SIVANANDA (SIVA) RAJANANDA

New York, NY · SivaRajananda.com · linkedin.com/in/sivananda-rajananda

Senior Data & Applied Scientist, 4.5 years building and operating production ML for e-commerce advertising and operations. Owned a conversion-rate system that served up to ~20M predictions a day at peak, and took an inherited LLM customer-service agent from pre-production to a live brand pilot.

EDUCATION

Harvard University · M.S., Data Science 2020 – 2021

UCLA · B.S., Cognitive Science, Computing specialization 2013 – 2017

EXPERIENCE

Pattern® (e-commerce accelerator)

Lehi, UT and Hong Kong · Feb 2022 – present

Senior Data & Applied Scientist Sep 2025 – present

- **Conversion-rate prediction system** (owned since Feb 2023, through two role changes): CatBoost models across Amazon, Walmart and Tmall (migrated from XGBoost), with production training moved from CPU to multi-GPU SageMaker training jobs (Airflow-orchestrated, artifacts in S3); daily batch runs forecasting conversion rates 36 hours ahead at hourly granularity, up to ~20M predictions/day on the largest pipeline, feeding live ad bids.
- **Data-quality gates that fail the pipeline closed**, built after tracing a data incident to an upstream join that turned zeros into NULLs; prediction-similarity check before model promotion.
- **Designed, validated, and delivered the promotion-lift model** for the demand-forecasting platform: hierarchical Bayesian negative-binomial uplift model, pre-registered convergence gates, coverage-checked 90% prediction intervals (87–90% empirical), served as a precomputed elasticity table; built a production LLM reasoning layer that generates a second, independent set of lift predictions with written rationale, covering sparse-history cases.
- **System review:** architecture-level review of Pattern's core advertising ML system (~30k lines); mapped its full data flow for the Operations team.
- **LLM tooling:** built the Model Context Protocol (MCP) integration exposing the inventory engine to the company's AI agent platform (now in production); shipped three internal LLM skills and the team's skills repo, and taught the team how to build their own MCP servers and wire them into Pattern's AI agents.

Domain Lead, East Asia Operations Hong Kong · Feb 2025 – Sep 2025

- **Production LLM agent:** took over a pre-production customer-service agent (LangGraph, RAG over a Chroma vector store with OpenAI embeddings, AWS) and drove it to a live brand pilot; built a failure-driven regression set re-run before every prompt iteration.
- **Trained a remote successor to full ownership of the conversion-rate pipelines**: daily 1–2 hour code walkthroughs over 1–2 months, paired with a live AWS Batch to Outerbounds migration as his own deliverable. He now handles incidents and ships extensions independently.
- **China Ads proof-of-value** (funded): sole ads-domain expert; built the prototype pipeline (PostgreSQL, AWS China) and ran it with a human-in-the-loop bidding process. Presented to brand advertisers in Mandarin.

Data Scientist Feb 2022 – Feb 2025

- **Excess-inventory decision engine:** built a ~3,000-line SQL v1, then evolved it into a parallel Python simulation costing 8 disposition actions across 36 scenarios per ASIN over 2-year horizons (~600k simulations, under 3 hours per run) on ~$50M of excess inventory. Handed to its domain team in 2024, in production since; brought back in 2026 to own and extend it.
- **Event-driven inference:** deployed PyTorch image-enhancement CNNs as scale-to-zero serverless SageMaker endpoints triggered by S3 events; migrated 6 ML pipelines to Airflow, across four orchestration generations (EC2 → Batch → Airflow → Metaflow).
- **Organic-lift attribution model** (decay-weighted pre/post) feeding a company-wide advertising-return metric; listing-impact and keyword-ranking tools used by SEO teams internationally.

ASML · Data Science Intern

Jun 2021 – Aug 2021

- CNN ensemble defect-detection models for inspection imagery (TensorFlow; 97% accuracy, balanced data), shipped as an executable; full-marks review and return offer.

UCLA · Research Associate

Jun 2019 – Aug 2020 (and Apr 2017 – Sep 2017)

- Modeled metacognition with recurrent neural networks in PyTorch; cleaned, linked and analyzed >13 million California birth and medical records in Stata.
- Built and open-sourced the first browser-based random-dot kinematogram (jsPsych); used by vision labs worldwide since 2017.

UC Riverside · Lab Manager & Programmer

Sep 2017 – Jun 2019

- Built the lab's research infrastructure end to end (Linux, MySQL, mTurk); published the first open-access PEST staircase in JavaScript; presented the tools at the Vision Sciences Society conference.

PUBLICATIONS

8 peer-reviewed articles, 203 citations; first-author on confidence and metacognition; contributing author, *Nature Communications* 2023, *Psychological Review* 2026.

SKILLS

Languages: Python · SQL · C++ (coursework) · Mandarin and Malay (professional working proficiency)

ML / LLM: CatBoost (multi-GPU) · XGBoost · scikit-learn · PyTorch · TensorFlow · pandas, NumPy, SciPy · feature engineering · hierarchical Bayesian modeling · uplift modeling · CNNs, RNNs · simulation and optimization · ads ML (bidding, attribution, ROAS) · Large Language Models (LLMs) · LangGraph · RAG (Chroma, OpenAI embeddings) · MCP · prompt engineering · LLM evaluation and regression test sets

Platform / MLOps: AWS (SageMaker, S3, Batch, Lambda, ECR) · Metaflow/Outerbounds · Airflow · MLflow · Snowflake · PostgreSQL · Docker · Terraform · GitHub Actions · CI/CD · model monitoring and data-quality gating · Streamlit · AI-assisted engineering (Claude Code, Cursor, Copilot)