

LOKESH MOHANTY

RESEARCH ENGINEER

(+91) 747-785-8933

lokeshm@iisc.ac.in

<https://lokeshmohanty.github.io>

github.com/lokeshmohanty

linkedin.com/in/lokeshmohanty

SUMMARY	Research engineer at the intersection of reinforcement learning and large-scale ML systems. PhD research on a dual-process, reward-gated architecture for real-time test-time learning under non-stationarity, with provable guarantees. Experienced in LLM pre-/post-training (SFT/RLFT) on TPU, distributed synthetic-data generation, and GPU-accelerated computing.	
EDUCATION	Indian Institute of Science (IISc) <i>Ph.D. in the Department of Computational and Data Sciences</i>	Bengaluru, India Aug, 2022 - Jul, 2027 (expected)
	• Advisor: Prof. Sashikumaar Ganesan • Research area: Reinforcement Learning, LLM Post-training • GPA: 8.01/10.00.	
	Indian Institute of Technology Kharagpur (IIT KGP) <i>B.Tech. in Mining Engineering</i>	Kharagpur, India Jul, 2015 - Jun, 2019
	• GPA: 7.10/10.00.	
RESEARCH	A Dual-Process, Reward-Gated Architecture for Test-Time Learning <i>Ph.D. Research, IISc — advised by Prof. Sashikumaar Ganesan</i>	2024 - Present
	• Real-time test-time adaptation of RL policies under non-stationarity: a fast-weight learner adapts continually while a deliberate controller escalates only on environment-reward failures • Reward-gated escalation via a drift-robust change detector (CUSUM) with provable gate optimality and dynamic-regret guarantees • Generator-free instantiation: quality-diversity behavioral repertoire with trust-region policy search and verify-before-commit	
PROJECTS	DSDG: Distributed Synthetic Data Generation <i>ZenteiQ AiTech Innovations — Collaboration, IndiaAI Mission</i>	Mar, 2026 - Present
	• Pre-trained LLMs on TPU using MaxText (JAX) • Built a scalable, distributed data-generation framework for SFT/RLFT over FastAPI, Kafka and YugabyteDB with Prometheus/Grafana observability • Ran supervised (SFT) and RL (RLFT) fine-tuning; served inference with vLLM	
	CAESAR: Combat Aircraft Engagement and Strategic AI Response <i>Defence R&D Organisation (DRDO) — Collaboration</i>	Mar, 2024 - Mar, 2026
	• ML framework for end-to-end emitter detection, identification and multi-sensor data-fusion for enemy recognition	
	SimInhale: Particle Deposition in Human Lung Airways <i>Indian Institute of Science (IISc)</i>	Sep, 2022 - Mar, 2023
	• GPU-accelerated particle deposition; flow simulation with ParMoon (PDE solver), parallelized with OpenMP and CUDA	
EXPERIENCE	ARC Document Solutions Kolkata, India	Jul, 2019 - Jul, 2022
	• Backend microservices in Node.js/TypeScript for a facility-management platform, with S3 storage and OCR-based deep search • Common IAM using OAuth2 and SAML with Redis for caching • CI/CD on AWS using Jenkins and Docker	
SKILLS	Programming: Python, C++, TypeScript/Node.js, Haskell	
	ML & Research: PyTorch, JAX, Reinforcement Learning, LLM Post-training, vLLM	
	Systems & Tools: Docker, Kubernetes, Kafka, YugabyteDB, AWS/GCP, ClearML	
	Open-Source Software: Neovim, Void Linux, LoopBack4	