

Christian Jeffrey Fonseca Rodriguez

AI Engineer | Data Architect | Data Scientist

Lima, Peru (-05:00 GMT)

(+51) 920202310

christian.fonseca.r@gmail.com

<https://linkedin.com/in/christian-fonseca-rodriguez>



Data & AI Engineer dedicated to solving enterprise challenges through advanced machine learning and cloud infrastructure. With a decade of experience across diverse industries, I specialize in automating data pipelines, deploying intelligent GenAI agents, and establishing MLOps standards. Leveraging deep expertise in Python, PySpark, and multiple cloud platforms, I consistently deliver reliable technical solutions that significantly reduce processing times and drive business growth.

EDUCATION

Master of Science in Artificial Intelligence (in progress) – Universidad Nacional de Ingeniería, Peru (2024 – 2026).

Diploma in Advanced Computing – Centre for Development in Advanced Computing, India (2015).

Bachelor of Science, Mechatronics Engineering – Universidad Nacional de Ingeniería, Peru (2011 – 2015).

PROFESSIONAL EXPERIENCE

Baufest, AIOps Specialist — May 2025 - Present

- Built AI chatbots for instant documentation retrieval and error analysis, reducing search and debugging time to seconds.
- Automated Boomi documentation generation, standardizing integration artifacts and accelerating knowledge delivery.
- Designed AI-driven unit test generation to automate validation and improve code quality and developer productivity.
- Implemented AI video documentation that converts walkthroughs into structured docs with conversational guidance.

Tech Stack: Python, Docker, LLM, RAG, OpenAI, Gemini, Bedrock, Boomi, React, AWS.

JobLeap AI, AI Solutions Engineer — October 2024 - April 2025

- Designed 10+ end-to-end data pipelines from multiple vendors using relational and vector databases.
- Established automated testing frameworks for data pipelines, guaranteeing high reliability and superior data quality.
- Applied advanced NLP and embedding techniques to enrich data transformations, boosting accuracy and relevance.
- Used LangChain and LLMs to generate relevance scores and create resume and interview simulations.

Tech Stack: Python, Docker, LLM, NLP, OpenAI, DeepSeek, PostgreSQL, Pinecone, AWS.

Antamina (via TIVIT), Senior Data Solutions Consultant — November 2023 - September 2024

- Migrated on-prem SQL to AWS RDS via DMS, enabling 3-min ingestion and cutting station refueling delays.
- Deployed RDS monitoring with automated alerts for failures/overloads, reducing production downtime.
- Built Databricks-based pipelines with Athena pipelines for ML-driven electricity-demand forecasts.
- Delivered recommendation dashboards with automated stakeholder notifications, boosting operational efficiency.

Tech Stack: Python, PySpark, SQL, PostgreSQL, Databricks, AWS.

BBVA, Manager Data Scientist — June 2021 - October 2023

- Achieved a 10% rise in debit card sales by creating a rule-based and ML-driven client profitability recommendation system.
- Improved expert evaluation of billing accuracy by 50%, using forecasting and elasticity models in retail budgeting.
- Boosted corporate customer acquisition by 15% through cost calculation system development.
- Implemented the RORC threshold selection, resulting in a 5% monthly decrease in the number of high-risk clients.

Tech Stack: Machine Learning, Python, PySpark, SQL, Oracle, Datio.

Banco de Crédito del Peru (BCP), Senior Data Scientist — November 2019 - May 2021

- Slashed lead evaluation time from two days to under one minute using credit-score and income forecasts.
- Improved customer acquisition accuracy by 10% by implementing advanced customer-transaction graph modeling.
- Cut data analysis time by 30% by designing a reusable univariate and multivariate analysis toolkit.
- Halved model metrics tracking time by 50% through the development of an IFRS-9 scoring toolkit.
- Evaluated models via portfolio projections to forecast shifts up to 6 months ahead.

Tech Stack: Python, PySpark, SQL, Oracle, SAS, Cloudera.

