Bhaskar Akkena

Rochester, NY 14623
585-710-8744 akkena.pavan@gmail.com linkedin.com/in/akkena/

Education

Rochester Institute of Technology
Master of Science in Data Science, GPA: 3.6 *Rochester, NY*

Rochester Institute of Technology
Bachelor of Science in Computing and Information Technologies, GPA: 3.6 *Rochester, NY*

Technical Skills

Programming: Python, Java, JavaScript, SQL, Bash, PHP
Software Engineering: OOP, Data Structures & Algorithms, SDLC, Debugging & Troubleshooting, Git
Data Science & AI/ML: Machine Learning, Statistical Modeling, Data Analysis, Model Evaluation, Data Visualization, PyTorch, Tableau, TensorFlow, Keras, scikit-learn, Pandas, NumPy, R, Matplotlib, Seaborn
Data Engineering & Databases: Apache Spark, Databricks, Kafka, MySQL, MongoDB, Neo4j
Cloud & Other Tools: AWS, Docker, Kubernetes, CI/CD, Infrastructure Automation
Systems: Linux/Unix, Windows Server, Active Directory, VMware, System Administration, Monitoring

Work Experience

Lowe's Companies, Inc **Jan 2026 – Present**
Software Engineer - ML *Charlotte, NC*

- Translated loosely defined business problems into robust **data science models**, converting requirements into measurable objectives, sampling strategies, evaluation criteria, and reproducible analytical workflows.
- Curated labeled datasets across **behavioral, text, and catalog domains**, designing sampling strategies, establishing **ground truth**, and enforcing label-quality standards for reproducible model training and evaluation.
- Evaluated **ML models and LLMs** across in-house training, fine-tuning, and external API approaches, comparing **accuracy, coverage, failure modes, and cost** to select production-ready solutions.
- Partnered with engineering teams to productionize ML solutions by implementing **reproducible testing, data-drift monitoring, and automated retraining triggers**, while documenting error costs and methodology for product and policy stakeholders.

Saunders College of Business **Jan 2023 – May 2024**
System Administrator *Rochester, NY*

- Engineered migration of applications from **Docker to Kubernetes**, implementing namespaces, deployments, and ingress configurations to create scalable, repeatable deployment workflows and improve deployment efficiency by **40%**.
- Developed and standardized automated provisioning workflows for **200+ systems** using **Ivanti** imaging and deployment tools, reducing repetitive configuration and improving consistency across environments.
- Implemented **Zabbix** monitoring across **500+ enterprise systems**, analyzing performance metrics and system logs to identify bottlenecks, troubleshoot failures, and reduce infrastructure downtime by **25%**.
- Resolved **2,000+ application and infrastructure issues** through systematic debugging, log analysis, root-cause analysis, and incident triage using **OTRS**, maintaining a **95% on-time resolution rate**.

Rochester Institute of Technology **Aug 2022 – Dec 2022 — Jan 2025 – Dec 2025**
Teaching Assistant *Rochester, NY*

- Mentored **150+ students** in **Python and Bash** for task automation, Linux scripting, file processing, and command-line programming through hands-on laboratory instruction.
- Conducted **weekly labs, code reviews, debugging, and grading** for **1,000+ programming assignments**, evaluating program logic, scripting implementations, and automation workflows.
- Developed supplemental exercises covering **Linux command-line operations, file handling, and cron automation**, providing practical experience in Python/Bash development and automated workflows.

Projects & Research

Bias Characterization in AI-Generated Travel Narratives **Jan 2025 - Dec 2025**

- Conducted **end-to-end AI research** on cultural and representational bias in LLM-generated travel narratives, analyzing how **model behavior, linguistic context, and evaluation methodology** influence perceived bias and toxicity.
- Architected a scalable **5-layer Python data pipeline** processing ∼**600,000 online travel comments**, including **Data API ingestion, text preprocessing, language identification, semantic relevance filtering, LLM generation, and automated evaluation**.
- Evaluated **100,000+ relevant samples** across **LLaMA 3.2, Gemma 2, GPT-2, and Sarvam** using controlled generation parameters and reproducible batch processing to conduct cross-model and cross-language experiments.

- Developed an evaluation framework combining **human annotations, Detoxify toxicity metrics, correlation analysis, distributional analysis, and Bland–Altman agreement testing**, identifying systematic gaps between **human and automated toxicity assessments**.
- Identified an **AI blind spot** in automated toxicity detection, demonstrating that implicit/contextual harm and linguistic framing can cause safety evaluators to **underestimate human-perceived toxicity**, with implications for LLM evaluation and responsible AI.

Evolutionary Search for Toxicity Control in Large Language Models (LLMs) **Aug 2025 - Dec 2025**

- Conducted **AI research** on automated toxicity control in LLMs by developing an **evolutionary search framework** that iteratively optimized generated text against toxicity objectives.
- Built an automated **LLM evaluation pipeline** integrating **Perspective API, HateBERT, and embedding-based metrics** to assess toxicity across **5,000+ generated responses**.
- Designed a **context-aware semantic toxicity scoring system** using **SentenceTransformers** to capture contextual relationships beyond surface-level toxic keywords, achieving **92% scoring consistency**.
- Implemented an **iterative evolutionary optimization loop** for response refinement, achieving a **37% reduction in mean toxicity scores** across experimental generations.
- Developed **Python-based experiment tracking and visualization** with structured JSON logging and Hugging Face models to monitor generations, benchmark model behavior, and analyze optimization trends.

Fraudulent Job Posting Classification Model **Aug 2024 - Dec 2024**

- Investigated the problem of fake job postings in online job markets by analyzing publicly available records and developed a **predictive model**, aiming to improve data reliability and protect job seekers from fraudulent opportunities.
- Engineered a robust data pipeline to perform **data cleaning**, and **preprocessing** on a dataset spanning 5 years of job postings across 50+ countries, with data on 100+ companies.
- Applied **text preprocessing**, **feature extraction**, and **model tuning** to analyze 10,000+ job postings, optimizing performance through 5-fold cross-validation to improve generalization and reduce overfitting.
- Implemented **Logistic Regression**, **Random Forest**, and **Gradient Descent** algorithms to predict fake job postings; achieved 72% F1-score, 80% Precision, and 67% Recall through **hyperparameter tuning**.

Tutor Scheduling Application **Jan 2024 - May 2024**

- Led the senior year capstone project at **Rochester Institute of Technology (RIT)**, collaborating with **GCCIS** to design and develop a **Tutor Scheduling Application** that automated tutor hour tracking and reduced administrative workload by **50%**.
- Developed scalable backend **APIs** with **Node.js** to handle scheduling, session tracking, and authentication; implemented **role-based access control** and **audit logging** in an existing **SQL** database.
- Enhanced system reliability by optimizing **database queries** and **stored procedures**, improving performance and enabling real-time synchronization for administrators and tutors.

Application Performance Monitoring Tool **Aug 2023 - Dec 2023**

- Developed a lightweight **Application Performance Monitoring (APM)** solution for **Linux** systems to collect and visualize key process and performance metrics in real time.
- Created interactive **Python**-based dashboards that transformed raw log data into actionable insights, helping diagnose **system bottlenecks** and improve decision-making.
- Improved overall **system uptime** by detecting anomalies proactively and deploying optimized solutions, reducing **downtime** and performance degradation by over **30%**.