehicle tracks, without manual annotation, cutting heading error from 176° to 1.3° and raising unseen-video detection from 46% to 91%
- Paired 2D and 3D detectors in congested traffic, using 2D detections to confirm low-confidence 3D boxes and retraining on the confirmed labels, extending test vehicle tracking from 13.1 s to 32.7 s

Trajectory Extraction and GPS Validation

- Built a Kalman filter with a backward smoothing pass (RTS) on true frame timestamps to handle dropped and frozen stream frames, cutting speed error against GPS from 12.6 to 0.27 mph on the test pass
- Built a held-out GPS evaluation over nine passes from two recording sessions, with a 2D-detection baseline and a method-independent coverage metric based on KITTI and nuScenes visibility rules

PUBLICATIONS

1. **[IV'27] Sunghwan Cho**, Hang Zhou, Bofeng Cao, Shixiao Liang, Keke Long, and Xiaopeng Li, "Extracting Vehicle Trajectories from Existing Roadside Traffic Cameras: Monocular 3D Detection and GPS Validation," *IEEE Intelligent Vehicles Symposium (IV)*, 2027 [in progress]
2. **[TR-C]** Hang Zhou*, Bofeng Cao*, **Sunghwan Cho***, Shixiao Liang, Keke Long, David Noyce, and Xiaopeng Li, "Vehicle Trajectory Reconstruction and Automated Vehicle Identification Using Roadside Cameras," *Transportation Research Part C: Emerging Technologies* [in progress]

*Equal contribution

Industrial Experience

MoDoc AI May 2026 – Oct 2026
Software Developer Intern Boston, MA (Hybrid)

- Led a 5-person team migrating 3,500+ pages from a no-code platform to Next.js and Payload CMS on GCP, preserving URLs for 20K+ users; track progress with a page-similarity score dashboard synced from GitHub
- Led 5-person AI orchestration and marketing teams, processing multimodal text, voice, and video data, and building partnerships with external organizations alongside the CEO

Contoro Robotics Feb 2026 – May 2026
Software Engineer Intern Austin, TX

- Built a robot control platform in Python, TypeScript, FastAPI, routing plain-English commands through GPT & Claude into MCP tool calls, hosted on AWS so anyone could drive robots in NVIDIA Isaac Sim with no setup
- Built the RobotMCP product website in React & JavaScript, adding GA4 event tracking to show where developers dropped off before signing up, and cutting page weight 90% to keep it usable on slow connections

PNS Networks Dec 2025 – Jan 2026
Software Engineer Intern Seoul, South Korea

- Built a freight rate platform in Python, Flask, React, and TypeScript with REST APIs, packaged as an Electron desktop app with an LLM chat panel, cutting manual processing time 95%
- Analyzed 100k+ multi-year shipment records in Python with regression modeling to find loss drivers, building a risk-scoring model to provide the client evidence to renegotiate carrier contracts

Deloitte Jun 2025 – Aug 2025
Technology Strategy and Transformation Intern Seoul, South Korea

- Supported senior engineers on a paperless transformation for a $50B financial institution, using TOGAF to map current systems and sequence upgrades across business units
- Scored AI-OCR, blockchain, and eForm vendors on a KPI and cost-benefit framework, benchmarking cross-industry cases and supporting the integration review of the client's Java applications

Republic of Korea Army Jun 2023 – Dec 2024
Sergeant, Network Engineer Seoul, South Korea

- Operated the PRC-999K tactical radio system and conducted geospatial signal analysis to maintain secure, stable brigade-level communications, supporting real-time coordination for 1000+ troops across diverse scenarios
- Served as the communication link between command and front-line units by integrating TICN (Tactical Information and Communication Network) to enable real-time data flow and enhance situational awareness

Teaching

University of Wisconsin–Madison 2025 – Present
Undergraduate Teaching Assistant, ECE 352: Digital System Fundamentals Fall 2025, Spring 2026, Fall 2026

- Instructed 100+ students through tutoring and office hours in RTL design and SystemVerilog, applying a reverse classroom model focused on guided debugging and systematic problem decomposition

PROJECTS

Embodied AI Competition for Manufacturing (ECM 2027) Oct 2026 – Present
Python, PyTorch, Isaac Sim, LeRobot, HTCondor, Docker

- Fine-tune AI policies (ACT, pi0.5 VLA) mapping camera images to robot actions for a mobile bimanual robot servicing a semiconductor chamber, in LeRobot and Isaac Sim on 227 human demos
- Build a data engine that adapts demos to moved and rotated parts (190 extra episodes), score policies on the same 28 sim seeds on HTCondor, and ship the best as an offline Docker image tested in CI

RTOS: Real-to-Sim Robot Deployment Aug 2026 – Present
Python, Isaac Sim, Omniverse, Isaac Lab, MuJoCo, Gaussian Splatting, RL rltosim.com

- Co-founded a 4-person physical AI startup that turns one site video into a robot-ready digital twin, rebuilding scenes with 3D Gaussian splatting in NVIDIA Isaac Sim and Omniverse for pre-deployment testing
- Simulated heterogeneous teams of humanoids and robot arms in MuJoCo and Isaac Lab, porting a multi-agent RL policy for two-humanoid cooperative carrying to 88% success and scaling to a 20-robot factory
- Ranked 1,219 robot deployment designs for a real car-door assembly line with discrete-event simulation and Pareto optimization, narrowing them to 3 options that trade off cost, cycle time and payback

