Prakul Sunil Hiremath

Machine Learning | Large Language Models | Reinforcement Learning

prakulhiremath@vtu.ac.in | GitHub | Google Scholar | LinkedIn | Website

Belagavi, Karnataka, India | +91 93800 72238

Research Summary

B.Tech. Computer Science researcher focused on language-model training and evaluation, reinforcement learning, and reliable machine-learning systems, with additional work in quantum information and communication. Conducts empirical research with GPT-2 and LLaMA-3-8B, including PPO-based RLHF experiments, inference-latency instrumentation with vLLM, controlled evaluations, and calibration analysis. Two papers accepted as posters at NeurIPS 2026; builds tested, open-source research software in Python and C++.

Education

Visvesvaraya Technological University (VTU) — Belagavi, Karnataka, India
B.Tech. in Computer Science and Engineering | CGPA: 8.02/10.00 — *Aug 2023 – Aug 2027 (expected)*
- **Relevant coursework:** Linear Algebra, Multivariable Calculus, Probability and Statistics, Design and Analysis of Algorithms, Artificial Intelligence and Machine Learning, Operating Systems.

Prerna Pre-University College, KSEAB — Belagavi, Karnataka, India
Pre-University, Science (Physics, Chemistry, Mathematics) | 91.00% — *Jun 2021 – Jul 2023*

Billimoria High School, CBSE — Panchgani, Maharashtra, India
Secondary School | 95.6% — *Mar 2011 – Apr 2021*

Technical Skills

Programming	Python, C++, C, Rust, SQL
ML / Deep Learning	PyTorch, Hugging Face Transformers, JAX, scikit-learn, NumPy, SciPy, Pandas, Polars
LLMs	Transformer-based language models, GPT-2/3/4/4.5, LLaMA-3-8B, RLHF experimentation, model evaluation, calibration analysis
RL & Deep RL	PPO, Stable-Baselines3, Gymnasium, rollout analysis, policy-gradient evaluation
LLM Systems / Compute	vLLM, inference-latency instrumentation, Ray, CUDA, GPU-based experiments, OpenTelemetry
Development	Git, GitHub, Docker, Linux, Jupyter, Slurm, Pytest, GitHub Actions, Weights & Biases
Methods	Controlled experiments, ablation studies, statistical analysis, linear algebra, probability, Markov processes, POMDPs
Exploring	LoRA/PEFT, knowledge distillation, agentic frameworks (LangChain, CrewAI), tool use and function calling, Triton, NCCL, Kubernetes, TensorFlow

Selected AI Research

Latency-Conditioned Selection Bias in RLHF — **NeurIPS 2026 (accepted)**
Sole author — Poster page
- Identified a failure mode in asynchronous RLHF in which discarding slow rollouts creates an implicit sequence-length filter: longer reasoning trajectories are under-sampled because generation time grows with token count, biasing policy-gradient updates.
- Instrumented a vLLM server running LLaMA-3-8B and quantified the latency–length relationship over 12,000 rollouts ($R^2 = 0.918$). Evaluated PPO under rollout truncation on GPT-2 (10 seeds) and LLaMA-3-8B (3 seeds); on LLaMA-3-8B, accuracy fell from 68.4% to 44.1%.
- Validated against a throughput-matched synchronous baseline and a synthetic latency-injection control. Proposed Effective Gradient Information (EGI) for early degradation detection and Selection-Probability Weighting (SPW), which recovers 97.8% of on-policy accuracy on LLaMA-3-8B with under 0.3% overhead.

When Does Sequential Detection Collapse to a Scalar? — **NeurIPS 2026 (accepted)**
First author — Poster page
- Derived a necessary-and-sufficient condition for exact scalar reduction in Bayesian sequential detection and characterised the lead-time loss when the condition fails.
- Validated on 72 synthetic configurations (lead-time prediction within 0.4 steps MAE) and on CICIDS2017, where the method achieved 14.3 steps of anticipation versus at most 1.5 for scalar deep-detector baselines.

