Daniel Poncelas Vargas

Valencia, Spain • (+34) 635 827 909 • danielponcelasvar@gmail.com • **danielponcelas.com** •
linkedin.com/in/daniel-poncelas-vargas

AI engineer finishing an M.Sc. in Artificial Intelligence. Two years of production computer vision at a startup, research on LLM-agent security being prepared for publication, and two products I built, launched and monetised on my own: an iOS app with 77k first-year users and a live SaaS at €65k ARR. Graduated in the top 1% of my Computer Engineering cohort, with exchanges at KTH Stockholm and Trinity College Dublin.

Education

M.Sc. in Artificial Intelligence, Pattern Recognition and Digital Imaging Sep 2025 – Sep 2026
Universitat Politècnica de València | GPA 8.7/10, 2 Honours distinctions *Valencia, Spain*

B.Sc. in Computer Engineering, Artificial Intelligence specialisation Sep 2021 – Jul 2025
Universitat Politècnica de València | GPA 9.1/10, ***top 1% of cohort****, 10 Honours distinctions* *Valencia, Spain*

- Best First-Year Academic Record award and University Excellence Scholarship (top 3 students of the year group). Final thesis graded 10/10.
- **KTH Royal Institute of Technology** (Aug 2024 to Jan 2025, 8.6/10): M.Sc. courses in Machine Learning, Deep Learning, Computer Vision and Applied GPU Programming. **Trinity College Dublin** (Jan to Jun 2024, 9.5/10).
- **CDTM** Honours Degree in Technology Management, a joint institution of TU Munich and LMU Munich (2025), admitting around 15 students per cohort. **ETH Zurich** summer programme in AI and Machine Learning (Jun 2025).

Research

M.Sc. Thesis: *A Systematic Evaluation of Prompt-Injection Defences in LLM Agents* 2026
Universitat Politècnica de València | paper submission in preparation *Valencia, Spain*

- Benchmarked **8 state-of-the-art defences** across **8 models from 3B to 70B parameters** (GPT-4.1, Claude, Gemini 2.5, Llama 3.x, Qwen 2.5) on the AgentDojo and InjecAgent benchmarks: **roughly 40,000 agent runs**, with utility, robustness and refusal metrics backed by statistical tests.
- Showed that the training-based defence is **Pareto-dominated by prompt-level defences**, refuting the starting hypothesis, and that **scale does not buy robustness**: within each model family, attack success rate climbs with parameter count (0.02 at 3B to 0.36 at 70B).
- Wrote a mechanistic interpretability chapter (activation patching, logit lens, TransformerLens) attributing defence failure to a **distributed circuit** rather than a few attention heads.
- Served open-weight models with **vLLM** (OpenAI-compatible, tool-calling) on a **SLURM** GPU cluster.

Experience

Solo Founder and Engineer, StudyOS and GenTotalAI 2025 – Present
Self-employed | engineering, product, growth and monetisation *Valencia, Spain*

- **StudyOS**, an iOS study assistant built on LLMs: **77k users** in its first year, peak **$42k ARR** and **$24k** revenue over the last 12 months at close to **100% gross margin** (€25/month of infrastructure). Reached **top 5** in the Education category of the Spanish App Store. Swift, Firebase, Supabase, LLM APIs.
- **GenTotalAI**, an AI video-generation SaaS launched in March 2026: **€65k ARR**, **190 paying subscribers** and **14,000** generation jobs served. Next.js, Express, Supabase/PostgreSQL, Redis and BullMQ job queues, Stripe, Cloudflare R2, CI/CD to a VPS.
- Set pricing and retention strategy from **cohort, churn and LTV analysis** joining Stripe, Supabase and PostHog data, and run paid acquisition end to end.

Computer Vision & AI Research Engineer (internships, then full-time) Sep 2023 – Oct 2025
Evestel Iberia | computer-vision startup: retail analytics and driver monitoring *Valencia, Spain*

- Person detection, tracking and re-identification for retail analytics and driver-monitoring systems. Models I trained and tuned **run in production**.
- Led the migration of the detection stack from **YOLOv5 to YOLOv11** with embedding-based re-identification, improving mAP@50 and mAP@50:95 by up to **5%**.
- Exported and optimised models through **ONNX and TensorRT**, and ran GPU-accelerated inference on both **NVIDIA (CUDA) and AMD (ROCm)** hardware.
- Automated dataset annotation and validation across roughly **15,000 images** (Roboflow), cutting manual data effort by about **40%**.
- Built the company's internal **computer-vision evaluation platform** as my B.Sc. thesis (graded 10/10): mAP, F1, precision-recall curves, benchmarking and data-drift detection.

Technical Skills

Languages: Python, Java, SQL, Swift, Bash, C/C++ (CUDA)
Machine learning: PyTorch, TensorFlow, Keras, scikit-learn, OpenCV, Hugging Face, TransformerLens, vLLM, ONNX, TensorRT
Product and infrastructure: Next.js, React, Node.js/Express, Supabase, PostgreSQL, Firebase, Redis/BullMQ, Stripe, Docker, Linux, Git, GitHub Actions, SLURM, PostHog
Spoken languages: Spanish (native), Catalan (native), English (C1, Cambridge certified)