Priyal Sharma

Los Angeles, CA | priyalsh@g.ucla.edu | (925) 532-3838 | linkedin.com/in/priyalsh | github.com/preelshh | priyal-sharma.com

Education

University of California, Los Angeles — Expected June 2027
Bachelor of Science: Statistics & Data Science, Minor in Data Science Engineering — *Los Angeles, USA*

- **GPA**: 3.99 – **Awards**: Dean's Honors List (x6) – **Involvement**: DataRes @ UCLA, ACM @ UCLA, Creative Labs @ UCLA
- **Coursework**: Data Structures & Algorithms, Computer Science I & II, Systems and Signals, Machine Learning, Linear Algebra with Proofs, Probability Theory, Mathematical Statistics, Computational Genomics & Probabilistic Models, Regression & Data Mining

Skills

Languages: C++, Python, Java, C#, SQL, R, JavaScript, Bash, Luau
GPU & Distributed: PyTorch DDP, NCCL collectives, AMP mixed precision / Tensor Cores, multi-GPU embedding sharding, `nvidia-smi`
Frameworks & Libraries: PyTorch, PyTorch Lightning, TensorFlow, scikit-learn, XGBoost, Hugging Face, NumPy, Pandas, SciPy
Tools & Platforms: Git, Docker, MLflow, Spark, Databricks, Airflow, CI/CD, AWS (S3, EC2, SageMaker), GCP, Linux/Unix

Work Experience

Nike | *AI/ML Engineering Intern* — Jun 2026 – Present
Global sportswear company applying AI/ML across product, consumer, and supply chain — *Beaverton, USA*

- Re-architected **multi-GPU training** from **DataParallel** to **DDP**, coordinating one process per GPU over NCCL across 4 A100s, cutting training from 4 hours to 20 minutes per run, raising **GPU utilization** 30% $\rightarrow$ 85%, and **speeding inference** throughput 4x
- Lifted the **behavioral sequence model**'s NDCG@10 from 0.847 to 0.910 across 20 automated distributed sweeps over embedding dimension, transformer depth, BERT-style masking ratio, and optimizer choice, producing learned embeddings for 2.8M users
- Shipped **64-dim Behavioral Sequential Transformer post-MMoE embeddings** into **production DCNv2 L2 ranker** of **Nike Search**, driving an additional **+2.19pp NDCG@10 lift** over the Stage-1 A/B baseline with a projected $\sim$**$7.4M increase** in annual revenue
- Designed a **hypothesis-testing** program to explain why the best sequence-model checkpoint did not produce the best ranker, isolating three candidate causes and testing each: user memorization via held-out **overfitting checks**, redundancy with **existing features via Ridge, CKA**, and **MLP probes** ($R^2 = 0.004$, near-orthogonal), and data scarcity via corpus scaling from 2M to 3M users (**val AUC flat**), ruling out redundancy and volume as explanations

UCLA David Geffen School of Medicine | *Undergraduate Researcher* — Jan 2025 – Present
Computational and machine learning analysis of gene-regulatory programs in cellular quiescence and cancer — *Los Angeles, USA*

- Built a **regulon-inference pipeline** (**pySCENIC**, AUCell) over **scRNA-seq**, pruning $\sim$**6M** candidate TF-target edges into a network
- Extracted **attention-derived gene regulatory networks** from a **transformer foundation model** (scPRINT2) across 4 T-cell states, applying **weight-thresholded pruning** to produce cell-state-resolved graphs of up to **1.4M edges** spanning around 4K genes
- Created an **S3-based data ingestion** pipeline (boto3) to pull a 10,282-gene Perturb-seq dataset and align it with internal quiescence data, **normalizing** heterogeneous gene-ID schemes (Ensembl/symbol detection, version stripping, deduplication)

Pacifico Bio | *AI/ML Engineering Intern* — Jun 2025 – Aug 2025
Food biotech company using mycelium fermentation to address global food scarcity — *Berlin, Germany*

- Built a scalable **ImageJ-based data engineering pipeline** enabling automated image preprocessing with **2s per image** latency
- Shipped a multi-modal classification model (image and numeric data) with **92%** accuracy, cutting manual review time by **80%**
- Delivered an end-to-end full-stack software product using **Java** (GUI) and **Python** (inference) into a usable workflow for R&D

Giant Music LLC | *Data Engineering Intern* — Sep 2024 – Jun 2025
Music company using data-informed A&R and market insights to discover and develop emerging artists — *Los Angeles, USA*

- Automated a manual daily data pull into **SQL ingestion pipeline**, replacing a multi-hour analyst task with an unattended job
- Designed an internal artist-ranking engine with rule-based logic and **NLP** for efficient data-driven scouting **(175K $\rightarrow$ 50 artists)**
- Reduced workflow steps from **10** to **3** by coding custom text-cleaning + language-detection scripts using **Google APIs** and **regex**

Projects

Apex Volleyball: Neural World Model | *Python, PyTorch, NumPy, pygame* — Jul 2026 – Aug 2026

- Hand-implemented full **transformer inference** (attention, matmuls, layer norm) in $\sim$100 lines of **NumPy**, matching PyTorch to **5e-7 numerical parity** at **2ms/call** and cutting a **2GB** runtime dependency to **3MB** with no framework on the hot path
- Replaced a volleyball game's physics engine with the resulting **671K-param** world model, trained on **30K** simulated rallies

Physics-Informed Models for Glucose Forecasting | *Python, PyTorch, torchdiffeq, NumPy, Matplotlib* — Jan 2026 – Mar 2026

- Engineered **Physics-Informed Neural Networks** (**RNN, LSTM, Transformer**) for glucose prediction, scoring **0.50 mg/dL RMSE**
- Built **differentiable physics model** using **MATLAB/Simulink** Bergman ODE theory, implementing Neural ODEs via **torchdiffeq**

OncoReg | *Python, AlphaGenome API, scikit-learn, PyPI* — Jan 2026 – Feb 2026

- Shipped an **open-source Python** package scoring cancer-relevant genetic variants via **Google DeepMind's AlphaGenome API**
- Maintained on **GitHub** and **PyPI** with versioned releases, full documentation, and example notebooks for researcher onboarding

Additionals

Awards & Certifications: Kaggle Competition Winner, Generative AI (Databricks), Probability Theory (HarvardX), Akuna Capital Options
Interests: Piano, Video Game Development, Quantum Mechanics, 2D Animation, Music Production, Tennis, Reading, Watercoloring