

Carlos Noe Lopez Mejia

+49 15258443438 • carlos.mejia@tum.de • carloslmescom@gmail.com • [Website](#) • [LinkedIn](#) • Munich, Germany

PROFESSIONAL SUMMARY

Cloud & Data Engineer with hands-on experience delivering AWS infrastructure-as-code, multi-environment data platforms, and batch pipelines (EMR/Spark) at enterprise scale, plus a background designing DevOps workflows, Terraform modules, and IAM/security hardening. Comfortable working independently with minimal supervision as an individual contributor, proven by shipping production infrastructure at BMW and by architecting, building, and maintaining a full-stack platform (React/Supabase/Vercel) solo. Currently completing a M.Sc. in Informatics at TUM Munich.

EDUCATION

Technical University of Munich, Munich, Germany **Oct 2023 - Expected Apr 2027**

Master of Science in Informatics

Coursework: Cloud Computing, Cloud Information Systems, Cloud-Based Data Processing, DevOps: Engineering for Deployment & Operations, Machine Learning, Requirements Engineering, Software Engineering

Instituto Politécnico Nacional, Mexico City, Mexico **Jan 2015 - Jul 2020**

Bachelor of Computer System Engineering

Coursework: Artificial Intelligence, Software Engineering, Business Administration, Distributed Computing, Microcontrollers and Databases.

Universidad de Caldas, Manizales, Colombia **Jan 2019 - Jul 2019**

Student Exchange in Computer System Engineering

Coursework: Machine Learning

EXPERIENCE

BMW, Germany **Jun 2025 - Sep 2026**

Cloud DevOps Engineer - Working Student

- Design and maintain AWS infrastructure as code across two production services: **Terraform** for the Road Ahead Backend (Pothole Map / Dangerous Curve Assistant) and **AWS CDK** for the Charging Enhancement Layer (CPCL), plus **AWS CLI** for operational tasks, building reusable modules/constructs and staged environment configurations (int/dev/prod) with controlled backend/state management and rollout steps.
- Build and support batch data pipelines, including **EMR/Spark** job orchestration, **AWS Glue**, **Python**, **Scala**, and **Spark** job artifacts publishing for enterprise-scale autonomous driving and electromobility data workloads.

- Configure **IAM** roles and policies for cross-service access, secure data publishing, and least-privilege permissions across environments; contribute to broader security hardening (**KMS, Secrets Manager, VPC**).
- Publish new data-product workflows, including semantic data distribution patterns for downstream consumers in the **US** and **EMEA**.
- Operate and improve **observability/monitoring** infrastructure, including **alerting** thresholds with **Athena, Step Functions, QuickSight**, and defect analysis, to strengthen platform reliability.
- Produce architecture **documentation**, implementation proposals, and technical handoffs; support incident/investigation work with root-cause analysis and evidence-based reporting.

Ocumeda, Germany Jun 2024 - Jun 2025

Data Engineer and Analytics - Working Student

- Collaborated closely with stakeholders, including C-level executives (**CTO**), to gather and translate business requirements into scalable data solutions.
- Designed and implemented a **scalable data architecture** on AWS (S3, Lambda, Redshift, Step Functions) using Airbyte and Metabase to support the startup's rapid growth.
- Developed around **10 ETL pipelines** using Python, connecting different data sources as external APIs and SharePoint, some resulting in a **400% reduction in execution time**, significantly improving data processing efficiency.
- Established data architecture best practices and standards, ensuring **data quality, consistency**, and governance across the organization.
- Monitored and optimized data pipeline performance, proactively identifying and resolving bottlenecks to meet data delivery **SLAs**.
- Managed the data infrastructure that supported data-driven decision-making for over **15 users**, including C-level executives.

ITESM (Instituto Tecnológico y de Estudios Superiores de Monterrey), Mexico Feb 2022 - Sept 2023

Part-Time University Lecturer

- Invited lecturer at the Master's in Applied Artificial Intelligence (MNA-V) at Tecnológico de Monterrey for **6 weeks in July-August**. Course: Introduction to MLOps.

- Participated in the proposal and design to register the first MLOps subject as part of the ITESM master's program called Master in Applied Artificial Intelligence (MNA-V).
- Part-time lecturer in programming subjects with **Python, OOP, Unit and Integration Testing with Pytest, Packaging with Python, API REST with FastAPI, Software Engineering, Monitoring, Observability, Docker, Docker Compose, Terraform, and GitHub Actions**, applied to a certification course in Deployment of Machine Learning Models in Production Environments (MLOps).

