

Saaketh Chaganty

GEN AI REASONING · SCALABLE BACKEND · FULL STACK AI

Boston, MA

(551) 280 6937 | chaganty.s@northeastern.edu | saakethypes.github.io | github.com/saakethypes | linkedin.com/in/saakethchaganty

Education

2022-24 **MS in Information Systems**, *Northeastern University* | Boston, MA 3.5/4.0
2017-21 **Bachelor's in Mechanical Engineering**, *Mahindra University* | Hyderabad, India 7.5/10

Courses: Deep Learning | Machine Learning | Algorithmic Digital Marketing | Neural Modeling | Data Science Eng. Methods | Research Methods in AI | Program Structures | Advanced Data Analytics | Application Data Modeling & Development | Enterprise Web Development |

Skills

Programming Python, JavaScript, C, C++, Go, Java, SQL, Clingo, HTML, TypeScript, FastAPI, Django, Autogen, PyTorch, OpenCV, DSPy, TensorFlow, Matplotlib, NodeJs, ReactJS, Snowflake, Supabase, GraphQL
Software Linux, Docker, Neo4j, Git, GitLab, GCP, Azure, MS Power Apps, AWS Lambda, MongoDB, Pinecone, PowerBi, DynamoDB, Tableau, PostgresDB, PySpark, PyReason
Soft Skills Time Management, Problem-solving, Holdership, Documentation

Experience

Devocratic Technologies, *Lead Backend Engineer* | Delhi, India *May 2021 - Aug 2022*

- Engineered a scalable NoSQL single-table database and created RESTful APIs for a prediction market trading platform (similar to nadex) in AWS DynamoDB, enabling secure & concurrent processing of 10K+ transactions for 40,000+ users, managing funds exceeding 20,000\$+.
- Designed and implemented a robust price change algorithm leveraging real-time data analytics to accurately reflect market demand and supply ensuring optimal market equilibrium. Led cross functional teams constructing CI/CD pipelines with GitLab.

UrbanKisaan, (Y Combinator 'S20) *AI Ops Intern* | Hyderabad, India *Feb 2021 - May 2021*

- Deployed a YOLO R-CNN model for a smart hydroponics farm (92% accuracy) to identify harvest stages in 12 crops, leveraging TensorFlow and PyTorch for training on rover-captured crop images, facilitating precise identification for 4 tonnes of produce.
- Led a team of 4 in streamlining end-to-end AI integration for dynamic yield-based discounted pricing for E-commerce marketplace; securing a feature in TechCrunch 2021 for pioneering autonomous large-scale farming, contributing to a 23% surge in sales using the ML pipeline.

OTSI, *Data Science Intern* | Hyderabad, India *April 2020 - Sep 2020*

- Deployed a robust face recognition OpenCV model for a facial attendance system at scale with a team employing professional ML lifecycles and developing 8+ SQL stored procedures, successfully prototyping within the company for 200+ employee registers.

Projects

Full-Stack LLM Reasoning Engine *Jan 2024 - Mar 2024*

- Developed a safe reasoning engine glued to an LLM framework facilitating safe execution following complex rules, facts and constraint optimizations with a proactive QA GPT. Additionally capable of querying and executing Py & SQL code.
- Augmented thought templates and utilized answer set programming, with reasoning and grounded solvers; Tech stack Docker, Neo4J, Clingo, PgVector & OpenAI APIs.
- Innovated knowledge engineering leveraging hierarchical graph RAG knowledge base and few shot prompting. Handling structured atomic fact extraction from unstructured website/pdf chunks achieving a 95% accuracy in complex task completion.

Multi-Modal CNN Transformer Git *Sep 2023 - Nov 2023*

- Modeled a CNN transformer architecture for satellite images & weather station data with Conv2D LSTM & GRU attention layers performing predictive forecasting of lake-effect rain at lake Michigan.
- Researched and designed 12 visual & statistical features for cloud formations additionally utilizing SHAP toolkit for FE; Trained the model within 100 epochs to generate 3 days forecast for light rainfall with a 76% accuracy.

LLM Ad Targeting on Social Media Personas Git *Mar 2023 - May 2023*

- Built a social media platform with personalized ad delivery using a fine tuned ViT-G image captioning model to extract context from user posted images, enabling targeted ad placement based on image liking behavior and maintaining an evolving persona for each user.
- Developed a high-performance reverse image search using ImageNet pretrained weights for efficient image search over 10K+ images within 21ms using Pinecone, DynamoDB & Streamlit stack.

E-commerce Customer Churn Prediction Git *Mar 2023 - May 2023*

- Applied Logistic Regression using Snowflake & Scikit-Learn to predict churn customers expenditure for an e-commerce platform.
- Performed GridSearchCV, XGBoost, oversampling, and SMOTE to enhance model performance and achieved a 90% accuracy on the sales data.

Agoda Clone Java MVC Git *Mar 2023 - May 2023*

- Architected a sophisticated role-based Java OOP backend system with advanced RBAC database design, ensuring modularity, scalability, and robust data integrity.