

Divyansh Agrawal

Applied AI Engineer | Production AI Systems | Agents & Evals | AI Infrastructure

9819144432 | divagr1925@gmail.com | [LinkedIn](#) | [GitHub](#) | divagr.com

Experience

BizDateUp

Jul. 2026 – Present

AI/ML Technical Lead

Mumbai, India

- **Leading applied AI development**, translating investment, diligence, document intelligence, research, and operations use cases into agent architectures, evals, integrations, and production workflows.
- **Build production LLM systems** spanning retrieval, tool use, structured reasoning, evaluation, and human-in-the-loop control from rapid prototypes through deployment.
- **Work directly with portfolio engineering teams** on APIs, distributed services, data pipelines, cloud integrations, and AI infrastructure.

Bajaj Finserv Health

Dec. 2025 – Jun. 2026

SDE Engineer Intern

Pune, India

- **Independently maintained and upgraded an AI adjudication system for health insurance claims** in a high-ownership engineering team reporting directly to the CTO, processing **50k+ claims and 1 million+ documents/day** at **97% accuracy** with HITL for suspicious/rejected claims.
- **Integrated adjudication into a 22-step production workflow** spanning ingestion, document processing, classification, routing, decisioning, and verification, using triple-check model passes and self-hosted models.
- **Built and productionized a real-time STT-LLM-TTS agent stack** with tool use, PII redaction, tone/prosody control, batching, KV-cache reuse, and Q8 inference; supported 10+ STT/TTS models at **76% lower serving cost**.

Achievements

- **Winner, Cerebras × Google DeepMind Gemma 4 Hackathon:** Built Trace, an autonomous multimodal debugging system taking support reports through reproduction, patching, verification, and handoff in ~12 seconds.
- **Winner, ISB Hackfest:** Built an on-prem voice clinical AI system spanning documentation, ICD-10 treatment coding, and insurance workflows.
- **National Rank 2 / 15,000+ teams, Bajaj HackRx 6.0:** Built a low-latency RAG and browser agent for retrieval, reranking, and web-grounded reasoning.
- **Finalist / 30,000+ teams, Meta × PyTorch × Hugging Face:** Built a GRPO-trained Kubernetes agent operating under cost and latency constraints.
- **Qwen Developer Ambassador, Alibaba Cloud:** Technical demos, open-source contributions, developer education, and model feedback across the Qwen ecosystem.

Projects

FactorioRL | *Reinforcement Learning, LLM Agents, Hierarchical Control, Game Environments*

- **Built Factorio into an agent-native playground and RL research environment** where LLMs, learned policies, and hybrid systems interact with the same persistent world, build factories, acquire reusable skills, and recover from failures.
- **Engineered deterministic game stepping, typed actions, structured partial observations, parallel rollouts, fast resets, and procedural tasks;** trained masked PPO agents with hierarchical skills to study sparse rewards, temporal abstraction, and generalization.

Fabric | [GitHub](#) | [Demo](#) | *WebMCP, Distributed Agent Runtimes, Edge Infrastructure, Human-in-the-Loop*

- **Architected a distributed browser-based compute fabric for LLMs**, enabling ChatGPT to dynamically synthesize WebMCP tools across heterogeneous client devices; executed workloads where data lives to eliminate cloud data transfer while routing human-in-the-loop tasks via native device capabilities.
- **Engineered self-healing tool infrastructure via Cloudflare Durable Objects:** built real-time topology tracking that detects device disconnects mid-execution, replans the task graph across surviving hardware, and hot-swaps WebMCP tool schemas on the fly without breaking agent context.

Trace | [GitHub](#) | *Autonomous Agents, Browser Use, Code Intelligence* | *Cerebras × Google DeepMind Winner*

- **Built a multimodal autonomous debugging agent** that turns vague bug reports into reproduced failures, source-localized patches, verified fixes, and GitHub PRs in ~12 seconds using Gemma 4 31B on Cerebras.
- **Designed the end-to-end agent loop** across browser reproduction, code reasoning, constrained patching, regression verification, and GitHub integration.

Pull Guard | [GitHub](#) | [Demo](#) | *Agent Evals, Adversarial Verification, Code Intelligence*

- **Built an AI-native evaluation and verification system for large PR queues** combining semantic clustering, source-grounded review, patch–claim consistency checks, and adversarial test generation.
- **Designed a self-verifying execution harness** that validates generated tests against clean, candidate, and deliberately broken revisions in isolated disposable runners.

Hyperion | [GitHub](#) | [hyperionhq.co](#) | *LLM Infrastructure, Routing, Reliability*

- **Built a multi-provider LLM gateway handling ~21k requests/sec** at ~5 μ s routing overhead with per-key budgets, semantic caching, model routing, and automatic failover.
- **Designed production reliability, privacy, and cost controls** including cost-aware routing, air-gapped PII redaction at ~35 μ s p99, and ~3 μ s failover.

Memlayer | [GitHub](#) | *Agent Memory, Retrieval, Python* | 250+ stars, 3k+ weekly PyPI downloads

- **Built a model-agnostic long-term memory and retrieval layer for LLM agents** with structured memory, hybrid vector/graph retrieval, and pluggable storage, adopted across thousands of installs.

Evaluation & Research

EchoNet | [GitHub](#) | [Preprint](#) | *Agent Evaluation, Epistemic Robustness, Retrieval*

- **Built a controlled synthetic-web benchmark for epistemic arbitration** across 600 matched counterfactual environments testing poisoning, manufactured consensus, source provenance, and legitimate belief updates.
- **Scored models across eight epistemic-arbitration metrics** measuring resistance to misinformation, legitimate belief updating, source fidelity, provenance sensitivity, calibration, retrieval behavior, consensus dependence, and final decision quality.

Agent Social Simulation | *Multi-Agent Systems, Emergent Behavior, Model Evaluation*

- **Built controlled societies of 200 LLM agents across four tribes** with distinct value systems and currencies but identical starting conditions, isolating how model behavior shapes emergent social outcomes.
- **Ran 100-round simulations across different models** to study trade, cooperation, migration, conflict, value drift, and long-run survival under the same incentives, resource pressures, and inter-group dynamics.

StateShift-JEPA | *Agent Reliability, World Models, Representation Learning*

- **Developed an action-conditioned latent world model** that predicts future state representations to detect stale agent beliefs after hidden environment changes and trigger repair.
- **Built StateShift-Bench and a predictor-repair planner** with hidden-shift TextWorld environments, uncertainty-aware action scoring, and frontier-model baselines.

Technical Skills

AI/ML: LLM Systems, Agents & Tool Use, Retrieval/RAG, World Models & Predictive Latent Models, Representation Learning, Computer Vision, Self-Supervised Learning, Multimodal Learning, Post-Training & RL, Evaluation Harnesses & Graders, Quantisation, Model Routing, Inference Optimisation & Serving, Prompt Caching & Continuous Batching

Backend & APIs: Distributed Services, API Design, Event-Driven Architectures, Background Workers, Data Pipelines, Service Integrations, FastAPI, Django, Node.js/Express, Kafka

Data & Infrastructure: PostgreSQL, pgvector, Redis, SQL, Vector/Graph Retrieval, Caching, Observability, Reliability & Recovery, PII Redaction, Kubernetes, Docker, CI/CD, AWS, GCP, Azure

Languages: Python, TypeScript, JavaScript, SQL, Bash

Education

Mukesh Patel School of Technology Management & Engineering

B.Tech in Computer Engineering with Honors in AI/ML

Mumbai, India

2022 – 2026