

Jayden Wang

zw6963470@gmail.com Portfolio jaydenw.me GitHub [Jayden23018](https://github.com/Jayden23018)
Shenzhen College of International Education +86 156 0296 4366



PROFILE

I like building things, but I am just as interested in finding out where they go wrong. My first introduction to programming was practicing algorithm challenges. I fell in love, and gradually expanded my boundaries to develop technology for real-world scenarios, specifically focusing on products that genuinely help people. A Minecraft assistant that kept losing its plan led to my first-author paper, now accepted at the NeurIPS 2026 NEMO Workshop and under review at AAAI-27. Since then, I have explored medical-image synthesis, built BlindRun for visually impaired runners, and designed a safer walking router. Each project began with something that puzzled me and became a way to understand what makes a system useful—and worthy of someone's trust.

EDUCATION

Shenzhen College of International Education (SCIE) / Grade 12

Shenzhen, China

GPA 3.95 / 4.0; SAT 1560; TOEFL 5.5; IELTS 7.0

A Levels: Mathematics (A, achieved); Computer Science (A, predicted); Physics (A, predicted); Economics (A, predicted); Further Mathematics AS (A, predicted)

RESEARCH & SELECTED PROJECTS

Relevance-Filtered Symbolic Context Remediates the Task-Planning Bottleneck in Embodied LLM Agents **AI RESEARCH**

First author · NeurIPS '26 NEMO Workshop · AAAI-27, in review June. 2025 — Aug. 2026

I built a Minecraft agent that performed well on simple tasks but often skipped prerequisites in multi-step plans. Testing five models from 8B to 671B parameters showed the same problem across model sizes. I then wrote 46 task procedures and designed a retrieval method using Minecraft's recipe graph. On 60 held-out tasks, Qwen2.5-Max completed 19 tasks without retrieval and 49 with ranked procedures. The locally runnable 8B model improved from 23% to 43% without additional training. I designed the study, ran the experiments, and released the code with the paper.

Virtual Contrast-Enhanced CT Synthesis **CLINICAL AI**

Held-out evaluation returned a negative

Feb. 2026 — Now

I built a pix2pix GAN to synthesize contrast-enhanced CT images from non-contrast scans. The model scored 0.76–0.81 on SSIM, but the unchanged non-contrast scan scored 0.774–0.832 and beat it on six of seven runs — most anatomy is shared. I added this identity baseline to the evaluation. On held-out chest scans, predicted enhancement correlated with real contrast at Pearson $r = 0.31$, below the 0.5 success criterion I set. The architecture had not learned the enhancement mapping. This baseline is now a standard check in the team's pipeline. I designed the evaluation and ran the experiments.

BlindRun, for visually impaired runners ENGINEERING

Backend, then the iOS client too · Pre-launch, in internal acceptance testing Oct. 2025 — Now

I began BlindRun after meeting visually impaired runners at a road race in Shenzhen. Running clubs relied largely on phone calls and group chats to match runners with certified volunteers. I first built the Spring Boot backend and later took over the iOS client. After joining a guided run myself, I revised several features that had assumed a sighted user. The app now uses phone conversations instead of complex profile forms, assigns volunteers through sequential offers rather than open claiming, and sends SOS alerts directly to emergency contacts. BlindRun is preparing for a closed beta with a local running club, with a public launch planned for November 2026.

SafePath, a walking router that optimizes for safety ENGINEERING

Solo · Penn ESAP final project · live at safepath-black.vercel.app July. 2026 — Aug. 2026

I built SafePath to compare shorter and potentially safer walking routes in Philadelphia using 106,881 reported incidents and OpenStreetMap data. An early version incorrectly identified many short crosswalk segments as high-risk because of how I normalized the data. Revising the metric improved the results and showed me how strongly a small modeling choice can shape what a map appears to say. The website displays both routes and their trade-offs, while noting cases in which the alternative route offers little clear improvement.

HONORS & AWARDS

USACO, Gold Division	2026	AIME Qualifier	2025
USA Computing Olympiad · third of four tiers		AMC 12B 2024 · Global Distinction (top 5%)	
Senior Physics Challenge, Silver	2025	National Economics Challenge	2024
National physics competition		National individual Silver · regional team award	

SPEAKING

Speaker, COSCon'26 — the 11th China Open Source Conference, invited and scheduled, not yet given: BlindRun and the accessibility decisions behind it.

Speaker, COSCon'25 — invited talk on the task-planning research above, for a general open-source audience.

BEYOND THE SCREEN

Photography and drone work (street, landscape, aerial) · video editing (Bilibili) · table tennis, nationally certified Level-2 athlete · Go, amateur 3-dan.

TECHNICAL SKILLS

Languages: Java · Python · JavaScript / Node.js · Swift

Frameworks: Spring Boot · PyTorch · MONAI · Flask · SwiftUI · Mineflayer · NetworkX / OSMnx · pandas / GeoPandas · GANs (pix2pix)

Systems: MySQL · Redis · SQLite · WebSocket · DICOM / NIfTI · OpenAPI / contract testing · VoiceOver & accessibility APIs · WCAG 2.2

Method: Experimental design · pre-registered pass conditions · applied statistics (McNemar, Bonferroni, cluster bootstrap) · patient-level cross-validation · image registration · de-identification