Jaiprasad R

REINFORCEMENT LEARNING (RL) LEAD · MACHINE LEARNING ENGINEER

Hyderabad, India | (+91) 99012 75669 | jprampure@gmail.com | linkedin.com/in/jaiprasadr | github.com/jaiprasadreddy | Scholar

Full-stack Machine Learning (ML) Engineer / Reinforcement Learning (RL) Engineer with 10+ years' experience leading AI teams across research and transitioning systems from research to production, in healthcare, commerce and post-training data for frontier models. Currently RL Lead at Deccan AI, leading a 20-person team across all RL initiatives, with systems such as Stark, an agentic RL gym, and Media Studio, a Large Language Model (LLM) gateway serving 2M+ requests a day; also built a hybrid-search engine over 2M+ expert profiles and an AI voice interviewer. Over the last three years, consulted for multiple early-stage and pre-seed startups in the US and India, shaping their AI roadmaps and advising on AI architecture and pipeline design.

SKILLS

Languages	Python, C/C++, Java, CUDA, Bash, Git
ML / RL / LLM	PyTorch, TensorFlow, verl, Hugging Face, vLLM, NVIDIA Triton, LangChain, LlamaIndex, LiveKit (voice agents), OpenAPI-based mock environments, OpenCV
Retrieval & Data	PostgreSQL, pgvector, FAISS, Redis, Celery, hybrid search & LLM reranking
Infrastructure	Docker, Kubernetes, AWS (Lambda, Step Functions, SageMaker, Batch, Bedrock), GitHub Actions, CI/CD, FastAPI
AI Domains	Agentic RL & RL Gyms, RLHF, LLMs & LLM Agents, Synthetic Data & Task Generation, LLM Evaluation & Trajectory Analysis, Voice Agents, Retrieval-Augmented Generation (RAG) / Hybrid Search, Generative AI, LLMOps, MLOps, Natural Language Processing (NLP), Computer Vision, Deep Learning

EXPERIENCE

- **Deccan AI** Hyderabad, India
  RL Lead (Reinforcement Learning / Machine Learning Engineering) *Apr 2025 - Present*
  - **Leadership & Roadmap:** Part of Deccan AI's leadership team, owning the **yearly RL roadmap**. Scaled the AI/RL engineering team from **1 to 10 engineers** and ran delivery teams across several projects, leading **about 20 people** overall.
  - **Stark (Simulated Environments, Trajectories and Agentic RL Kit) [Blog]:** Led the team that built the RL Gym Studio platform and designed its research roadmap for **multi-app RL gyms**. Hands-on, built the **Model Context Protocol (MCP) build framework** that automatically converts **80%+ of API docs into mock connectors**. The team shipped a **human-in-the-loop annotation platform** for increasing task complexity and collecting golden trajectories, plus a kit covering **synthetic data creation**, **automatic task and prompt generation**, a **training wrapper** for agentic RL runs, and a **trajectory analysis pipeline** for grading agent behavior. Delivered **100+ environments averaging 30 tasks each to a top-five AI lab**.
  - **Coding RL Gym:** Built an RL gym of **500+ verifiable software-engineering tasks mined from real private repositories**, spanning dataset families such as **DeepSWE** and **Terminal-Bench**, for training and evaluating **coding agents**. Built the **codebase screening** and **Personally Identifiable Information (PII) redaction** pipelines that qualify repositories for use, and **Manthan**, a **task-creation pipeline** that mines, authors and repairs tasks and automatically creates **model-breaking DeepSWE tasks** at a **70% automated yield**; for the remaining 30%, led the team to build an **IDE-driven, human-in-the-loop fix pipeline**. Frontier models such as GPT-6.1 Sol have a **pass rate of less than 40%** on these tasks.
  - **Caption Bench [Research]:** Led the team that designed the complete research framework for evaluating **LLMs on video understanding across 7 properties**, and published it as Deccan AI's **video-captioning benchmark**. Benchmarked the **top five Vision-Language Models (VLMs)** on it; the best model reaches only **about 57% accuracy**.
  - **Media Studio:** Tech architect and sole backend lead for Media Studio, an OpenRouter-style **LLM gateway** that fronts multiple model providers behind one API with routing, rate limiting and failover, serving **about 2M requests per day** and sustaining **burst loads above 10,000 requests per second**.
  - **Expert Talent Search Platform:** Shipped a hybrid retrieval system over **2M+ profiles** to source domain experts for RLHF projects. **Dense embeddings**, structured filters and an **LLM query-understanding and reranking pipeline** reach **nDCG approx. 0.70** on expert-labeled relevance sets. A **Celery** resume-ingestion pipeline parses **1,000+** resumes per batch into structured signals that feed the index.
  - **AI Interviewer at Scale:** Built a configurable, end-to-end **voice-agent interviewer** that generates role-specific questions and adaptive **follow-ups driven by the candidate's responses**. A **proctoring model** (**voice classification, gaze detection and tracking**) flags suspicious sessions so interviews run unattended at volume, and an **evaluation engine** scores candidates across **multiple dimensions** with a structured summary for the interviewer.
  - **Browser-Use & Computer-Use Gyms:** Built the initial stage of the **browser-use and computer-use gyms**, and advise the teams now extending them along with other RL environment, training and evaluation projects.
  - **Research:** Researching recipes for **synthetic data pipelines for coding benchmarks**, including DeepSWE-style tasks, **formal verification** in coding benchmarks, **computer-use gyms** and **design-to-code (UI/frontend) gyms**.

