Chun Chet Ng

ngchunchet95@gmail.com | linkedin.com/in/chunchet-ng | github.com/chunchet-ng | chunchet.com | Kuala Lumpur, Malaysia

SUMMARY

Chief AI Officer with a Ph.D. in Artificial Intelligence (Computer Vision) specialising in OCR. Lead a 16-person team building LLM-powered document intelligence and credit scoring platforms for multiple financial institutions in Malaysia, processing 50,000+ documents per month with seven-figure ARR. First author of PRISM (EMNLP 2026 Industry Track; Poster Paper Award, FinLLM Symposium @ IJCAI 2026). Experienced in taking AI from research prototypes to on-premise and cloud production.

EXPERIENCE

AI Lens Group — Kuala Lumpur, Malaysia

Chief AI Officer — *April 2024 – Present*

- Lead a 16-person AI and engineering team delivering credit scoring and document intelligence platforms to multiple financial institutions in Malaysia, with seven-figure annual recurring revenue (ARR, MYR).
- Designed and delivered platforms for bank statements, payslips, ICs, invoices, financial statements, and bureau reports, processing 50,000+ documents per month at >95% extraction accuracy and <5s per page on on-premise infrastructure.
- Built production pipelines for document extraction, validation, financial analysis, credit scoring, and report generation using FastAPI, Python, React, PostgreSQL, Docker, RabbitMQ, and Azure.
- Provisioned and maintained Azure production infrastructure and cloud resources, and set up GitHub Actions CI/CD with self-hosted runners, branch gating, containerized deployments, and release readiness checks.
- Lead engineering delivery through sprint planning and reliability targets; mentor engineers across backend, DevOps, OCR, and AI system design.
- Drove the R&D track: first-author PRISM paper at EMNLP 2026 (Industry Track), which won the Poster Paper Award at the FinLLM Symposium @ IJCAI 2026; transitioned research into production; and helped the team secure 3rd place in the ACM ICAIF 2025 Agentic Retrieval for Financial Documents Grand Challenge on Kaggle.

Etch Media — Kuala Lumpur, Malaysia

Vice President of Engineering — *October 2023 – March 2024*

- Integrated an LLM-based classification pipeline mapping CV-detected object and person labels to contextual advertisement categories for real-time targeted ad delivery.
- Engineered a C# media playback service for custom LED display boards, handling advertisement scheduling and real-time content switching.
- Designed backend microservices bridging AI inference engines for vehicle recognition (make, model, colour) and facial analysis (gender, age, emotion) into production ad delivery workflows.
- Set up GitHub Actions CI/CD and containerized deployments; improved web security via Cloudflare and maintained Azure cloud resources.

Tibame Image Recognition — Remote

Part-time Instructor — *2023 – 2025*

- Open-sourced course materials for license plate recognition and signboard OCR.
- Taught hands-on OCR classes covering license plate recognition and signboard text spotting using PaddleOCR and MMOCR. [video]

AkiraKan Limited — Kuala Lumpur, Malaysia

Part-time ML/AI Research Scientist — *January 2023 – September 2023*

- Worked on 3D and AI projects; handled OCR-related projects including lamp post ID extraction and license plate recognition based on traffic scene images.

ViTrox Corporation Bhd. — Penang, Malaysia

Part-time Research Intern — *March 2020 – December 2022*

- Worked on industrial OCR methods for detecting and recognizing defective characters on chip images under low contrast, blur, broken-stroke, and other aesthetic-defect conditions.
- Co-organized ICDAR 2021 RRC - ICText with Baidu and partner universities, contributing to the largest Integrated Text on Chips Dataset (ICText).

- Proposed a curriculum-learning training strategy to improve character recognition performance on ICText; produced two publications from the research project.

University of Malaya Kuala Lumpur, Malaysia
Research Assistant (Ph.D. research) *September 2018 – December 2023*

- Researched scene text detection, recognition, spotting, and low-light image enhancement under Prof. Chan Chee Seng; contributed to benchmark datasets ArT (ICDAR 2019), EST-VQA (CVPR 2020), and ICText (ICDAR 2021).

SELECTED PUBLICATIONS

(1) **Chun Chet Ng***, Jia Yu Lim*, Zhen Hao Chu, Yixi Zhou, and Wei Zeng Low. "PRISM: Prompt-Refined In-Context System Modeling for Financial Retrieval." EMNLP 2026 (Industry Track). **Poster Paper Award**, FinLLM Symposium @ IJCAI 2026; **Oral**, Workshop on Machine Learning in Finance @ KDD 2026. [arxiv] [code]

(2) **Chun Chet Ng***, Zhen Hao Chu*, et al. "AI-BAAM: AI-Driven Bank Statement Analytics as Alternative Data for Malaysian MSME Credit Scoring." Workshop on Advances in Financial AI @ ICLR 2026; earlier version **Oral** at ACM ICAIF 2025 workshop. [arxiv]

(3) **Chun Chet Ng***, Che-Tsung Lin*, et al. "When IC Meets Text: Towards a Rich Annotated Integrated Circuit Text Dataset." Pattern Recognition, 2024. [paper] [code]

(4) **Chun Chet Ng***, Che-Tsung Lin*, et al. "Text In the Dark: Extremely Low-Light Text Image Enhancement." Signal Processing: Image Communication, 2025. [paper] [code]

(5) Xinyu Wang, Yuliang Liu, Chunhua Shen, **Chun Chet Ng**, et al. "On the General Value of Evidence, and Bilingual Scene-Text Visual Question Answering." CVPR 2020. [paper]

**Equal contribution. Full list (11 papers incl. ICPR, ICDAR, ICIP, ACM MM) at Google Scholar.*

EDUCATION

University of Malaya Kuala Lumpur, Malaysia
Ph.D. in Artificial Intelligence (Computer Vision) *September 2018 – December 2024*

- Thesis: "Optical Character Recognition Under Challenging Lighting Conditions". Submitted December 2023; passed the viva December 2024.

Tunku Abdul Rahman University of Management & Technology Kuala Lumpur, Malaysia
Bachelor of Computer Science (Honours) in Software Engineering; ***CGPA: 3.97/4.00*** *May 2015 – May 2017*

- President's Award (best graduate of the year, 2017/2018).

AWARDS

FinLLM Symposium @ IJCAI 2026: Poster Paper Award (PRISM).

ACM ICAIF 2025 Agentic Retrieval for Financial Documents Grand Challenge: 3rd place (Kaggle).

Peer reviewer: CVPR, ICCV, ECCV, ICLR, AAAI, ACM MM, BMVC, ICDAR; Pattern Recognition.

SKILLS

Programming: Python, C#, Java, SQL, JavaScript/TypeScript

AI/ML: LLMs, RAG, OCR, document intelligence, credit scoring, computer vision

ML Libraries: PyTorch, TensorFlow, NumPy, Pandas

Backend & Cloud: FastAPI, React, PostgreSQL, Docker, RabbitMQ, Azure

DevOps: GitHub Actions, CI/CD, on-premise deployment

Certification: NVIDIA Certified Professional: Generative AI LLMs (2026)

Languages: Mandarin (Native), English (Professional), Malay (Fluent)