

Yuvraj Singh

 [Github](#)  [Website](#)  [Linkedin](#)  yuvraj.mist@gmail.com  [HuggingFace](#)  [X](#)

EXPERIENCE

alphaXiv | *Research intern* Oct. 2025 – March 2026

- Co-authored an [article on further](#) ablating the Evolutionary Strategies for fine-tuning LLMs paper .
- Reimplemented and reproduced results from seminal and recent ML papers (e.g., Attention Is Not All You Need, TRM/HRM) from scratch in PyTorch under the [alphaXiv paper-implementations](#) repository.
- Built evaluation pipelines and benchmarking infrastructure for LLM/VLM systems; evaluated models such as DeepSeek-OCR and OlmOCR2 on OmniDocBench and related OCR benchmarks.
- Deployed large models (e.g., DeepSeek-OCR, OCR2, LLMs like Ouro) using vLLM on Modal/Baseten, and generated a 100k-document OCR dataset from arXiv PDFs for large-scale document understanding research.

University of Maryland | *Research intern* Dec 2024 – Feb 2025

- Worked on UI/UX2Code generation with primary focus on creation of a robust dataset for the VLM models to be fine-tuned upon.
- Scraped 100+ websites of static web pages and collected their commit history and corresponding webpage layouts to create a dataset of 200+ records.

IISER, Kolkata | *Summer Research intern* May 2024 – May 2025

- Co-authored (first author) with Prof. Kripabandhu Ghosh, for creating a dataset of 40k scraped YouTube comments, humanely-verified, stance-analysed famous sports controversies (cricket and football) with applied stance detection.
- Utilized LLMs like Llama-3.1/3.2, Mistral-7b, Qwen-2.5 under zero/few-shot prompt to create the dataset. Benchmarked and fine-tuned existing open-sourced LLMs Reasoning (distilled) and Non-Reasoning LLMs on the humanely verified dataset. Submitted to COLM 2025.

EDUCATION

International Institute of Information Technology Bhubaneswar 2023- 2027
BTech Computer Science Engineering

PROJECTS

NeatRL | *RL Algorithms* 2024

- **Description:** A 200+ starred Github repository of one-page, clean implementations of RL algorithms from scratch for everyone to use in PyTorch. Tested on ALE, OpenAI Gym, PettingZoo.
- Includes algorithms from value-based methods like SARSA, DQN, Q learning, to policy-based based like PPO, SAC, A2C etc

SmolHub | *Website showcase of my deployed models* 2025

- **Description:** A website/playground to showcase model cards for all of my RL From Scratch and Paper Replication repositories, mainly for using them on low-compute devices.
- Includes a 400+ starred GitHub repository section of my 30+ from-scratch reimplementation attempts of seminal/SOTA papers from NLP to Vision and Audio, in Pytorch.
- RL from Scratch section, NeatRL, for every RL algorithm I have re-implemented with visualisations of trained agents on them.

Smolcluster | *An educational library for distributed training/inference of neural networks* 2026

- **Description:** Built smolcluster, a distributed deep learning framework enabling training and inference across heterogeneous hardware (Raspberry Pis, NVIDIA GPUs, Apple machines) using PyTorch and socket-based communication.
- Implemented algorithms including FSDP, Classic DP, Elastic Distributed Parallelism, Synchronous Parameter Server, MP and EP.
- Developed distributed inference infrastructure with model-parallel LLM serving, streaming token generation, FastAPI backend, and React-based chat interface, integrated with Grafana/Loki monitoring .

ACHIEVEMENTS

- 2nd place @GeoHack '24 Finale** 2024
Secured 2nd position at finals of GeoHack 2024 with our project- FarmGenie organised by IEEE GRSS Kolkata and SAADRI.
- Finalist @YESIST12 '24- Special Track** 2024
Led the winning project PlogPayouts at YESIST12 '24 under Special Track category
- Geek-o-thon (D3 @IIIT-BH)** 2023
Led the winning project MovieMania at D3 2023, inter-college hackathon in AI/ML domain