- **Mensa Brands** — Bengaluru / Abu Dhabi, UAE
  ML Engineer (IC) — *Nov 2022 - Dec 2024*
  - **Pricing Central:** Engineered a dynamic pricing system using Facebook's **Prophet** for demand forecasting and **Mixed Integer Programming (MIP)** for portfolio-wide price optimization, delivering a **5-7% EBITDA uplift**.
  - **Infinity:** Architected a serverless MLOps platform on **AWS CDK** integrating **Lambda, Step Functions, Batch and GitHub Actions**, enabling automated model deployment, training, monitoring and data-ingestion pipelines.
- **Johnson Controls International (JCI)** — Bengaluru, India
  Principal ML Engineer — *Mar 2021 - Aug 2022*
  - **Business Logic Module:** Re-engineered the business-logic layer that processes events from the **RabbitMQ** message broker and computes derived metrics, achieving a **10x faster response time** and a **15% reduction in deployment cost**.
  - **ACVS Engine Revamp:** As a member of the AI Architecture team, redesigned the object-detection module and upgraded it to **YOLOv8**, extended it to multi-class detection and multi-object tracking, and moved tracking to **DeepSORT**.
  - **JCI AI Round Table, India:** Key contributor to the AI Leaders Committee, driving Applied AI strategy and prioritizing solutions that shaped the long-term roadmap in the smart ACVS sector.
- **Mfine (Novocura Pvt Ltd)** — Bengaluru, India
  Senior ML Engineer — *Jan 2019 - Mar 2021*
  - **Enigma ML Platform:** Led design and build of the company's **MLOps infrastructure** from the ground up on **Metaflow, MongoDB, Docker, Kubernetes and DVC**, cutting **model re-training time by 50%** and **AWS costs by 75%**, with one-click production deployment and end-to-end model tracking.
  - **Maya AI Assistant:** Rebuilt the doctor-assistant tool with **Bayesian inference** for symptom matching and **Siamese networks** for differential diagnosis, cutting **physician workload by 30%** and achieving an **80% click-through rate on top-5 predictions**.
  - **Medical Embeddings:** Trained domain-specific **Word2Vec, GloVe and Bio-BERT** models on proprietary clinical data (SNOMED-CT, ICD-10), improving performance across multiple in-house clinical NLP tasks.
  - **Other Contributions:** Built a proprietary Explainable AI library (Grad-CAM, CAM, attention inference); mentored interns on an MFCC-based heartbeat-classification system.
- **SigTuple Pvt Ltd** — Bengaluru, India
  ML Engineer IV — *Dec 2016 - Jan 2019*
  - **Shonit (Blood Analyzer):** Engineered a computer-vision algorithm for White Blood Cell (WBC) extraction using segmentation, adaptive thresholding and morphological operations, improving **accuracy from 67% to 99%**. Trained 200+ deep-learning models for WBC classification to reach **96% F1**, leading to **2 successful clinical studies**.
  - **Drishti (Fundus/OCT Screening):** Re-engineered the Diabetic Retinopathy (DR) screening pipeline with a **CNN using feature-pyramid pooling**, producing a model that significantly outperformed previous benchmarks.
  - **Other Contributions:** Designed the Dipstick automatic-analyzer pipeline using **ORB feature descriptors** and an **SVM classifier**; contributed to the in-house ML platform for segmentation, structured data and a scikit-learn wrapper.

PROJECTS

- **QA Co-pilot LLM Agent [Video]** — Drizz Dev, Bengaluru
  Architected a **multimodal LLM agent** on **LangChain** that executes QA tests from predefined flows using custom **agentic tools**. Powered by **OpenAI GPT-4o**, it validates mobile-app functionality with high accuracy, automating test execution.
- **Personalized AI Video Generation** — Unscript.ai, Bengaluru
  Modularized the **DeepFaceLab** codebase for **5x faster inference** in personalized video generation, enabling near-real-time synthesis, and integrated **Wav2Lip** for precise lip-sync.

EDUCATION

Scaler Neovarsity, Master of Science (MS) in Computer Science | India — June 2024
National Institute of Technology, Surathkal, Bachelor of Technology (B.Tech) in Electrical Engineering | India — May 2016

PUBLICATIONS AND PATENTS

Jaiprasad R., Ahmed M., et al. *"Method and System for Generating a Structure Map for Retinal Images"*. Indian Patent 201841049496 (granted).

Jaiprasad R., Gupta K., et al. (2020). *"Dual Stream Network with Selective Optimization for Skin Disease Recognition in Consumer Grade Images"*. ICPR 2020.

Mundhra, D., et al. (2017). *"Analyzing Microscopic Images of Peripheral Blood Smear using Deep Learning"*. DLMIA Workshop, MICCAI 2017.