

Aiesha Lanba

(425) 420-3642 — aiesha.lanba@berkeley.edu — linkedin.com/in/aiesha-lanba

Education

University of California, Berkeley

Expected Graduation: December 2026

Bachelor of Arts in Data Science | Concentration: Business and Industrial Analytics

Berkeley, CA

Relevant Coursework: Introduction to Machine Learning, Data Engineering, Principles & Techniques of Data Science, Data Structures, Structure and Interpretation of Computer Programs, Data Mining and Analytics, Concepts of Probability, Linear Algebra and Differential Equations, Principles of Engineering Economics

Technical Skills

Languages & Tools: Python (Pandas, NumPy, Scikit-learn, PyTorch, TensorFlow), SQL, Java, C++, Git, Jupyter, AWS (EC2, IAM), Spark, Databricks

Machine Learning & AI: Regression, Classification, K-Means, PCA, CNNs, RNNs, Transformers, GANs, LLMs, Statistical Modeling, MLE, Backpropagation, Optimization

Product & Program Management: Technical PRD Writing, Jira, Confluence, Cross-Functional Stakeholder Management, Product Strategy, Go-to-Market Analysis

Data Analysis & Visualization: EDA, Statistical Analysis, Tableau, Power BI, Bias Detection, Simulation Modeling

Experience

AI Product Strategy Intern

May 2026 – August 2026

T-Mobile

Bellevue, WA

- Built business case for drone edge-compute over cloud/on-device alternatives by quantifying GPU infrastructure savings per fleet, shifting product strategy to prioritize edge computing alongside 5G connectivity.
- Authored technical PRDs defining edge-compute requirements across 4 drone use cases in 3 weeks, translating one into an active Skydio demo and the other into an executive-ready value proposition for leadership.
- Presented edge-compute strategy to T-Mobile's Chief Data and AI Officer and VP of Data and AI, generating alignment, and sourced BRINC's VP of Engineering as a prospective customer to expand partner pipeline.

Data Analyst/AI Intern

June 2025 – August 2025

Wipro

Remote

- Modeled and optimized 6 agentic AI deployment architectures across compute, memory, and orchestration dimensions; identified adaptive designs that cut projected costs by 25–40%.
- Quantified enterprise scalability constraints across 10 operational dimensions, including compute utilization (50% static vs 80-95% adaptive), failure recovery costs, and human oversight burden, providing data-driven thresholds to prevent common pilot-to-production failures documented in industry.
- Designed governance framework with measurable SLOs for coordination budgets and auditability standards, informing operational controls for enterprise AI adoption.

Projects

Autodifferentiation Library (BearTensor) | Python, NumPy, PyTorch Architecture

April 2026

- Built custom n-dimensional tensor library in Python/NumPy, replicating PyTorch's autograd engine with dynamic computation graphs, operator overloading, and support for arbitrary tensor shapes.
- Implemented backpropagation via topological sort (Kahn's Algorithm) to propagate gradients through computation graphs, computing adjoints for vectorized operations using the chain rule.
- Developed SGD, Momentum, and Adam optimizers from scratch to automate parameter updates, applying adaptive learning rates and momentum terms to minimize training loss across deep networks.

Predicting Housing Prices in Cook County | Python, Pandas

November 2025

- Built linear regression model on 250K+ property records, engineering features through one-hot encoding, standardization, correlation analysis, and polynomial terms, reducing RMSE from 0.85 to 0.79.
- Designed custom fairness-focused error metrics beyond RMSE, uncovering that the model overestimated lower-value properties 58% of the time versus 31% for high-value properties, and recommended stratified evaluation to address valuation bias.