Delta Smith (Remote, USA) Sep 2021 - Sep 2023

Data Engineering and Architect - Freelancing

- **Designed Data Roadmap for 2022 (Q2 - Q4) and 2023**, for Security, Observability, Monitoring, Orchestration, Data Governance, DataOps, Data Quality, Privacy, and Compliance.
- **Lead data projects** from requirements mapping to database modeling and delivery.
- Improve the design, documentation, maintenance, monitoring, and optimization of the existing code and data tools.
- Generated three workflows orchestrated by **Airflow Google Composer**.
- **Architected and implemented** an end-to-end batch data pipeline with **ScreamingFrog, Google Cloud Storage, GSheets, Docker, and Terraform**.
- Architected and developed **Data Observability, Logging, and Monitoring** infrastructure using **Google Cloud Functions, BigQuery, Data Studio, and Shipyard**.
- Designed and implemented **data security baselines**, such as a training course on best practices and authentication (MFA, 2FA).
- Implemented a batch **pipeline orchestration** with **Shipyard, Cloud Functions, and BigQuery**.
- Supported and improved Analytics Engineering team processes, such as ETL and ELT.

Wizeline, Mexico City, Mexico Jan 2022 - April 2023

Data Engineer

- Developed **one streaming** and **two batch pipelines** to extract, transform, and load (ETL) data from multiple data sources to a system recommendation using **Python**, both clouds **Google Cloud Platform** (Cloud Functions, Cloud Storage, Scheduler, Cloud SQL, Pub/Sub, Big Query, Dataflow, Scala) and **Amazon Web Services** (EMR, Pyspark, S3, Athena, Lambda functions, CLI).

- **Documentation of data requirements** from stakeholders and business areas such as Operations, Data Science, and Analytics, among others.
- Taught Terraform content about data engineering in an external boot camp.
- Taught Kafka fundamentals in an external course on streaming content.
- Designed and implemented multiple **unit and integration tests** with **pytest**.
- Created a **REST API** for a **machine-learning model** for the Wizeline internal project, using **FastAPI, Python, Docker, and AWS EC2**.

iVoy - On-Demand & Same-Day Deliveries, Mexico City, Mexico Feb 2021 - Jan 2022
Data Engineer

- Collaborated in the DE architecture design and implementation. Multiple proofs of concepts about data **batching** and **streaming** (**Apache Beam**), using cloud platforms (**GCP and AWS**), orchestration (**Airflow**), data preprocessing (**Spark**), data requirements from stakeholders, and data migration.
- Creation of a **streaming pipeline** and **production deployment** that processes all the packages per day, which feeds internal microservices from other business areas, using **Apache Beam, BigQuery, and Dataflow**.
- Consolidate data from multiple databases such as **MySQL, Postgres, and MongoDB** into a **Kafka** cluster (500M+ records and continuous data streams).
- Creation of a **Data Lake** for raw data in **AWS S3** and consumption with **Glue, Athena, and QuickSight** to analyze data.
- **Documentation of data requirements** from stakeholders and business areas.
- **Data migration** from an **AWS S3** bucket to another endpoint bucket using Python and **Boto3**.

Uber Technologies Inc, Mexico City, Mexico Oct 2019 - Jan 2021
Driver and Rider Operations Intern

- Responsible for **data collection, cleaning, exploration, and analysis** on large and complex data sets **about documents, trips, and profile information of drivers, partners, and vehicles** in all 56 Mexico Rides operational cities, used for regulatory processes using **SQL (Presto), Python, Pandas**, and internal tools.
- Productized a fully automated end-to-end **GUI application** to automate the **pre-processing of driver, vehicle, and partner documents** (massive image validation, compression, and conversion). The data was delivered monthly to governments

through a REST API included in the application, **saving 40+ hours** per month of manual work using **Python, Django, and SQL**.

- Converted a manual spreadsheet into an **automated dashboard** that reports the regulatory status of Mexican cities for **drivers and vehicles** and derives relevant insights every month to monitor and support the implementation of business strategies in the compliance status using **Python, SQL, and Google Sheets**.
- Developed a **document processing automation prototype** to avoid document rejection for **drivers and partners** for internal audits and increase information legitimacy of National Identification documents using **Deep Learning** (supervised learning, ResNet-50), **Computer Vision**, and **OCR**. (81% accuracy across four categories).
- Generation of **high-level documentation** from Engineering teams in San Francisco for non-technical people in Rides Ops Mexico, Regulatory, and Marketplace, among others.

