

Kabir Murjani

+91-98255-26632 | kabirmurjani@gmail.com | kabirmurjani.com | [GitHub](#) | [Google Scholar](#)

EDUCATION

Nirma University

Bachelor of Technology, Electrical Engineering (Minor in Management)

Ahmedabad, India

July 2023 – May 2027

- **Relevant Coursework:** Control System Theory, Probability & Statistics, Electrical Design, Network Analysis, Power Systems

RESEARCH EXPERIENCE

Indian Institute of Management Bangalore

Research Intern — Prof. Abhay Sobhanan

Bangalore, India

May 2026 – July 2026

- Developed neural combinatorial optimization solvers for NP-hard knapsack, routing, and synchronization problems. Derived exact solvers over massive combinatorial spaces and introduced a novel NP-hard problem with large-scale benchmark instances.
- Designed a neural hybrid solver combining a deep RL policy for solution construction with metaheuristic refinement, narrowing the solution-quality gap to a strong metaheuristic baseline at a fraction of its computational cost.
- Formalized the Traveling Thief Problem with Drone (TTP-D) using high dimensional state/action encodings, enabling zero-shot transfer, instance-size generalization, and faster solutions than exact solvers across varying problem scales.

Singapore Institute of Technology NVIDIA AI Centre

Research Collaborator — Dr. Timothy Liu

Remote

Jan 2026 – April 2026

- Designed a sequential RL framework for multi-objective decision-making over hierarchical knowledge graphs, reformulating conflicting objectives to produce consistent and deterministic decisions across high-throughput textual data.
- Developed anomaly detection methods for continuous audio-visual streams, combining spectral and semantic representations to detect abnormal patterns in real-time data. Deployed models through NVIDIA Inference Microservices (NIM).
- Benchmarked transformer representations across 12 layers and 1,696 segments, identifying the optimal layer with 0.736 ROC-AUC and a $\sim 10\%$ relative AUC gain over the final layer under identical evaluation conditions.

Xtin Capital, Inc.

Machine Learning Researcher

Ahmedabad, India

May 2025 – Aug 2025

- Built a high-frequency ingestion engine ingesting 800+ articles/min using sentiment models and hierarchical knowledge graphs to capture semantic relationships, converting unstructured real-time news into structured data for downstream decision making.
- Developed a backtesting environment with adversarial data generation to quantify qualitative data and evaluate portfolio rebalancing strategies across volatility regimes, enabling systematic stress testing under high-fidelity simulated conditions.
- Integrated sentiment signals to rank and cluster high-impact news, quantifying their marginal contribution to portfolio rebalancing performance against a price-only baseline across multiple volatility regimes.

PUBLICATIONS & PROCEEDINGS

Peer-Reviewed

- **Murjani, K.**, Bhavsar, M., Patel, M., & Talukdar, J. (2026). AlphaRoute: LLMs as Semantic Optimizers for Multi-Objective Routing. In *IEEE International Conference on LLM-Aided Design (LAD)*, Stanford University, CA* [\[arXiv\]](#) [\[Code\]](#)
- **Murjani, K.** (2026). Zero-Copy Semantic Contagion: An In-Memory Streaming Architecture for Evolving Attention Graphs. In *ACM SIGMOD Workshop on Data Management for the Modern Financial Systems, Bengaluru, India*[†] [\[arXiv\]](#) [\[Code\]](#)
- **Murjani, K.**, & Vyas, P. (2026). LURE: Bayesian Signaling Game in Multi-Step Agent Interactions. In *Conference on the Mathematics of Artificial Intelligence, Sochi, Russia** [\[PDF\]](#) [\[Code\]](#)

Under Review

- **Murjani, K.**, & Sobhanan, A. (2026). Drive, Pack, Fly: The Travelling Thief Problem with Drone. *Under review, European Journal of Operational Research (EJOR)* [\[arXiv\]](#) [\[Code\]](#)
- **Murjani, K.**, & Sobhanan, A. (2026). Neural Combinatorial Search Needs Rollouts. *Under review, Eighteenth International OPT Workshop on Optimization for Machine Learning, NeurIPS 2026*
- **Murjani, K.**, & Patel, N. (2026). The Combinatorics of Trust over Financial Knowledge Graphs[‡]. *Under review, Eleventh Workshop on Financial Technology and Natural Language Processing (FinNLP 2026), EMNLP 2026*
- **Murjani, K.**, Vyas, P., Abrol, A., Gupta, R., & Wang, Z. (2026). The Variance of a Judge: Zero-Variance Rewards by Deterministic Regression. *Under review, ACL Rolling Review (ARR)*

PROJECTS

Selective Coverage over Temporal Graphs [Code]

- Recast long-context QA as budgeted evidence selection over temporal fact graphs of QMSum transcripts (succession, contradiction, relatedness edges). Trained a heterogeneous graph-transformer policy with PPO/GAE under verifiable rewards.
- Improved token-F1 from 25.0% to 27.7% while using $28\times$ fewer input tokens. Proved the coverage objective submodular and NP-hard via Maximum Coverage, giving a $1 - 1/e$ greedy guarantee. [‡]Shared methodology, with Nisarg Patel (Google).

