Yanlong Yang

PhD Researcher — Computer Vision & Machine Learning

yanlong.yang@my.jcu.edu.au • +61 434 774 843 • Cairns, QLD, Australia

yyxr75.github.io • github.com/yyxr75 • Google Scholar • ORCID 0009-0000-8870-5993

RESEARCH INTERESTS

- Diffusion / generative models for inverse problems and computational imaging
- Multi-sensor fusion (radar, LiDAR, camera) for autonomous-driving perception
- 3D object detection, tracking, and point-cloud enhancement
- World models for perception

EDUCATION

PhD, Engineering — Computer Vision & Machine Learning

Apr 2024 – Oct 2027 (expected)

James Cook University, Cairns, Australia · Advisor: Dr. Tao Huang (Intelligent Computing & Communications Lab).

MEng by Research, Optical Engineering

2017 – 2020

Beijing Institute of Technology · Thesis: 3D human thermal-surface reconstruction (86.8 / 100).

BEng, Opto-Electronics Engineering

2013 – 2017

Tianjin University · Honours thesis on BST–silicon metamaterial THz spectra (86 / 100).

RESEARCH & PROFESSIONAL EXPERIENCE

Doctoral Researcher — James Cook University

Apr 2024 – Present

Generative priors for radar super-resolution point-cloud generation; world models for perception. Contributed computer-vision models to a JCU × Queensland Government (DETSI) AI crocodile-detection system, trained on thousands of hours of video.

Algorithm Engineer — Vanjee Technology, Beijing

2020 – 2024

Radar–LiDAR perception for autonomous driving: 3D object detection and sensor fusion in adverse weather (basis of RaLiBEV). Led a 3-person team building a data closed-loop system (shipped v1.0); built a MIMO radar signal-processing toolkit.

Graduate Researcher — Beijing Institute of Technology

2018 – 2020

Designed the hardware, software, and mechanical system for human 3D thermal-surface reconstruction.

PROJECTS

Australian Skilled-Migration Toolbox — ausimmitool.com (solo, end-to-end)

Aug 2026 – Present

Public Chinese-language data site covering ~490 occupations. Scrapes the government SkillSelect dashboard (Qlik engine over WebSocket) and rebuilds automatically via GitHub Actions when a monthly release lands; LLM-based CV-to-occupation matching served through a Cloudflare Worker API; email login with synced watchlists; interactive treemap heatmap and per-occupation dashboards.

World Models for Autonomous-Driving Perception

2025 – Present

Ongoing PhD work extending diffusion-based generative priors toward predictive world models of the driving environment; maintains a curated survey of occupancy world models (Awesome-Occupancy-World-Model).

AI Crocodile-Detection System — JCU × Queensland Government (DETSI)

2024 – Present

Computer-vision models that distinguish crocodiles from other floating objects, trained on thousands of hours of footage and deployed on a mobile, trailer-mounted camera platform for real-time public-safety monitoring in North Queensland.

DeblurSDI — Self-Diffusion Driven Blind Imaging (first author)

2025 – 2026

Zero-shot, self-supervised blind image restoration: recovers sharp images under unknown blur from a single degraded observation by driving a self-diffusion process, with no pre-training or external data. Accepted at CVPR 2026.

Self-Diffusion for Solving Inverse Problems (co-author)

2025

A self-diffusion framework that alternates noising and denoising with an untrained convolutional self-denoiser, adapting to arbitrary forward operators without pretrained priors. arXiv:2510.21417.

RadarINV — Unsupervised Radar Point-Cloud Enhancement

2023 – 2025

Began at Vanjee as radar super-resolution point-cloud generation from raw radar signals, continued in the PhD. Densifies sparse, noisy automotive-radar point clouds toward LiDAR-like quality with a diffusion prior guided by arbitrary, unpaired LiDAR data — no paired radar/LiDAR supervision. arXiv:2505.09887 · code: github.com/yyxr75/RadarINV.

RaLiBEV — Radar & LiDAR BEV Fusion for All-Weather 3D Detection

2022 – 2025

Fuses radar range-azimuth heatmaps with LiDAR point clouds in bird's-eye view via an interactive transformer fusion module and consistency-aware label assignment; anchor-box-free. Outperforms prior methods by 13.1% (clear) and 19.0% (fog) at IoU 0.8. Published in IEEE TCSVT · code: github.com/yyxr75/RaLiBEV.

Data Closed-Loop System for Vehicle–Road Cooperative Perception — Vanjee (team lead)

2022 – 2024

Led a 3-person team building a data closed-loop pipeline that feeds field data back into training to improve the generalization of perception models for vehicle-to-infrastructure (V2I) cooperative autonomous driving; version 1.0 rolled out to customers. Related roadside-LiDAR work: SSP, small-object detection from sparse roadside point clouds (ITS World Congress 2023).

MIMO Radar Signal-Processing Toolkit — Vanjee

2020 – 2022

Key developer of signal processing and algorithms for an in-house MIMO radar toolkit built from the TI reference design, laying the groundwork for radar–LiDAR fusion; hardware and software deployed on an engineering sample unit.

3D Human Thermal-Surface Reconstruction

2018 – 2020

Master's project at BIT. Built a ROS-based capture system that time-synchronizes an Intel RealSense D435i depth camera with a lab-built FLIR thermal camera; reconstructed full 3D human body models with point-cloud registration and global optimization, then UV-mapped thermal imagery onto the mesh to output textured 3D thermal body models for analyzing surface heat distribution. Published in Infrared Technology (2022).