Reinforcement Learning for Humanoid Locomotion Oct 2026 – Present
Python, PyTorch, MuJoCo, PPO, HTCondor, ONNX, Weights & Biases

- Fine-tuned a reinforcement learning walking policy for the Booster K1 humanoid with PPO in MuJoCo, raising the fall-free top speed in simulation from 2.4 to 4.5 m/s while keeping push and delay robustness
- Building a library of skill policies for walking, running, kicking and goalkeeping, combined through multi-agent RL orchestration and connected to a full computer vision pipeline that triggers each behavior

Neural Network Chip Design for Wearable Activity Detection (7nm Neuron) Sep 2026 – Present
Cadence Virtuoso, Calibre, HSPICE, ASAP7

- Designing a neural network neuron for wearable activity recognition in ASAP7 7nm, building the multiplier, adders, and ReLU multiplexer as separate transistor-level cells in Cadence Virtuoso
- Verifying each layout with Calibre DRC, LVS, and parasitic extraction, simulating the extracted netlist in HSPICE at 0.9 V to maximize operating frequency per area

RISC-V Processor Design Sep 2026 – Present
Verilog, RISC-V, ModelSim

- Designing a RISC-V RV32I processor in Verilog, growing a single-cycle datapath into a 5-stage pipeline with a hazard detection unit, data forwarding, and branch prediction
- Verifying every phase against a reference model with trace-based tests, tuning the pipeline to reduce stall cycles and raise instruction throughput

LifeVision: Pool Drowning Detection Sep 2026
Python, PyTorch, RF-DETR, ByteTrack, Isaac Sim, MuJoCo

- Built a CCTV drowning detection pipeline in Python with RF-DETR and a custom ByteTrack tracker that keeps swimmer IDs across frames, cutting false alerts 86% on real footage, served through a FastAPI review app
- Simulated drowning scenarios in NVIDIA Isaac Sim and MuJoCo to create ground truth that real footage can't provide, testing the full detection pipeline against it on unseen cases before fine-tuning raised recall from 25% to 100%

Embedded System Console Aug 2026 – Present
C, FreeRTOS, ARM Cortex-M, SPI, I2C

- Built a card game in C on an ARM Cortex-M microcontroller, scheduling display, input, and sensor work as concurrent FreeRTOS tasks over SPI and I2C drivers so the board runs the full game standalone

Local AI Orchestration Platform Jun 2026 – Aug 2026
Python, FastAPI, React, MCP

- Built the chat and connector layer of a local-first AI agent platform with a team, a FastAPI streaming backend with a React frontend, plus MCP connectors and local model backends so the agent ran on a laptop without cloud APIs

Vision-Action for Robotic Manipulation May 2026 – Jul 2026
Python, OpenCV, MuJoCo, ROS 2 github.com/sunnycho100/robot-arm

- Built a Python robot-arm pipeline with OpenCV tracking and a MuJoCo simulation twin that turns a demonstrated motion into a repeatable workflow, centering the gripper by visual servoing and sensing grip from servo feedback with no force sensor

Visual Research Paper Discovery Platform Apr 2026 – Jun 2026
React, TypeScript, Express, Gemini

- Built a research exploration tool that generates literature-grounded knowledge graphs from 200M+ works through the OpenAlex and Semantic Scholar APIs, using Gemini to organize topics into expandable trees

AI-Agent Productivity Platform Jan 2026 – Mar 2026
Next.js, TypeScript, Spring Boot, FastAPI, PostgreSQL

- Built a full-stack platform where multiple AI models debate to reduce hallucination, with multimodal document ingestion through PDF parsing, OCR, and GPT-4o vision, and PostgreSQL run tracking

Segway Self-Balancing Controller Aug 2025 – Jan 2026
SystemVerilog, FPGA, Synopsys DC github.com/sunnycho100/segway-control-system

- Built a Segway balance controller in SystemVerilog with PID control and gyro/accelerometer fusion over SPI, UART, and PWM, deployed to FPGA and synthesized on a SAED 32nm library to close timing at 333 MHz

AWARDS

- Dollmeyer Scholarship, College of Engineering, University of Wisconsin–Madison, 2026–2027
- SAIL (Summer of AI Laboratory) Scholarship, N+1 Institute and UW–Madison, 2026
- Summer Internship Scholarship, University of Wisconsin–Madison, 2026
- Undergraduate Scholarship for Summer Study, University of Wisconsin–Madison, 2026
- Dean's Honor List (all 5 of 5 semesters), University of Wisconsin–Madison
- Valedictorian, Class of 2021, Australian International School, Singapore, 2021

TECHNICAL SKILLS

Languages: C, Python, SystemVerilog, Assembly, Java, TypeScript, JavaScript, MATLAB
AI & Vision: PyTorch, OpenCV, YOLO, RL (PPO), ROS 2, Isaac Sim, MuJoCo, MCP, RAG
Frameworks & Systems: FastAPI, Flask, Next.js, React, Electron, Spring Boot, PostgreSQL, AWS, Linux, Docker, Git, Synopsys DC, Cadence Virtuoso, CosmosScope