Spectral Dominance in Sequential Drug Control [Code]

- Modeled drug scheduling under resistance evolution as a zero-sum game on a within-host compartmental graph, deriving a spectral-radius condition for when sequential treatment dominates myopic dosing (with Prof. Connor Jerzak, UT Austin).
- Proved single-stage NP-completeness by reduction from 0-1 Knapsack and bounded the myopic-optimal gap across 3 regimes. Validated on 4 calibrated modules (N up to 30 nodes) via a bilevel MILP and a spectral-greedy heuristic.

Hierarchical Reinforcement Learning for Cost-Emission Optimal Dispatch [Code]

- Decomposed 24-hour cost-emission dispatch on IEEE 118-bus (54 generators, 4,242 MW) into a metaheuristic macro layer for a 15-point trade-off front and a decentralized RL micro layer (Undergraduate project, advised by Prof. Kuntal Bhattacharjee).
- Reduced per-zone action space from 5^{21} to 7 discrete dispatch levels, making tabular Q-learning tractable while maintaining $< 0.05\%$ daily fuel-cost variation. Benchmarked against value-based, policy-gradient, and neural architectures.

Aniate: Terminal Command Synthesis via Graph Search [Code] [Website]

- Built Ant, a terminal system mapping natural-language intent to executed commands by searching a graph of candidate command paths, pruned by iterative guardrails and dynamic model routing.
- Designed its context layer, Engram, as bounded multi-hop traversal over a typed relation graph, ordering candidate nodes before expansion. Rust-accelerated and compiled for sub-50 ms retrieval, released as the open-source python library **aniate**.

ChessNano: Autoregressive Search over the Legal-Move Graph [Code]

- Trained a sub-50M-parameter autoregressive transformer to predict moves directly from Standard Algebraic Notation, with no explicit board representation, so a game is learned as a walk over the position graph induced by the rules.
- Used quantized training (PolarQuant) and tree-of-thought constrained decoding, expanding only edges present in the legal-move graph at each ply so sampling stays on legal continuations. Awarded OpenAI compute grant through OpenAI Parameter Golf.

SKILLS

Languages: Python, C/C++, Rust, MATLAB

ML & Frameworks: PyTorch, PyTorch Geometric, TensorFlow, Hugging Face (SFT, PEFT), LangGraph, NetworkX, RustworkX

Optimization & Solvers: Gurobi, CPLEX, HiGHS, OR-Tools, Pyomo, CVXPY, Z3, Lean/Mathlib

Tools & Infrastructure: L^AT_EX, Git, Docker, Linux, CUDA, Weights & Biases, Ray, Redis, Weaviate

Research Interests: Deep Reinforcement Learning, Neural Combinatorial Optimization, Game Theory, Language Models

HONORS & AWARDS

ACM SIGMOD/PODS 2026 India Fellowship | *Recipient* | 2026

- Selected from a competitive nationwide applicant pool for excellence in advanced development and open-source research.

IIT Kanpur Chess Masters Premier League 2.0 | *3rd Place* | 2023 – Present

- Placed **3rd** among 60+ universities representing the Nirma University Chess Team. Head of Chess representative for EE.

LEADERSHIP & SERVICE

IEEE Computer Society, Nirma University | *General Secretary* | Dec 2025 – Feb 2026

- Founded the chapter and served as its first General Secretary, building the organizing team and launching an NLP workshop series and recurring research seminars to build campus research culture.

Association for Computational Linguistics (ACL) 2026, San Diego, CA | *Volunteer* | July 2026

- Session Chair, Main and Findings Track, 64th Annual Meeting of the Association for Computational Linguistics (ACL), alongside Dr. Jingbo Shang and Dr. Mengdi Huai, with over 10,000 attendees.

Team Dyaus, Nirma University | *Technical Mentor* | Jan – May 2024

- Mentored the avionics subsystem for a CanSat entry in the Türksat Satellite Competition 2024, advising PCB architecture for the flight electronics stack across schematic and layout revisions.

[†] Oral presentation ^{*} Accepted (Oral)