

Ron Medina

📞 (+63) 962 831 6730 • ✉️ ron.mdn@gmail.com
🌐 linkedin.com/in/ron-medina-5a411218a/ • 🌐 particle1331

Senior Machine Learning Engineer with a maths & physics background. I design, build, and ship ML and deep learning systems that optimize operations and surface useful insight — with attention to the infrastructure and engineering that keeps them running in production.

Employment History

Sr. Machine Learning Engineer | Afni

Jan 2025 – Present

- Developed backend predictive models and evaluation pipeline for the hiring platform, with a recruiter-facing UI serving model predictions — drove a 6% increase in lifetime value and significant KPI improvements vs BAU hiring in A/B tests
- Team lead for multiple AI team initiatives: owner of cloud infra & code repositories, contractor coordination on task delivery, stakeholder meetings, roadmap development, code review, mentoring, and individual contribution. Drafted the team's MLOps framework and roadmap
- Collaborated with multiple VPs on feature discussions, roadmap, capabilities, and communications
- Direct report to the VP of Software Engineering & AI on day-to-day operations and progress
- Built the core backend models for the hiring platform that evaluates 5,000+ monthly candidates across all PH sites — predictive models for customer and employee retention risk, and ranking/recommender-style models for employee performance across program placements
- Designed a document analysis pipeline processing up to 30K documents daily with failure tracking, retries, and sufficient throughput; documents are attributed to the organizational hierarchy (employee, coach, manager) supporting RBAC, tagged using NLP and LLM tools, with a monitoring and analytics dashboard that only surfaces results users have access to based on their org hierarchy
- Implemented a RecSys-style solution for determining ideal program placement of candidate hires, by combining LOB description, contractual obligations, headcount requirements, candidate history, and interview/assessment scores
- Designed a repeatable feature engineering, preprocessing, training, and monitoring pipeline with a clean deployment interface for ML models served via Azure Functions; trained per-program per-KPI models, overcoming sparsity through extensive EDA and domain understanding

Machine Learning Engineer | Ubiquity

Aug 2022 – Jan 2025

- Reduced cost by 6X and increased accuracy of speech recognition system by +10%
- Developed and deployed a transcription service which transcribed 1.2M+ production calls
- Within 3 months after hiring, presented a POC for extending the call recorder system from mono to stereo recording using socket programming and Avaya APIs — became the basis for a major project for the Telco team and subsequent realtime transcription efforts
- Developed and designed a fault-tolerant distributed offline task queue service to scale speech recognition
- Developed a Transformer-based semantic search engine that runs performant on CPU
- Extended an existing open-source annotation tool for human data labeling to serve our internal use-case
- Developed AI services for downstream processing, modeling, and analytics of call transcripts
- Helped develop the backend application for searching and filtering transcriptions
- Contributed to a masking service for images (screenshots) containing sensitive data using Tesseract

Junior Data Scientist | Sitel

June 2021 – Aug 2022

- Creation and deployment of APIs for integrating ML algorithms with existing products
- Creation and deployment of Power BI dashboards
- Design of KPIs and metrics for various business processes
- Works directly under the Director of District Operations Quality Management
- Ensure integrity and accuracy of reports

Machine Learning Engineer | BrewedLogic, Inc.

Mar 2020 – June 2021

- Collaborative filtering RecSys written in NumPy, Scikit-Learn, and Pandas and served via Django — deployment increased average ticket count from 3.38 to 4.86 and average ticket value from \$11.49 to \$14.21 after the first two months in production. Became the RecSys platform of Crisp deployed on 26 US franchises each with multiple stores
- Worked with a senior data scientist on customer segmentation and sales forecasting and in developing a fraud detection model for fraudulent VoIP transactions, drastically improving over previous rule-based approaches
- Preprocessing, feature engineering, training, and monitoring of deployed models on 1M-5M row datasets; service migration from Django to FastAPI

Skills

Data Analysis

- Data visualization, data wrangling, and EDA using Pandas, seaborn, matplotlib, and NumPy.
- SQL, probability modelling, statistics, clustering

Machine Learning

- Deep neural networks in TensorFlow and PyTorch
- Machine learning models in scikit-learn
- Recommender systems, anomaly detection / imbalanced learning
- Weak supervision for training noise-aware models
- Gradient Boosting (Catboost, XGBoost, LightGBM), ensembling/stacking
- LLMs: prompting, RAG, fine-tuning, MCP, agentic workflows (LangChain, LangGraph)
- Embedding-based retrieval and semantic search

Model Deployment and MLOps

- REST APIs (Django, FastAPI, Flask); CI/CD (Gitlab CI/CD, GitHub Actions); uv, Typer, Makefiles
- Experiment tracking & model management with MLflow; task queues with Celery, SQS, RabbitMQ
- Containerization with Docker; unit/differential/regression testing with pytest; version control with git
- AWS: Lambda, SQS, RabbitMQ, S3, RDS, EC2 / Auto Scaling groups; PostgreSQL, MySQL, Redis
- Microsoft Azure: Azure ML / Foundry, Azure Container Registry, App Service, Blob Storage, Azure Functions

Others

- Gold level in Problem Solving and Python @ Hackerrank
- Author of OK Transformer — a collection of notebooks and articles on deep learning, ML engineering, and MLOps. Auto build / deploy via GitHub Actions + tox. Featured in the Gallery of Jupyter Books.
- Contributed to the Appendix: Mathematics for Deep Learning of Dive into Deep Learning — a widely used open-source DL textbook; acknowledged as a contributor.

Education

University of the Philippines - Diliman

06/13 – 12/18 (courses), 09/23 – 01/24 (thesis)

Bachelor of Science

Major in Mathematics

Awards: University Scholar, 2nd Semester 2013-2014. GWA: 1.23

Thesis: *An Intro. to Finite Frames and a QR Factorization Approach for Constructing MB Frames*

Relevant courses: Intro to Computer Science (Python), Numerical Analysis, General Relativity (MS / PhD level), Linear Algebra, Advanced Calculus

Eskwelabs

July 2019 – Oct 2019

Data Science Bootcamp

Attended a 10-week bootcamp which included 160 hours of in-class learning in addition to coursework. At the end of the bootcamp, I presented my capstone project about modeling nonlinear chaotic systems using neural networks implemented in TensorFlow 1.x to industry leaders in Makati City, Philippines.

Early Employment

Bootcamp Associate | Eskwelabs

Oct 2019 – Jan 2020

- Developed curriculum materials on machine learning algorithms. Facilitated live hackatons
- Presented a talk on artificial intelligence and deep learning at the National Youth Congress, UP Diliman School of Economics, Nov 2019

Data Analyst | Tita's Groceria (E-commerce)

June 2017 – June 2019

- Helped grow the shop's follower count from 30,000 to 100,000+ w/ hundreds of daily transactions
- Analyzed frequently-bought-together items using graphs, Markov chains, and correlations, and clustered customers by RFM criteria