Neimpetu S.A.S, Manizales, Colombia May 2019 - Sep 2019

Software Developer

- Analysis of user requirements with clients.
- Designed, developed, and tested more than **30 modules, CRUD frontend and backend**, search/sort/filter functionalities, and general page navigation using **Angular, Typescript, Java, Java Spring Tool, and SQL Server**.

T&M Consulting, Mexico City, Mexico Jun 2018 - Nov 2018

Web Application Developer

- Survey creation for company clients in different stages, like design and testing in production with JavaScript, CSS, HTML, Pepperi, and Visual Studio Code.
- Experimentation with Progressive Web Apps for internal testing using Service Worker and Manifesto.

Propulsar MX, Mexico City, Mexico Dec 2016 - Jul 2018

AI Intern Developer

- Created **16 chatbots** to identify and save customer intentions in SQL Server in written messages using **Natural Language Processing, C#, Microsoft Bot Framework, SQL Server, and Azure Cognitive Services (API LUIS)**.

TECHNICAL SKILLS

- **Cloud Platforms:**
 - AWS: S3, Lambda, Redshift, Step Functions, EMR, Athena, Glue, EC2, QuickSight, IAM, KMS, VPC, Secrets Manager, Systems Manager.

- GCP: BigQuery, Cloud Functions, Cloud Storage, Composer, Dataflow, Pub/Sub, Compute Engine, Google Domains
- **Data Engineering & Processing:**
 - Languages/Core: Python (Pandas, NumPy), SQL
 - Big Data Tools: Spark (PySpark), Kafka, Apache Beam
 - Orchestration & ETL/ELT: Airflow, ETL/ELT Design, Airbyte, dbt
 - Concepts: Data Architecture, Data Modeling, Data Governance, Data Quality, DataOps
- **MLOps & Development:**
 - Infrastructure & CI/CD: Docker, Terraform, AWS CDK, Git, GitHub Actions
 - Testing & APIs: Pytest, FastAPI, Django
 - Monitoring & Observability Practices
- **Machine Learning & AI:**
 - Libraries: Scikit-learn, TensorFlow, Keras, OpenCV
 - Domains: Deep Learning (CNN, ResNet), NLP, Computer Vision (Image Analysis, OCR, Segmentation), nnU-Net
 - Concepts: Supervised Learning (Regression, Classification, SVM)
- **Web Technologies & Other:**
 - Languages: Javascript, Java, C#, Typescript
 - Frameworks: React, Angular
- **Databases:**
 - Relational: SQL Server, MySQL, PostgreSQL, Redshift
 - NoSQL: MongoDB

PROJECTS

Salsuki, Community Management Platform (Full-Stack, Solo Engineer) 2024 - Current

- Sole full-stack developer and product owner for a community management platform serving a non-profit salsa school with 60+ active registered users (bookings, skill

levels, wallet balances tracked as system of record).

- Built the frontend with **React 19, TypeScript, Vite, and Tailwind CSS 4** (PWA, mobile-first) and the backend on **Supabase** (Postgres with Row-Level Security, serverless Edge Functions in Deno, Auth, Storage).
- Deployed on **Vercel** with per-PR preview environments, **GitHub Actions** for lint/build/migration validation, and **Sentry** for production error monitoring.
- Designed a credit-based internal economy (“wallet” balances) with asynchronous bank-transfer reconciliation via CSV processing, replacing manual, spreadsheet-based payment tracking.
- Built class/event booking with capacity limits, waitlists, and skill-level gating, plus an integrated learning-resource library and an interactive knowledge-graph view (XyFlow) of dance-figure prerequisites.
- Run an AI-agent-orchestrated engineering workflow: scope structured tickets, dispatch implementation to Claude Code agents on isolated branches, and review every PR before merge, sustaining ~15-20 shipped tickets/month as a solo maintainer.
- [Website](#)

DevOps Course - Development & Deployment of a GenAI-Powered Web Application Apr 2025 - Aug 2025

- Designed and implemented the Retrieval-Augmented Generation (RAG) and Generative AI subsystem for LECTure-bot, an intelligent course material assistant.
- Built a Python-based LangChain microservice to process and embed uploaded documents, store/query vectors in Weaviate, and integrate with cloud-based and local LLMs for context-aware Q&A and flashcard generation.
- [GitHub Repository](#)

Seminar Course - Natural Language Processing: Methods and Applications (Summer 2025) Apr 2025 - Aug 2025

- Conducted independent research on adversarial attacks targeting large language models, surveying vulnerabilities, and evaluating existing and proposed mitigation strategies.
- Produced a [seminar paper](#) and delivered a presentation summarizing findings on securing LLMs against malicious exploitation.

