Hector Andres Acosta Pozo

ML Engineer, Computer Vision & Edge AI

Daegu, South Korea
hector.acosta.pozo@gmail.com
hect.dev · github.com/hectorandac
linkedin.com/in/hector-acosta

I build computer vision models and ship them to the edge. I lead the AICU Gait Scanner, from pose and gait models to optimized on-device builds for iOS and Android, and the services behind them. 8+ years of software engineering, 3+ years of applied ML.

EXPERIENCE

AICU Nov 2024 – Present
Machine Learning Researcher, Gait Scanner lead

- Lead the AICU Gait Scanner: computer vision models for real-time gait analysis used by clinicians.
- Own the path from training to shipping: optimized edge inference with Core ML on iOS and CUDA/C++ accelerators, native Android and iOS apps.
- Connect the on-device stack with Python and Ruby on Rails services for regulatory-ready rollouts.

BetterUp Sep 2021 – May 2025
Software Developer, part-time contractor

- Built internal integrations with Microsoft Teams, Slack and other third-party platforms (Ruby on Rails, JavaScript).
- Helped build Microsites, the alternative payment platform that handled thousands of partner transactions.

Kaleyra (formerly mGage) Sep 2019 – Sep 2022
Software Engineer Specialist

- Designed a multimedia transcoding system and high-throughput messaging applications at mGage.
- After the Kaleyra acquisition, moved PHP monoliths to multithreaded Go services (Java, Rails, Go, PHP).

Pierpont Global Nov 2018 – Nov 2019
Co-founder, lead software developer

- Co-founded an online car auction platform and led a team of four building the full stack (Rails, JavaScript, Java).
- Rebuilt in 2026 as a live demo with ML pricing and photo-assisted listings.

RESEARCH

- **Gated YOLO** (MSc thesis). A dynamic gating mechanism inside YOLOv6 with scene-specific pruning, to cut compute for real-time detection on edge devices without losing accuracy. Code
- **A scene-specific object detection system utilizing the advantages of fixed-location cameras** (JICCE 2023, co-author). Edge AI detection for fixed cameras on embedded hardware. Paper
- **ROAD-SEC dataset** (Dataset, 2023). 13,221 images and 166,192 labeled objects from fixed security cameras over highways and roads, in varied weather and light. Dataset

SELECTED PROJECTS

- **Pierpont**. Live car auctions with ML pricing (LightGBM, conformal intervals), CLIP photo checks in int8 ONNX and weekly retraining behind a promotion gate. github.com/pierpontglobal/pierpont
- **finance-dashboard**. Self-hosted finance app with an MCP endpoint, so an AI agent can manage it through task-level tools behind a token and a NetworkPolicy. github.com/hectorandac/finance-dashboard
- **CRAFT Line Extractor**. Document line segmentation built on CRAFT word detection: deskewing, polygon grouping into lines, and a web API. github.com/hectorandac/CRAFT-Line-Extractor

SKILLS

- **ML & CV:** PyTorch, TensorFlow, YOLO, Pose estimation, OpenCV, Core ML, CUDA/C++
- **Platforms:** iOS, Android, Edge inference, Ollama, MCP
- **Backend:** Python, Go, Ruby on Rails, TypeScript, Java
- **Infra:** Kubernetes (k3s), Helm, Docker, VictoriaMetrics, Grafana, GitHub Actions

EDUCATION

Kyungpook National University, MSc, Computer Science — Daegu, South Korea
INTEC, Undergraduate studies — Santo Domingo, Dominican Republic