

Tianliang Zhang (Jay)

Canberra, ACT, Australia

tianliangjay@gmail.com | tianliangzhang.com | github.com/xingkongliang | linkedin.com/in/zhangtianliang

Self-introduction

AI Engineer focused on LLM agents and production AI products, with a Ph.D. and six years shipping ML systems end-to-end. Independently built and shipped four AI products — a 2.2K-star open-source agent-skills manager and an AI IELTS coach with 1,400+ users — covering LLM scoring with strict-JSON output contracts, RAG, multi-provider fallback and evaluation harnesses. Currently Senior Software Engineer at LexisNexis Risk Solutions, building computer-vision systems for global identity-fraud detection. Published in CVPR and IEEE Transactions; ICCV LVIS Challenge Champion and Tencent Outstanding Employee.

Work experience

Independent AI Products | Founder & Solo Developer

2024 – Present

Built and shipped four AI products end-to-end:

- **Skills Manager** (github.com/xingkongliang/skills-manager): open-source desktop app (Rust/Tauri 2 + React 19, SQLite) that keeps one central AI-agent skill library and syncs it to 16+ coding agents (Claude Code, Cursor, Codex, Copilot...) by symlink or copy; presets toggle whole skill groups per agent/project scope, a background updater polls Git sources on an interval and auto-applies upstream changes, and a file watcher keeps each agent's view live. 2.2K GitHub stars, 190+ forks.
- **EasyIELTS** (easielyts.cn): AI IELTS speaking & writing coach — scores spoken answers on the four official IELTS criteria via a strict-JSON contract, returning per-criterion bands plus targeted fixes and vocabulary suggestions; a context-aware chatbot answers from the current page content and the user's selected text; production iterated with a Python analytics suite (retention/churn cohorts, recall campaigns, scoring evals). 1,400+ users. Next.js + Supabase.
- **ACT Learner** (passmate.com.au): bilingual ACT road-rules coach grounded in the official ACT Road Rules Handbook via a custom RAG pipeline — handbook PDF parsed and chunked along its own outline, each chunk fact-checked against the embedded PDF text layer with quality gates flagging unverified spans so the tutor never cites them as official; IDF-weighted coverage maps every practice question to its source chunk. Prompt-hash seed caching cuts repeat-question cost/latency; OpenAI/Gemini multi-provider fallback. Next.js.
- **Sticker It** (sticker.it.com): generative-AI sticker SaaS — generates a 4×4 sticker grid, then a CV pipeline slices it and removes backgrounds into transparent cut-outs. Next.js, OpenAI, Stripe.

LexisNexis Risk Solutions | Senior Software Engineer

Mar 2025 – Present

- Led the development and production deployment of an end-to-end AI identity verification suite, encompassing automated document preprocessing and multi-layered physical forgery detection systems.
- Designed and implemented vision models for card/passport segmentation, perspective transformation, and layout-based anomaly scoring to detect sophisticated forgeries.
- Re-architected the inference pipelines using ONNX and refined data processing workflows, reducing latency to sub-second levels to meet high-volume production requirements.
- Launched multiple core algorithmic systems into production with industry-leading accuracy (>99%), strengthening global fraud prevention capabilities and document processing efficiency.

PICEA (3irobotix) | Senior Algorithm Engineer

Mar 2023 – Apr 2024

- Built the vision stack for a voice/gesture-controlled robot vacuum — pose estimation, object detection, and LiDAR-fused 3D user-ranging; 95.06% spatial accuracy (0.17m average error), +18% detection in dynamic scenes.
- Optimized and deployed models on edge hardware (RK3566) with ONNX/RKNN, raising YOLO inference from 6 to 30 FPS.

Tencent, YouTu Lab | Senior Researcher

Jul 2020 – Mar 2023

- Built 14 CV algorithms (9 classical/OpenCV, 5 deep-learning/PyTorch: detection, segmentation, transformers) for lithium-battery surface-defect and electrode-tab inspection — 100% recall on critical defects at 0.5% false positives, cutting miss rate 65% versus the prior system.
- Deployed across 48 inspection machines processing 24,000 cells/day (~188GB), driving AUD 620K in industrial-AI contracts.

ICCV 2021 LVIS Challenge – Long-Tail Instance Segmentation

- Won Champion (Rank 1) and the Best Innovation Award — +4.9% AP on rare categories and +2.5% overall AP versus competitors.

Education

University of Chinese Academy of Sciences

Ph.D. in Signal and Information Processing

Sep 2017 - Sep 2020

Master of Science in Industrial Engineering

Sep 2013 - Jul 2017

Wuhan University of Technology

Bachelor of Electronic Information Engineering

Sep 2009 - Jun 2013

Skills

AI / LLM Engineering: LLM apps & agents (OpenAI / Gemini / Claude), RAG, function-calling & structured output (strict-JSON), evaluation harnesses, multi-provider fallback

Languages & ML: Python, TypeScript, C++; PyTorch, scikit-learn, OpenCV

Web & Product: Next.js, React, Supabase, Stripe, Tauri (Rust)

Infra & Deployment: Docker, Git, Linux, Vercel; ONNX / RKNN on edge (RK3566); ROS

Selected Publications and Patents

1. Ke W*, **Zhang T*(Co-first author, equal contributions)**, Huang Z, et al. Multiple Anchor Learning for Visual Object Detection. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
2. **Zhang T**, Xu H, Ye Q, et al. Feature Calibration Network for Pedestrian Detection in the Wild. *IEEE Transactions on Intelligent Transportation Systems*, 2020.
3. Gu C, Wang C, Gao B, Liu J, **Zhang T (Corresponding author)**, HDNet: A Hierarchically Decoupled Network for Crowd Counting, *ICME 2022*.

Awards and Honors

- 2023 PICEA Spark Award
- 2021 Tencent Outstanding Employee
- 2021 ICCV LVIS Challenge Champion and Best Innovation Award
- 2019 National Scholarship (PhD)