Felipe Chavarro Polanía

AI product development and solutions architecture

Bogotá, Colombia | felipechavarropolania@gmail.com
linkedin.com/in/felipechavarro | felipechavarropolania.com
github.com/felipechavarropolania

Professional profile

AI process automation specialist with HPE and Hewlett-Packard experience since 2012 across product strategy, hybrid cloud architecture and enterprise delivery. Current focus on AI products and services. Independent technical work includes AI evaluation, Python research tooling and reproducible experiments.

AI software factory

Created an AI-powered software factory that automates software development and enables Product Management to build robust prototypes in weeks, accelerating time-to-market and time-to-value.

Core capabilities

AI process automation · AI products and services · Product strategy and roadmaps · Solutions architecture · Hybrid cloud · Technical presales · Service lifecycle management · AI evaluation · Harness engineering · Experimental design · Python research tooling

Professional experience

Hewlett Packard Enterprise / Hewlett-Packard
August 2012 to present

WW Service Offer Manager AI / AI Expert
January 2025 to present

- Work across AI product development, service offerings and customer enablement, connecting AI capabilities with enterprise needs.
- Define product strategy and communicate roadmaps for AI initiatives, including LLMs, AIOps and intelligent automation.

Earlier roles within HPE
WW Business Segment Manager; Practice Delivery Leader for LATAM; Cloud Engineer / Hybrid Cloud Solutions Architect; Lead Technical Consultant.

- Managed service lifecycle planning and coordination across global teams; used business analysis to identify improvement opportunities.
- Shaped delivery strategies with presales teams, aligning customer expectations, available resources and regional opportunities.
- Conducted cloud maturity assessments, solution architecture, cost-benefit analysis and technical sales demonstrations.
- Worked with executive stakeholders on technology strategy and migration roadmaps; handled customer escalations and technical enablement.

Independent technical work and research

Governed AI Systems

Python research toolkit using JSON and CSV inputs and deterministic decisions from declared policies.

- Recipe Discovery generates fractional-factorial designs and analyzes replicated results.
- Compute Governance applies promotion rules and checks complete result matrices.
- Judge Audit preserves response identity, reconstructs requests and distinguishes decision changes from technical failures.

Repository examples are synthetic demonstrations. Empirical evidence is reported in the preprints below.

github.com/felipechavarropolania/governed-ai-systems

Public preprints

Staged Factorial Screening for Budget-Constrained Micro-Pretraining

Sole author, arXiv, 2026. Bounded methods study with 613 experiments, repeated checks and continuations on two hosts. Examines staged screening and confirmation under constrained compute.

arxiv.org/abs/2606.05186

Small Experiments, Cheaper Decisions: A Case Study in Staged Promotion for Micro-Pretraining

Sole author, arXiv, 2026. Studies auditable promotion rules with 12-hour confirmation across two hosts. The full protocol records 169.2 training GPU-hours. Findings are limited to the tested setting.

arxiv.org/abs/2606.11387

Independent voice AI projects

Produced more than 100 English and Spanish audiobooks using ElevenLabs since 2024. Tested v2 voices and checked generated audio for artifacts.

Additional experience

Udacity | Tech Mentor | 2021 to 2025

Reviewed code and projects; mentored learners in cloud computing, data science and programming.

IBM | IT Specialist / Presales Engineer | 2011 to 2012

Technical work across AIX, Linux, Citrix, VMware, virtualization and enterprise infrastructure.

Education and languages

Master of Science | University of Liverpool | 2013

Telecommunications Engineering | Universidad Piloto de Colombia | 2008

Spanish: native · English: full professional proficiency