Rimac Seguros y Reaseguros, Data Scientist — August 2018 - November 2019

- Halved fraud investigation times from six to three months using graph modeling for vehicle policies.
- Drove a 20% sales increase by implementing personalized, end-to-end pricing models for vehicle insurance.
- Slashed model pre-processing time by a factor of 20x by effectively leveraging PySpark on AWS.

- Automated the monthly and on-demand execution of predictive models, improving both efficiency and reliability.
- Tech Stack: Machine Learning, Python, PySpark, SQL, Oracle, AWS.

Accenture, Application Development Associate — January 2018 – August 2018

- Cut daily back-office data collection time from four hours to just ten minutes using Python.
- Developed AI proof-of-concept business cases and custom applications across the full project lifecycle.

Tech Stack: Python, SQL, Automation Anywhere, AWS.

University of Engineering and Technology (UTEC), Research Assistant — April 2016 – January 2018

Project: Post-Stroke BCI — State-funded EEG interface enabling communication for hemiplegic patients.

- Enhanced Brain-Computer Interface performance by 10% using advanced digital signal processing and ML algorithms.
- Built a custom C++ user interface for brain signal acquisition based on ERP and sensorimotor rhythms.
- Optimized EEG data acquisition and channel choice by applying Python's statistical and CSP methods.

Tech Stack: Python, TensorFlow, Torch, C++, MATLAB, EEG Signals Analysis.

Universidad Nacional de Ingeniería, Undergraduate Research — April 2014 – December 2015

Project: Mining Robot — Autonomous system to detect toxic gases in underground mines.

- Improved the robotic exo-arm's movement precision by 5% by implementing Fuzzy Logic and Advanced Controllers.
- Developed a real-time object recognition system for polygons and hand gestures using Python and OpenCV.
- Designed and implemented robust navigation and route planning algorithms using both MATLAB and C++.

Tech Stack: Python, Genetic Algorithms, C++, MATLAB, Robotics.

CONSULTING & TEACHING EXPERIENCE

Coderio, Data & AI Consultant — October 2025 – Present

- Built an end-to-end finance system for Andes Logistics, optimizing 9-country data consolidation from days to minutes.
- Established Banco Basa's medallion lake governance by developing data lineage and automated quality/security KPIs.
- Engineered an ML logistics forecasting model for Santillana, saving ~\$50,000 in printing by predicting quote sales.

Tech Stack: Python, DAMA, SageMaker, AWS.

Tottus (via Sapiens Digital Lab), Machine Learning Engineer — September 2025 – January 2026

- Built Marketing Forecast pipelines and enhanced Demand models, refactoring code to automate production releases.
- Engineered a GenAI error notification agent using Gemini and Google ADK, optimizing code for scalable integration.
- Established MLOps standards in Vertex AI for Model Monitoring and Registry, ensuring observability of deployed models.
- Modernized Datahub/VPX workflows and resolved CI/CD critical bugs, confirming Churn/CLV models ensuring stability.

Tech Stack: Python, Docker, LLMs, Gemini, Vertex AI, GCP.

Apple (via Turing), LLM Data Scientist & Engineer — November 2024 – April 2025

- Developed over 300 reasoning notebooks for training LLMs on complex subjects like ML, robotics, and control systems.
- Covered prompt engineering, data structures, and data science with validated code and solution steps.
- Implemented robust QA frameworks ensuring code and reasoning correctness across hundreds of generated datasets.
- Built content generation pipelines using OpenAI, DeepSeek, and Gemini to ensure linguistic accuracy.

Tech Stack: Python, LLMs, OpenAI, DeepSeek, Gemini.

DELOSI, Staff Data Architect — October 2024 – March 2025

- Implemented monitoring and alerting for cloud state machines, reducing review and error-handling effort by 60%.
- Automated the backup process for over 5,000 on-premises tables to the cloud, reducing development time by over 100%.
- Created Cloud Formation templates for new pipelines using Step Functions, Glue, Lambdas, and other AWS services.
- Implemented over 20 different architectures for different business units.

Tech Stack: Python, PySpark, SQL, AWS.

