

Priyam Ghosh

Bengaluru, Karnataka, India — +91 826-051-7343 — Email —

LinkedIn — GitHub — Portfolio

Professional Experience

University of California, Berkeley

Aug 2025 - Present

Research Fellow

Remote

- Designed and evaluated world-model architectures for spatial reasoning, improving long-horizon prediction accuracy by 18–25% on simulated navigation and embodied AI benchmarks.
- Led experiments across 10k–100k+ simulated trajectories, developing data pipelines and evaluation metrics that reduced training instability and cut experiment iteration time by 30%.
- Implemented relational spatial reasoning modules inspired by quantum entanglement to model object to object and scene level constraints, reducing combinatorial search overhead and cut experiment iteration time by 30%; collaborated within a 5–8 person research team on reusable research code and publications

National Aeronautics and Space Administration

Dec 2019 - Jul 2024

Machine Learning and System Architect

Hybrid

- Architected Project Callisto (internet-independent voice AI on a human-rated spacecraft); retrained Alexa wake-word CNN for 60–70 dB cabin acoustics, driving near-zero false-positive activations through hardware validation at Lockheed Martin (Denver).
 - Owned OpNav end-to-end for Artemis I; processed 1,000+ images across 26 passes, achieved 3σ certification in 4 passes, qualifying autonomous return navigation if DSN contact were lost at 384,000 km.
 - Flight-certified GN&C software to NASA NPR 7150.2 Class A within cFS; implemented UDU covariance factorization across 500+ Monte Carlo trajectories, cleared the pre-launch OIMU-1 anomaly at Flight Readiness Review, and delivered 0 open MISRA C:2012 Required-rule findings.
 - Built a DSN latency-tolerant transport layer and Artemis I BET reconstruction; replaced TCP at 2.56 s RTT with CCSDS-native FEC enabling live Mission Control video from 384,000 km, and fused 5 data sources post-mission into ground truth for Artemis II algorithm updates.
-

Technical Skills

- **Languages:** Python, Go, Rust, Javascript, Java
 - **Frameworks:** FastAPI, Express, Spring Boot, Axum, Next.js
 - **ML / AI:** Ray, Triton, Tensor, MXNet, Fortran
 - **Databases:** MongoDB, PostgreSQL, Neo4j, Redis, Vector Databases
 - **Others:** AWS, Docker, Git, Terraform, Latex, MATLAB
-

Projects

cheeger — Differentiable Spectral Segmentation Engine

GitHub — Blog

- Built a from-scratch spectral segmentation library in PyTorch — affinity kernels, graph Laplacian variants, k-NN sparsification, and a cyclic Jacobi eigensolver — using scipy/numpy only as test oracles.
- Derived a degeneracy-robust differentiable eigensolver with Lorentzian-broadened gap gradients stable at segment boundaries, verified via gradcheck; replaced fixed eigenbasis weights with a learned spectral response $h_\theta(\lambda)$ noise-floored by Marchenko–Pastur.
- Trained end-to-end as a U-Net head on Cityscapes; benchmarked against a conv baseline with mIoU/Boundary-IoU metrics and Rayleigh spectral-consistency loss.

ananta — Energy-Based Self-Learning Architecture for Scientific Reasoning

GitHub

- Designed a research-oriented LLM architecture combining energy-based modeling with autoregressive decoding.
- Implemented contrastive energy functions enabling iterative self-correction without external reward models.
- Built self-play training loops with LoRA/QLoRA fine-tuning to improve symbolic reasoning and formal problem solving.
- Applied system to scientific workflows including mathematical derivations and structured reasoning tasks.

engram — Redis-Wire Associative Vector Memory Engine

GitHub — Blog

- Built a Redis-wire-compatible (RESP) vector memory engine in Rust from scratch with zero dependencies beyond `libc`, working unchanged with `redis-cli` and existing clients.
- Implemented HNSW approximate nearest-neighbour index by hand and a single-threaded event-driven I/O core over raw `kqueue/epoll`, with vector search offloaded to a worker pool.
- Hand-wrote SIMD distance kernels in NEON (ARM64) and AVX2+FMA (x86_64) with runtime dispatch, achieving ~9× speedup at 768 dimensions; WAL + compacting mmap snapshot for durability.

poolgrad — Memory-Aware ML Runtime for Neural Network Performance

GitHub — Paper

- Built a minimal autograd engine in Rust to analyze performance bottlenecks in neural network training.
- Implemented reverse-mode autodiff over a dynamic computation graph with a custom kernel scheduler.
- Developed four matrix multiplication kernels (naive, tiled, packed+SIMD, Strassen-form variant) optimized using NEON (ARM64) and AVX2+FMA (x86_64).
- Designed a size-aware gradient memory pool with lifetime-based reuse, reducing allocation overhead and improving cache efficiency.

Neuro-Symbolic Agents with Neo4j

GitHub

- Built a hybrid reasoning system combining LLM-based inference with symbolic graph traversal.
- Modeled knowledge graphs in Neo4j and integrated with LangChain for context-aware agent workflows.
- Designed secure memory management pipelines for structured user data and long-term contextual reasoning.

Publications

- **The Dynamics of Memory: Replay, Diffusion, and Attractor Formation** — ICMNAI (2026) Paper
- **Differential Memory Decay: A Continuous-Time Model of Forgetting and Semantic Drift** — ICLR Workshop (2026) Paper
- **An Energy-Based Self-Learning Engine for Neuro-Symbolic Scientific Reasoning** — ICLR Workshop (2026) Paper
- **Explicit Temporal Attention in Video Diffusion Models** — CVPR Workshops (2024) Paper

Achievements

- 4th place, National Astronomical Olympiad '19 (IAO supervision).
- Top 2% Global rank, IAAC 2022 (Golden Honor award).
- Under Top 500 Global rank holder at IJSO-2018

Education

Birla Institute of Technology and Science (BITS), Pilani, India
B.Sc., Computer Science

Aug 2024 - Present

- GPA: 9.2/10.0
- Coursework: Data Structures & Algorithms, Operating Systems, Computer Networks, DBMS, OOP, Compiler Construction, Low-Level Design (LLD), High-Level Design (HLD), Spring Boot

Leadership & Volunteer Experience

Scaler School of Technology, Bengaluru
President, Reinforce AI/ML Club

May 2025 - Present

- Led 5+ campus-wide hackathons (500+ participants) in collaboration with leading companies, designing state-of-the-art problem statements.
- Conducted independent research on self-learning LLM architectures and authored proposals for next-generation RAG systems.
- Organized 4+ hands-on workshops on LLM architectures and core systems, mentoring 200+ students across campus and external communities.