Calibration Drift Under Reasoning — 2026
Preprint — arXiv:2606.11211
- Analysed how confidence and calibration shift under reasoning and chain-of-thought budgets, evaluating Llama-3.1-8B and Llama-3.3-70B across 21 benchmarks.

Additional submissions: *Evaluation Collapse* (ICLR 2027, submitted); *Information Is Not Enough* (AISTATS 2027, submitted).

Research Experience

Research Intern – Capability Dynamics Sep 2026 – Present

ITISMU, VTU Department of Computer Science and Engineering *Belagavi, India*

- Designing controlled PyTorch experiments to test whether the initial spectral structure of a one-block Transformer predicts the onset of generalisation on modular addition (mod 97); pilot runs span 30 random seeds and 20,000 training steps.
- Audited the evaluation protocol and uncovered commutativity leakage: at the first checkpoint reaching 95% training accuracy, test accuracy was 91.3% for examples whose reversals appeared in training versus 0.32% otherwise. Superseded the original crossing-time estimates accordingly.
- Evaluated an early-loss predictor and a permutation control, documented the negative results, and designed a leakage-free pair-orbit split with a pre-specified censored time-to-event analysis; corrected experiments are in progress.

Researcher – AstroSat Data Analysis Sep 2024 – Present

Cosmic Calculus and Astronomy Club, VTU *Belagavi, India*

- Built an unsupervised anomaly-detection pipeline for AstroSat SXT light curves, using synthetic transient injections into real background data to estimate detection probability and false-trigger rates; presented at the AstroSat conference, U R Rao Satellite Centre, Bengaluru (Jan–Feb 2026).

Founder and Research Lead Feb 2024 – Present

Aliens on Earth (AoE) Research Collective *Belagavi, India*

- Lead a student research collective working on machine learning, latent dynamical systems, and epistemic uncertainty. Group projects have secured more than INR 20 lakh in competitive funding through MSME Idea Hackathon 4.0, Startup Karnataka NAIN 2.0, and FKCCI Inno Manthan.

Selected Research Software

Latency Gym Jun 2026

C++20 reinforcement-learning environment GitHub | Zenodo

- Implemented a Pybind11-bound environment for latency optimisation with a reward combining mean latency, variance, and dropped-request penalties; trained PPO with Stable-Baselines3 against a random baseline. Includes 50+ deterministic tests.

ChronosDefrag May 2026

Contrastive regime discovery for time series GitHub | PyPI

- Built a PyTorch framework for time-series regime discovery using Temporal Entropy Regularisation and Similarity Persistence Decay; CPU inference benchmarked at approximately 1.14 ms mean and 2.31 ms p99, with 19 deterministic tests.

Temporal Leaks May 2026

Look-ahead bias detection for time-series ML GitHub | PyPI

- Developed a typed, CI-tested, published package that perturbs future observations and flags features whose historical values change; supports Pandas and Polars.

Research Service

- **Reviewer:** NeurIPS 2026 Workshop on Neural Representations (NeurReps).
- **Invited presentations:** OpenSearchCon Japan 2025 (Tokyo); CommunityOverCode Asia 2025 (Beijing).
- **Sponsored delegate:** Linux Foundation Open Source Summit 2025 and 2026.

Volunteering

Research and Development Lead Nov 2026 – Present

GDG on Campus, Belgaum *Science and Technology*

- Lead research and development initiatives covering emerging technologies, AI/ML, and applied computing; guide students in turning research ideas into working prototypes and foster a culture of peer learning.

Career Mentor – Career Buddy Jun 2026 – Present

Wadhwani Foundation *Education*

- Mentor students and early-career learners one-on-one on career planning, higher education, and skills development; completed 10+ sessions and received a Certificate of Recognition.

Languages & Interests

- **Languages:** Kannada, Marathi, English, Hindi; working understanding of Gujarati and Telugu.
- **Interests:** Quantum computing, football, chess; 4× debate competition winner.