India
+91 9176586545
kaushi@alumni.ncsu.edu
kaushikrajan.me

Kaushik Rajan

Summary

I am a scientist with nearly a decade of experience in the tech industry, including four years at Amazon where I developed production-level statistical and time series forecasting models that impacted billions of dollars in advertising revenue. I am seeking to join a productive research group that focuses on applying (deep) reinforcement learning (DRL) techniques to build multi-agent systems. My research interests center around applying game-theoretic principles, such as mechanism design, and neuroeconomic principles, such as neural and computational aspects of decision making, to create more efficient and adaptive agent interactions.

Education

Aug 2015 – May 2017 **Masters in Computer Science**, *North Carolina State University*, Raleigh, NC, USA

Aug 2009 – May 2013 **Bachelors in Computer Science and Engineering**, *VIT University*, India

Selected Publications

Jan 2025 *SSRN pre-print* - Multi-Agent AI: From Isolated Agents to Cooperative Ecosystems. This is also going to be published in the Theoretical Computer Science and Games and Political behavior eJournal by Mar 2025

Jan 2025 *Medium (blog) publication* - Advancing AI Reasoning: Meta-CoT and System 2 Thinking. This was article was accepted into the most popular data science publication in Medium (TDS). It has 800k readers.

Jan 2025 *Medium (blog) publication* - Multi-Agent AI: From Isolated Agents to Cooperative Ecosystems. This was article was accepted into a publication that has 75k readers

Dec 2024 *Medium (blog) publication* - Neural Fictitious Self-Play (NFSP) for Imperfect-Information Games. This was article was accepted into a publication that has 16k readers.

Dec 2024 *Medium (blog) publication* - COCONUT: Redefining Reasoning in Large Language Models. This was article was accepted into a publication that has 33k readers.

Dec 2024 *Medium (blog) publication* - Learning by Doing: An Introduction to Reinforcement Learning. This was article was accepted into a publication that has 16k readers.

Sep 2024 *Econometric Modeling: Capital Markets - Forecasting eJournal Vol 13, Issue 47* Integrating Econometrics and Deep Learning: An Explainable Approach to Portfolio Prediction.

Aug 2024 *SSRN pre-print* - Temporal Diversity in Music Recommendations: A Budget-Aware Approach to Enhance User Engagement and Content Discovery.

Patent

2024 [Provisional] System for Financial Portfolio Management using Batch Bayesian Optimization

Code artifacts (GitHub gists)

Note: These are the code artifacts associated with the publications mentioned above (experiment design, simulations, and results).

1. *End-end pipeline for simulating multi-agent RL (MARL) using the PettingZoo interface.*
2. *Simulation of mechanism design applied to chain-of-agents (CoA).*
3. *Simplified example of training a reinforcement learning agent.*
4. *Simulation of mechanism design applied to the public goods economic scenario.*
5. *Simulation of chain-of-agents applied to book recommendations*

Professional Experience

May 2024 – Present **Applied Scientist / Researcher**, *Self-employed*, India
- Building a solo business for automated content creation (text and video) through deep reinforcement learning agents, and software development principles.
- As a proof of concept, I launched an automated YouTube channel on artificial intelligence and game theory, growing it to 15,000 subscribers and 200,000 views in three months using organic strategies and targeted Google Ads.
- Authored 19 technical articles on Medium (technical blog) exploring deep reinforcement learning, game theory, and neuroeconomics, accumulating 25,000 reads in three months.
- Filed a provisional patent at the intersection of applied econometrics and Bayesian reasoning.
- Providing specialized expertise in auction design and game theory to the US government for a confidential investigation.

Sep 2020 – May 2024 **Data Scientist**, *Amazon*, Arlington, VA, USA
- Utilized statistical, econometrics (time-series forecasting) and deep learning techniques to forecast billions of dollars in ad revenue for Amazon's stores and digital teams. Tuned the models to achieve a worldwide mean absolute percentage error (MAPE) of 1%.
- Led multiple CFO-level goals and initiatives, serving as the science expert to senior economists, finance directors, and the VP.
- Published multiple internal scientific documents detailing the research that I did and the econometric models I developed. Focus areas: Time-series forecasting, deep learning, and observational causal inference.

Sep 2019 – Mar 2020 **Data Engineer**, *EAB*, Richmond, VA, USA
- Developed REST API-based real-time data pipelines to ingest millions of records daily, enabling research and data science in the education technology domain.

Jul 2017 – Sep 2019 **Machine Learning Engineer**, *E3 Retail*, Raleigh, NC, USA

- Developed a multi-layered LSTM network for time-series forecasting, comprising 3 LSTM layers (128, 64, 32 units) and 2 dense layers, reducing Mean Absolute Percentage Error (MAPE) from 18.5% to 10.85% on a 30-day forecast horizon.
- Developed a 4-block CNN (64, 128, 256, 256 filters) with 2 fully connected layers for item detection and inventory forecasting.
- Integrated a spatial attention module in the CNN, boosting mean Average Precision (mAP) from 72% to 85%, and incorporated item-level embeddings, further reducing 30-day MAPE from 10% to 9.5%.

Sep 2016 – May 2017 **Research Analyst (Part-time role - 20 hours/week)**, *Intersect Inc*, Durham, NC, USA

- Developed a thematic investing algorithm that outperformed market benchmarks by 15% in pilot tests.

Jun 2016 – Aug 2016 **Software Engineer Intern**, *Inmar*, Winston-Salem, NC, USA

- Developed regression models to forecast product sales for retail customers across 21 US states, analyzing $15MM in yearly revenue.

Aug 2013 – May 2015 **Engineer (Business Technology Analyst)**, *Deloitte*, India

- Developed data infrastructure and pipelines to process and store sensitive healthcare data, handling 2 million records daily.

Technical and Research Skills

Programming	Python (PySpark), SQL (SparkSQL)
ML/Science	Regression, decision trees/random forests, deep learning (CNNs, RNNs with LSTM, transformers, deep reinforcement learning), PyTorch, TensorFlow
Statistics	Econometrics (time series forecasting), game theory (mechanism design, incentive compatibility), experimentation, hypothesis testing, bayesian inference and reasoning, observational causal inference
Cloud/MLOps	AWS (Sagemaker, RDS, S3, EC2), Docker, GIT, Kubernetes
Logging/Viz	Elastic Search, Logstash, Kibana (ELK) stack, Grafana
LLM agent frameworks	LangChain, Semantic Kernel

Certifications

2024	Technical Fundamentals of Generative AI (Offered by Stanford Engineering — Online)
2024	Mathematics for Machine Learning and Data Science (Offered by DeepLearning.AI)
2023	Probability & Statistics for Machine Learning & Data Science (Offered by DeepLearning.AI)
2023	Calculus for Machine Learning & Data Science (Offered by DeepLearning.AI)
2022	Linear Algebra for Machine Learning & Data Science (Offered by DeepLearning.AI)

Professional Affiliations

Member, Association for Computing Machinery (ACM)

Member, The Econometric Society

Member, Association for the advancement of artificial intelligence (AAAI)