Autonomous Smart-Car — NXP Cup, Tianjin University (hardware lead)

2015 – 2017

Designed the car's circuit boards in Altium Designer around a Freescale K60 MCU — power-supply, motor-driver, and CCD camera-acquisition modules — and co-developed the vision and control software: Otsu-threshold binarization for road-edge detection and closed-loop steering/speed control. Team won the National Second Prize (24th of 200+ teams).

PUBLICATIONS

★ Self-Diffusion Driven Blind Imaging (DeblurSDI)

2026

Y. Yang, G. Luo. *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Zero-shot, self-supervised blind imaging — no pre-training.

Self-Diffusion for Solving Inverse Problems

2025

G. Luo, S. Huang, **Y. Yang**. *arXiv:2510.21417*.

★ Unsupervised Radar Point-Cloud Enhancement via Arbitrary LiDAR-Guided Diffusion Prior

2025

Y. Yang, J. Liu, G. Luo, H. Li, E. Ahn, M. Rahimi Azghadi, T. Huang. *arXiv:2505.09887*.

★ RaLiBEV: Radar & LiDAR BEV Fusion for Anchor-Box-Free Object Detection

2025

Y. Yang*, J. Liu*, T. Huang, Q.-L. Han, G. Ma, B. Zhu. *IEEE Transactions on Circuits and Systems for Video Technology*, 35(5):4130–4143.

SSP: Small-Object Detection from Sparse Roadside-LiDAR Point Clouds for V2I Systems

2023

H. Liang, **Y. Yang**, J. Li, Y. Deng. *ITS World Congress*.

Fusion Reconstruction Method for 3D Temperature Fields on the Human Body Surface

2022

Y. Yang, C. Xu. *Infrared Technology*.

Observation of Phase Transitions of Ba0.6Sr0.4TiO3–Silicon Hybrid Metamaterial by THz Spectra

2020

L. Wu, Y. Fu, N. Xu, Q. Sheng, X. Zhang, **Y. Yang**, et al. *ACS Applied Electronic Materials*.

PATENTS

9 granted Chinese invention patents (first inventor on 4) · assignee: Vanjee Technology · filing year shown

Method for Multiple-Target Detection and Related Product

2022

Y. Yang, J. Li, Y. Deng. CN116863325B, granted.

Assisted Parking Method, Device, Computer Equipment and Storage Medium

2021

M. Liu, **Y. Yang**, Y. Deng, Y. Fang. CN114299146B, granted.

Point-Cloud Data Enhancement Method, Device, Computer Equipment, System and Storage Medium

2021

Y. Yang, J. Li, F. Meng, Y. Deng, L. Wu. CN116359908B, granted.

Automatic Parking Method, Roadside Equipment, Vehicle and System 2021

K. Cao, **Y. Yang**, Y. Fang, Y. Deng. CN113997932B, granted.

Noise-Point Processing Method, Device, Equipment and Storage Medium 2021

D. Liu, J. Li, **Y. Yang**. CN116047477B, granted.

Method and Device for Generating Data, Electronic Equipment and Storage Medium 2021

Y. Yang, Y. Deng, J. Li, L. Wu. CN116027288B, granted.

Multi-Sensor Signal Fusion Method, Device, Electronic Equipment and Storage Medium 2021

Y. Deng, **Y. Yang**, J. Li, L. Wu. CN116027318B, granted.

Assisted-Driving Method for Vehicles in Parking Lots, Intelligent Terminal and Roadside Equipment 2021

Y. Yang, J. Li, K. Cao. CN115731738B, granted.

Target-Object Identification Method, Device, Electronic Equipment and Storage Medium 2021

L. Wu, **Y. Yang**. CN113687358B, granted.

Plus 13 pending applications (2021–2023) on millimetre-wave radar signal processing, multi-sensor target detection and tracking, and point-cloud generation.

TECHNICAL SKILLS

Programming	Python (primary), C/C++, MATLAB · PyTorch, PCL, OpenGL, Qt
Tools	Git, LaTeX, COMSOL, Altium Designer, Visio, Linux
Domains	diffusion / generative modeling · inverse problems · sensor fusion · 3D reconstruction · autonomous-driving perception
Languages	Chinese (native), English (IELTS 6.5)

HONORS & AWARDS

Postgraduate Research Stipend Scholarship (HDR) — James Cook University 2024 – Present

Second-Class Academic Scholarship (awarded each year) — Beijing Institute of Technology 2017 – 2020

National Second Prize (ranked 24th of 200+ teams) — NXP Cup, China National University Smart-Car Competition 2017

Golden Boot & team champion — Vanjee Technology company football tournament 2020 – 2024

LEADERSHIP & ACADEMIC SERVICE

JCU Student Association — Cairns Campus Officer (2025–2026); Postgraduate Officer (2024–2025) 2024 – 2026

Elected student representative delivering campus services and initiatives that build student engagement on campus. Represented JCU at the Australian University Student Voice Conference, Gold Coast (2025).

Reviewer — AAAI Conference on Artificial Intelligence (AAAI 2025, 2026) 2025 – 2026

Reviewer — International Conference on Learning Representations (ICLR 2025, 2026) 2025 – 2026