DAKSHA LADIA

+1(413)466-1092 | dladia@umass.edu | linkedin.com/in/dakshaladia | github.com/dakshaladia | dakshaladia.me

EDUCATION

University of Massachusetts Amherst, United States: Masters in Science in Computer Science **May 2026**
Relevant Coursework: Information Retrieval, Machine Learning, Natural Language Processing, Algorithms for Data Science, Theory & Implementations in Advanced Software Engineering, Reinforcement Learning, Trustworthy and Responsible AI
Manipal Institute of Technology, India: Bachelor of Technology in Computer Science & Engineering **June 2020**

SKILLS & CERTIFICATIONS

Technical skills: Java, Python, C#, C++, Machine Learning, Data Structures & Algorithms, PyTorch, Agile Methodologies (CI/CD), Python libraries (Pandas, Numpy, Scikit-learn), Object Oriented Programming, Retrieval-Augmented Generation (RAG), Large Language Models, Generative AI, AWS, Hugging Face, PySpark, Deep Learning, R, Docker, DBT, MLX, Differential Privacy, Opacus, SQL, MySQL, MongoDB, PostgreSQL, Git, GitHub, Snowflake

WORK EXPERIENCE

System1, Machine Learning Engineer Intern May 2025 - Aug 2025

- Built a multi-agent AI pipeline for processing natural language queries that use business-specific data, using Text-to-SQL generation with a RAG-based feedback loop, with a 98% success rate, and an evaluation framework for the pipeline

Microsoft, Bing Ads Team, Software Engineer June 2020 - July 2024

- Optimized selective sampling of Bing Ads' top advertiser listings, working on recommender systems, to optimize ad simulation result calculations, used in ad search and ranking, improving runtime by 15%, cost savings of 9%, reducing data simulations by 5% and enhancing the efficiency of the ad ranking system by 14%
- Engineered a Regression Tester Tool from scratch to automate the validation of the impact of new integrations or features against alternative code branches efficiently, reducing manual effort and saving developer hours by 90%
- Enhanced accuracy of ad performance analysis by improving the Machine Learning model used in ad-click prediction analysis, resulting in a 2% and 4% AUC gain in international and US markets, respectively
- Executed Strategic Global Market Optimizations through comprehensive hyperparameter tuning, overcoming challenges such as unanticipated market-specific behaviors, and increasing the average user engagement by 4%

FarmGuide, Data Science Intern Jan 2020 - May 2020

- Developed and deployed a Machine Learning model and forecasting algorithm for crop identification, health monitoring, and prediction, achieving 96% accuracy in recommending crops based on farm features and geo-locations

Microsoft, Bing Ads Team, Software Engineer Intern May 2019 - July 2019

- Analyzed seasonality patterns and anomalies in advertiser campaigns, using time-series modelling and forecasting techniques, focusing on auction dynamics and auto-bidding algorithms to optimize bid pricing based on advertiser goals

PROJECTS

Recomm AI - Privacy-Preserving Shopping Recommendation Assistant, github.com/dakshaladia/Recomm-AI

- Developed a privacy-preserving shopping recommendation system leveraging user behavior and experiences to suggest personalized products. Secured datasets using Differential Privacy and validated robustness via Membership Inference Attacks. Experimented with Federated Learning for on-device deployment and evaluated system resilience against misalignment, jailbreaking, and prompt injection attacks on LLMs.

Multi-agent AI Browser Automation Platform, github.com/dakshaladia/multiagent-browser-automation

- Built a vision LLM-powered browser automation system that navigates web applications using screenshot analysis and DOM parsing to execute multi-step workflows autonomously. Architected a modular framework with multi-LLM provider support, hybrid state detection, and automated dataset generation for workflow capture and analysis.

Secure, On-device, Privacy-first Personalized AI Companion

- Developed a multimodal (voice + text) AI companion that runs entirely on-device, ensuring all conversations remain private. Implemented federated learning for personalized adaptation to the user's tone and mood over time, enabling secure, privacy-preserving personalization without any data leaving the device.

Automation of Information Extraction from Medical Documents, github.com/dakshaladia/Doc-Process-Automation

- Built a pipeline that uses GPT-4o models and Google Vision API to classify different healthcare documents, performed OCR utilizing a combination of Large Language models like GPT and Claude, and converted unstructured text to structured text. Automation here reduces the time to read and parse these docs by 99%

Complex Query Synthesis for Enhanced Information Retrieval, github.com/dakshaladia/QueryGenerationAndRetrieval

- Worked on a research project for building a new pipeline for enhanced Document Retrieval using document segmentation, pseudo query generation, fine-tuning large language models, and semantic matching to mitigate issues countered by short, ambiguous queries through semantic matching methods, gaining 4% precision improvement over BM25.

Research Study on Bias in Large Language Models, link to the paper (preprint)

- Conducted in-depth analysis of pronoun and occupational biases in LLMs using targeted prompts and statistical tests, and published a detailed research study report outlining mitigation approaches.

AI SlackBot, github.com/dakshaladia/AI-SlackBot

- Developed a personal Slackbot to help summarise and answer any questions discussed on a specific Slack channel