

# RAMJI PURWAR

☎ +91 7275681030

✉ ramjipurwar48@gmail.com

🌐 [Portfolio](#)

📄 [in/ramji-purwar](#)

🔗 [Ramji-Purwar](#)

## EDUCATION

**Indian Institute of Technology, Gandhinagar**

*B.Tech in Computer Science and Engineering | CPI: 8.57/10*

Aug. 2024 – Present

*Gandhinagar, Gujarat*

**Coursework:** Data Structures & Algorithms, Theory of Computation, Software Tools and Techniques for AI, Introduction to Data Science, Machine Learning, Discrete Mathematics

## ACHIEVEMENTS

- Achieved a maximum Codeforces rating of **1603** max (**Expert**), placing in the top 4% of rated users globally. [Profile](#)
- Ranked in the **top 15 teams** nationally at the **upGrad School of Technology** hackathon, selected to present the solution at the Mumbai finals. [Deployed-Link](#)
- Ranked **1st** in the Game Design problem statement at **HackRush 2025**, IIT Gandhinagar's annual technical hackathon.

## OPEN SOURCE WORK EXPERIENCE

**libpysal | 2 Merged PRs | [PRs Link](#)**

- Implemented `Graph.from_networkx()` in libpysal, enabling direct conversion of NetworkX graphs into PySAL Graph objects using an optimized sparse matrix representation.
- Benchmarked against `textttmx.to_pandas_adjacency()` and demonstrated a 2–3× speedup on moderately large graphs.

## PROJECTS

**Attention is not Explanation | [GitHub](#)**

Jan'26 – Apr'26

*Mentor: Prof. Manish Padala — IIT Gandhinagar*

- Reproduced Jain & Wallace (2019) on BiLSTM (86.2% accuracy) and extended the study to DistilBERT (90.8% accuracy) on IMDB sentiment classification, testing whether attention weights faithfully explain model predictions.
- Designed adversarial attention attacks (gradient-based optimization, entropy, and random sampling), finding adversarial attention with average cosine similarity  $\sim 0.4$  yielding the same predicted class in 94.9–95.3% of 25,000 test samples.
- Measured rank correlation (Kendall's  $\tau = 0.15$  BiLSTM, 0.40 DistilBERT) between attention and gradient-based feature importance, falsifying attention as a faithful explanation.
- Developed Attention Rollout and gradient-weighted Layer-wise Relevance Propagation (LRP) from scratch as alternative, multi-layer attribution baselines for the Transformer.

**Scalable Graph Retrieval Pipeline | [GitHub](#) | [Demo](#) | [Report](#)**

Feb'26 – Apr'26

*Mentor: Prof. Anirban Dasgupta — IIT Gandhinagar*

- Trained a GIN-based graph embedding model using triplet loss with a beam-search (approximate) GED oracle, improving retrieval precision over untrained embeddings by up to 14% (e.g., 0.647  $\rightarrow$  0.790 on PROTEINS).
- Built an LSH-based approximate nearest neighbor retrieval pipeline that matched brute-force precision within 1–2% on most datasets while searching only 40–60% of the corpus.
- Identified an LSH failure mode on IMDB-Binary, where excessive hash selectivity caused candidate starvation (2.9% of corpus searched), dropping approximation quality to 0.44 versus  $>0.85$  on other datasets.

**Ctrl-Agent | [GitHub](#)**

May'26 – Present

- Built a two-pass email triage pipeline monitoring Gmail on a scheduled interval, fetching up to 30 unread messages per cycle and classifying each as action required / info / ignore via a deterministic rules.
- Architected a native MCP (Model Context Protocol) server enabling external LLM clients like Claude Desktop to autonomously invoke local filesystem, Git, and Gmail tools over stdio, with live UI syncing.
- Designed a multi-service architecture combining Flask, APScheduler, and system tray integration for persistent background operation, with single-instance handling and scoped `BASE_DIR` file access for secure, reliable execution.

**Chess Engine | [GitHub](#)**

Oct'25 – Jun'26

- Built a chess engine in C++ that plays at  $\sim 1700$  Elo (benchmarked against Chess.com engine), using negamax search with alpha-beta pruning to look ahead and pick the best moves efficiently.
- Designed an evaluation function that adapts across opening, middle, and end game phases, using move ordering, position hashing, quiescence search, and null move pruning to maximize search speed and accuracy.

## TECHNICAL SKILLS

**Languages:** C++, Python, C, HTML, CSS, JavaScript

**Libraries & Frameworks:** NumPy, pandas, scikit-learn, NetworkX, GeoPandas, Matplotlib, Seaborn, Flask, Streamlit, Pytorch

**Developer Tools:** Git, GitHub, Docker, Linux, Hugging Face, Vercel, Jupyter Notebook

## POSITIONS OF RESPONSIBILITY

**Core Member — GRASP, CP Group, IIT Gandhinagar**

Aug'25 – Apr'26

- Mentored peers on algorithmic problem solving, competitive programming techniques, and contest strategies.

**Webmaster — Student Academic Council, IIT Gandhinagar**

Aug'25 – Apr'26

- Designed and maintained the official Student Academic Council website serving the IIT Gandhinagar student community.