Practical Course - Recommender Systems, TUM Connected Mobility Lab (Summer 2025) Apr 2025 - Aug 2025

- Designed and implemented full-stack architecture for an AI-powered travel journal assistant, integrating React frontend, FastAPI backend, and PostgreSQL database with Dockerized deployment workflows.
- Led development of DevOps cycle and GenAI microservice for image analysis and metadata enrichment, enabling automatic timeline generation and personal information refinding based on photo uploads.
- [GitHub repository](#).

Interdisciplinary Course - LLMs (Skills: Python, LLMs, Software Engineering) Oct 2024 - May 2025

- Creating a framework that reduces the manual expenditure to **take a tool from its research stage to a professional level using LLM models (Python and Gemini)**.
- [Github repository](#).

Practical Course - Applied Deep Learning in Medicine (Skills: Python, DL, Segmentation, nnUNet) Apr 2024 - Aug 2024

- Performing research practical course/project with the Chair for AI in Medicine and Healthcare at TUM for Analysis of Resolution Differences on a Segmentation Task and state-of-the-art techniques.
- [Github repository](#).

Responsive Website (Skills: Javascript, Node.js, GCP, GitHub, HTML, CSS) Feb 2023 - Mar 2023

- Designed and implemented a responsive website with multi-language support, using **HTML, CSS, Javascript, Node.js, GitHub, GCP Compute Engine, and Google Domains**.

Bachelor's Thesis (Skills: Python, Django, Machine Learning, Image Analysis) Jul 2020

- A methodology that allows the detection of three types of diseases in tomato leaves by **image analysis** and **pattern recognition**. Experimentation and generate an embedded algorithm with supervised learning algorithms: Neural Networks, Support Vector Machines, Random Forest, K-Nearest Neighbors, Naive Bayes, and Convolutional Neural Networks.

Classification of types of food (Skills: Python, Flask, Flutter, CNN) Jul 2019

- Created a production-ready end-to-end system that predicts food scenarios, developed with Python and CNN, and
- deployed it on a Heroku server. This one was also connected to an Android app developed in Flutter.

PUBLICATIONS

- LATE BLIGHT SEGMENTATION IN TOMATOES USING SPATIAL DOMAIN METHODS AND COLOUR ANALYSIS **Oct 2020**
[INTERNATIONAL JOURNAL OF SCIENTIFIC RESEARCH](#)

AWARDS

- Second place in Social Data Challenge 4.0 by DataLab Mexico. **2021**
- GIST Tech-I Semifinalist, Mexico City, Mexico. **2016**
- Babson Build for Santander Universities Entrepreneurial skills course. Boston, Massachusetts, US **2015**
- TrepCamp Summer Program Class of 2014. Silicon Valley, California, US. **2014**
- Winner IDEup Spin Banco Santander for best entrepreneur idea. Mexico City, Mexico. **2014**

LANGUAGES

Spanish: Native

English: Fluent - IELTS Academic Score 6.5

German: A2 (self-assessed); TUM institutional coursework completed: A1.1, A1.2

VOLUNTEERING

Salsa Instructor & Founder, Salsuki (Munich, Germany) Nov 2023 - Current

- Founded and run a free community salsa school (currently 87 members across in-person classes every Friday), sharing Latino culture and fostering networking through dance. Digitized and scaled the operation with a self-built platform (see **Salsuki** under Projects).

LATUM (TUM Munich, Germany) May 2024 - Current

- Proud member of the inaugural staff generation at Latin American at TUM (LATUM). Contributed to numerous activities to support Latin American students and other regions of the world. From organizing engaging workshops and providing essential guidance to facilitating intercultural exchanges and social gatherings.

Buddy Program (IPN Mexico City) Jul 2018 - Jul 2021

- Support exchange students from countries such as France, Belgium, Colombia, Chile, and Spain in integrating into the university and local cultural lifestyle.

COURSES

Introduction to Machine Learning in
Production Machine Learning Data
Lifecycle in Production Machine Learning
Modeling Pipelines in Production MLOps
Essentials: Model Development and
Integration Deploying Scalable Machine
Learning for Data Science Wizeline
Bootcamp Data Engineering Artificial
Intelligence Foundations: Machine
Learning

2023 2023 2023 2023 2022 2021 2018