Fivvy, Senior Data Scientist Consultant — July 2023 – October 2023

- Led data modeling for the "Contextual Profiler" app, leading to its acquisition by two businesses.
- Developed a local Streamlit application to effectively visualize and present the performance of various models.
- Formulated and deployed solvency and customer acquisition models for financial products using AWS Glue.
- Modified complex budgetary models to accommodate new and requested functionalities within the application.

Tech Stack: Python, Docker, MySQL, AWS.

BPC Business School – La Molina National Agrarian University (UNALM), Professor — November 2020 – Present

- Directed and taught over 10 Advanced Analytics Specialization programs, developing all course materials.

Christian Jeffrey Fonseca Rodriguez

AI Engineer | Data Architect | Data Scientist

Lima, Peru (-05:00 GMT)

(+51) 920202310

christian.fonseca.r@gmail.com

<https://linkedin.com/in/christian-fonseca-rodriguez>



- Successfully trained more than 100 professionals from diverse industries in data analytics and Big Data concepts.
- Provided personalized, in-depth feedback to guide professionals in advancing their unique analytical projects.
- Delivered specializations in Big Data & Analytics, Data Science for Business and AI for Business.

Tech Stack: Python, TensorFlow, Torch, LLMs, LangChain, LangGraph, SQL.

Data Mining Consulting (DMC), Professor — October 2021 – Present

- Authored and instructed over 20 advanced analytics programs, creating all original course material for each edition.
- Established the institutional baseline for supporting and reviewing MLOps tests and cloud deployment projects.
- Key programs delivered include Data Engineering (AWS & Azure), MLOps & ML Engineering, Deep Learning, Time Series & NLP, Big Data, and Advanced Python (ETL, Web Scraping).

Tech Stack: Machine Learning, Deep Learning, Python, PySpark, TensorFlow, Torch, Git, MLFlow, Docker, AWS, Azure, GCP.

CERTIFICATIONS

Amazon Web Services (AWS) – 13x Certified.

- **Professional:** Solutions Architect (2025) | DevOps Engineer (2025).
- **Specialty:** Machine Learning (2025) | Advanced Networking (2025) | Security (2025).
- **Associate:** Machine Learning Engineer (2025) | Solutions Architect (2025) | Data Engineer (2025) | CloudOps Engineer (2025) | SysOps Administrator (2025) | Developer (2025) .
- **Foundational:** AI Practitioner (2025) | Cloud Practitioner (2024).

Microsoft – 4x Certified.

- **Associate:** Azure AI Engineer Associate (2025) | Azure Data Scientist Associate (2025).
- **Fundamentals:** Azure AI Fundamentals (2025) | Azure Data Fundamentals (2024).

Oracle – 4x Certified.

- **Professional:** AI Vector Search Data Science (2025) | OCI Data Science Professional (2025) | OCI GenAI Professional (2025).
- **Associate:** OCI AI Foundations Associate (2025).

Specializations & Nanodegrees

- **Udacity:** Computer Vision Nanodegree (2021) | Deep Learning Nanodegree (2021).
- **Coursera:** Mathematics for Machine Learning Specialization (2021) | TensorFlow: Data and Deployment Specialization (2021) | Deep Learning Specialization (2018).

SKILLS

GenAI, LLMs & Agents: RAG Architectures, LangChain, LangGraph, Google ADK, Prompt Engineering, Vector Databases (Pinecone, FAISS), Ollama, Hugging Face. LLM APIs & Platforms: OpenAI, Gemini, Anthropic, Bedrock, DeepSeek, xAI.

Machine Learning & MLOps: TensorFlow, PyTorch, Scikit-learn, MLflow, Vertex AI, Amazon SageMaker.

Programming Languages: Python, TypeScript, R, MATLAB, C/C++, C#, Rust.

Web & API Development: React, Next.js, FastAPI, RESTful APIs, Microservices.

Cloud Platforms: AWS, Azure, GCP, OCI.

Data Engineering & Databases: SQL, PySpark, Databricks, PostgreSQL, Oracle, Redshift, SQL Server, MongoDB, RDS/Aurora.

DevOps & Tools: Docker, Kubernetes, Git, CI/CD (Bitbucket/GitLab/GitHub Actions), Terraform, CloudFormation.

Languages: Spanish (Native), English (Advanced).