Ángel Gabriel Zeledón Fernández

DATA SCIENTIST & MLOPS ENGINEER — SAN JOSÉ, COSTA RICA

angel.zeledon.fernandez@gmail.com · +506 7186-2820 · linkedin.com/in/angel-zeledon-data-science · github.com/Angel-Zeledon

Dear Hiring Team,

I build forecasting systems that run on measured evidence — and, just as importantly, the pipelines that keep them running after the demo is over.

Most forecasting work fails in the same two places: a model that fits history beautifully and predicts nothing useful, and a model that works in a notebook and nowhere else. Four years of forecasting and agentic systems for supply chains have made me stubborn about both.

At **Softtek, for Walmart's agricultural supply chain** (Hortifruti Coris and CEDI Coyol), I modelled distribution-centre inventory, store demand, supplier constraints, lead times and raw-material availability so procurement had something better than a moving average behind it. Each series got the model its demand pattern justified — LSTM, ARIMA, ETS or Croston — and I contributed to the administrative shrinkage-detection work aimed at a problem estimated near $8M a year in Costa Rica.

At **Accenture**, for HP, Philips and Cisco's reverse supply chain, I built delivery-time forecasting, bottleneck detection and cost models — and the Azure ML, MLflow, Docker, Kubernetes and CI/CD platform that retrained, monitored and served them through REST APIs. I also designed the Generative AI side: RAG, multi-agent orchestration with LangGraph and CrewAI, and Model Context Protocol servers that let agents query supply-chain databases directly.

Since June 2026 my main work has been my own two platforms: **StockAI** (stockai.es), an inventory and purchasing SaaS that tells Latin American businesses what to order today, how much and from which supplier; and **WorkforceAI** (workforceai.es), where small teams give AI agents work in plain language while the backend caps spend and holds sensitive actions for human approval.

What drives all of it is a genuine love of mathematics and science. I wrote a 2,080-page mathematics course to understand it all the way down, and I keep coming back to **fintech and quantitative finance** because that is where probability starts costing or saving real money — the reason I also consult part-time on betting-risk models at Photon Rush. I want to spend my career on the technology the next decades will run on.

What I would want you to check first is how I report results. One of my published benchmarks was wrong; I retracted it publicly and left the retraction in the repository next to the fixed number. My neural-interface work leads with its least convenient finding — across 22 subjects, no decoder wins — and my mathematics textbook is published next to an audit of what it still owes the reader. I would rather hand a stakeholder a number with its uncertainty attached than a confident one that quietly breaks their inventory.

If your team has forecasts that planners silently override, risk that nobody has priced yet, or models that never made it out of a notebook, that is exactly the work I want to do. I would be glad to walk through any of the above in detail.

Ángel Gabriel Zeledón Fernández

Data Scientist & MLOps Engineer · Spanish (native) · English (professional working proficiency)