

RAHUL NANDAKUMAR

4600 Monterey Oaks Blvd, Austin, Texas, US

✉ rahul.nandakumar@utexas.edu | 🏠 Webpage | 🌐 GitHub | 🔗 LinkedIn | 📞 +1 (737) 707-8242

Education

- The University of Texas at Austin** Aug. 2024 - Present
PhD in Information, Risk and Operations Management
- The University of Texas at Austin** Aug. 2022 - Aug. 2024
Master of Science in Operations Research and Industrial Engineering (Thesis)
- National Institute of Technology, Andhra Pradesh, India** Aug. 2018 - May 2022
Bachelor of Technology in Chemical Engineering (Thesis)

Professional Experience

- Moloco** May. 2024 - Aug. 2024
Data Science Intern Redwood City, CA
- Increased win rates by 5% and actions by 30%, improving campaign performance for AMR consumer advertisers by adjusting discount rates dynamically for high-value users. Boosted CPA predictability by 20% through statistical research on creative diversity, funnel category, and budget mode, leading to actionable A/B test recommendations.
 - Achieved a 3-4% increase in ROAS for Activision by correlating pLTV with purchase events, leading to better revenue performance in D3 pLTV campaigns.
 - Developed DMA-based strategies for Bumble that reduced CPA by 15%, identifying optimal spend levels and high-performing DMA markets, enhancing profitability for UA campaigns.

- Sabre Corporation** Jan. 2024 - May. 2024
Data Science Intern Austin, TX
- Lead the development of a robust LLM-based solution, leveraging LLaMA 2 to achieve a unified address structures, by processing 1.35 million records of lodging properties from diverse aggregators. Achieved an initial accuracy score of 82.7%
 - Utilized a BERT-based model to generate contextual embeddings from categorical data in over 6 million shop requests to Sabre IntelliSell, improving accuracy and performance for cache rate prediction.
 - Utilized generated contextual embeddings for dynamic price prediction as a downstream task.
 - Leveraged the HPC capabilities of the TACC-Lonestar 6 supercomputer for parallel computing and optimization.

- Twimbit** Apr. 2022 - Jul. 2022
Data Science Intern: AI & Data Team Remote
- Established and managed an extensive data pipeline using Algolia, Heap, Matomo, and Segment to monitor website user interactions. Leveraged A/B testing insights to increase daily user engagement by approximately 5%.
 - Extracted raw HTML data from over 700 web pages using BeautifulSoup and trained a Decision Tree model to automatically classify records into specific categories (such as instabits, podcasts, articles).
 - Tracked read time and page depth scrolled to assess optimal content length, leading to a 20% increase in user retention.
 - Applied text processing and topic modeling techniques with gensim and spaCy-transformers to refine the record structure in Algolia, which led to a 64% improvement in search query response time and enhanced search recommendations.

- AMTDC, Indian Institute of Technology - Madras** Aug. 2021 - Feb. 2022
Research Intern Remote
- Led the development of an autonomous rover, using a 3D camera system for real-time semantic segmentation.
 - Utilized transfer learning with pre-trained models from TensorFlow 2 Model Zoo (SSD MobileNet, Faster R-CNN, Mask R-CNN) to improve the accuracy of object detection and segmentation tasks by 12%.
 - Collaborated with other teams to create a digital twin of the robot using Unity3D, ROS, and Gazebo Simulator, allowing for comprehensive testing of autonomous navigation strategies.

Projects

Reinforcement Learning for multi-node pricing and inventory management.

- Implemented and compared the performance of RL methods (A2C, DQN, PPO) against conventional Mixed IP optimized pricing and (S, s) re-order policies for single/multi-echelon environments with stochastic demand.
- Achieved an impressive 25 % increase in profit with dynamic pricing, and a 1.09 times increase in profit compared to traditional (S, s) order policy.

A graph-based big data framework using Hidden Markov Model and Constraint Satisfaction Problem.

- Designed a novel Bayesian Machine Learning framework to integrate a Constraint Satisfaction Problem into a Hidden Markov model for optimizing state space in financial time series analysis. Utilized Twitter sentiment analysis (Jan 20, 2017 - Jan 20, 2020) to introduce constraints into the model, differentiating from conventional non-CSP methods.
- Successfully applied this approach to predict DJIA, SNP500, and NASDAQ closing prices, achieving a MAPE reduction of 0.59%, elevating accuracy to 90%, and cutting computational time by 0.02s.

Pre-vizArt: Enabling Personalized Artistic Exploration.

- Developed a prompt-based generative artwork preview application using Hugging Face models and Flask, allowing users to choose artwork style. Deployed using Docker, ECS, and K8s on AWS for scalability and performance.

Skills

Languages: Python (Scikit-learn, Keras, Pandas, TensorFlow, PyTorch, Plotly, Dash), MySQL, R, MATLAB, JAVA, C++.

Technologies: Machine Learning, Computer Vision, NLP, Deep Learning, Generative Modeling, LLMs, Optimization.

Frameworks: AWS, GCP, Azure ML, Flask, Django, Docker, Kubernetes, Google Analytics, Tableau, PowerBI, Linux, Git.