

Abdul Basith Mohammed

Data Scientist

abdulbasith.mohd02@gmail.com

<https://www.linkedin.com/in/abdulbasithmohammed/>

<https://github.com/AbdulBasithMohammed>

+1 (902) 449-5102

Toronto, CA

Data scientist with experience building and deploying production-grade AI systems that deliver measurable real-world impact. Developed scalable solutions across computer vision, generative AI, and data-intensive applications, including a high-accuracy ANPR system and a multi-agent AI assistant. Experienced in designing robust model-driven systems, optimizing retrieval and inference workflows, and integrating machine learning solutions with cloud and distributed data infrastructure.

Work Experience

Data Science Intern

Jan 2026 - Apr 2026

Manulife

- Architected and implemented core systems for a high-visibility, production-grade RAG chatbot serving employees across 15+ countries, delivering features on a 2–3 day release cycle under direct executive oversight — including a multi-index hybrid retrieval framework, parallel query preprocessing pipeline, and acronym enrichment module.
- Conducted exploratory research and technical scoping for a Predictive Attrition Model, translating stakeholder requirements into a structured PoC roadmap for early employee flight-risk detection using internal datasets.
- Collaborated with ML Engineers and Data Scientists in an AI Context Engineering workshop to design a natural-language-to-analytics chatbot that translates plain English queries into structured Pandas workflows, enabling non-technical users to perform data exploration and visualization without writing code.

Associate ML Engineer

Jan 2023 - May 2024

Escapade Technologies

- Developed and deployed an Automated Number Plate Recognition (ANPR) system along with training a license plate object detector, achieving a 5x increase in correctly detected vehicles, improving plate-read accuracy from ~30% to 97% in production.
- Developed interactive dashboards using Power BI and IBM Cognos and built ETL pipelines (Python, Pandas, SQL, Regex, API integrations) to extract transactional and operations data, transform it for consistency, and load it into reporting systems, enabling automated insights and improved client decision-making.
- Focused on data collection using SQL queries and REST APIs, preprocessing datasets with Python (Pandas, NumPy, Regex), and optimizing ETL workflows to improve data reliability and processing efficiency.
- Collaborated on client-facing analytics projects such as income prediction and sales performance, applying multiple linear and logistic regression with one-hot encoding, grid search, and k-fold cross-validation, and visualizing results with predicted vs. actual comparisons and KPI dashboards in Power BI/Seaborn.

- Developed and deployed a machine learning model, 'Average Income Predictor', leveraging Census Bureau datasets, achieving 94% prediction accuracy on test data.
- Enhanced model robustness using PCA for dimensionality reduction, grid search tuning, and XGBoost, ensuring consistency via k-fold cross-validation and F1-score.
- Utilized Matplotlib and Seaborn to design and present interactive dashboards and visualizations, highlighting patterns and insights from large-scale socioeconomic datasets.

Projects

AskHR Virtual AI Assistant

- Redesigned a production RAG pipeline by consolidating 6-8 sequential preprocessing agents into a unified module with async parallel LLM execution, reducing end-to-end query latency from 14s to ~2s (85% reduction).
- Worked on a multi-index hybrid retrieval system implementing a 3-stage search pipeline across Azure AI Search that fuses vector and BM25 keyword scoring with semantic reranking, region-hierarchical filtering, and manager-aware chunk pairing to serve employees across 15+ countries.
- Implemented a hybrid acronym enrichment agent combining hash-based region matching with selective LLM fallback, reducing unnecessary model calls and improving multilingual query grounding.

MIMIC-III Medical RAG Document Encoder & Retrieval System

- Developed a multi-encoding medical document retrieval system that processed 44,915 MIMIC-III patient records using 4 different text chunking strategies (document-level, recursive character, semantic chunking, and entity-based section parsing).
- Implemented a vector database using Qdrant to store medical documents, supporting 4 distinct collections with method-specific chunking strategies (full-document, 1000-char recursive with 20% overlap, semantic percentile-based, and entity-section parsing).
- Built a 3-tier retrieval evaluation framework combining cosine similarity search, medical keyword matching, and LLM reranking (Google Gemini), demonstrating 7.2x improvement over keyword matching at meaningful similarity thresholds across 118 medical query evaluations.

Recursive Reasoning in Tiny Networks: Architectural Analysis & Evaluation

- Conducted a deep architectural analysis of the Tiny Recursive Model (TRM) and Hierarchical Reasoning Model (HRM) to evaluate recursive reasoning mechanisms and parameter efficiency, synthesizing findings across Sudoku-Extreme, Maze-Hard, and ARC-AGI benchmarks as measured by reported accuracy and ablation studies.
- Deconstructed and compared recursive supervision, gradient flow strategies, and latent state design in TRM vs HRM to understand performance gains from full backpropagation through recursion, analyzing reported improvements such as 87.4% vs 56.5% Sudoku-Extreme accuracy under different gradient approximations.
- Delivered a technical presentation explaining architectural differences between TRM and HRM, translating complex recursive reasoning concepts into clear visual and conceptual explanations for peer evaluation.

Traffic Management Automatic Number Plate Recognition (ANPR)

- Trained and fine-tuned a YOLO-based license plate object detector on a custom dataset, enabling 94% detection accuracy in real-world traffic scenarios, competitive with state-of-the-art models that report 95% accuracy in similar conditions.
- Developed a lightweight ANPR system deployable on any IP camera, reducing hardware costs and installation complexity by eliminating the need for specialized equipment.
- Integrated real-time image processing and OCR-based license plate recognition, achieving over 90% accuracy in varied lighting and angle conditions, aligning with leading systems that achieve 96.9% end-to-end recognition rates across multiple datasets.
- Built a scalable backend database using Pandas and CSV file management with automatic data organization, storing license plate reads, timestamps, images, and metadata for easy querying and analysis.
- Designed an intuitive user interface using CustomTkinter with CtkTable components, enabling switching between camera feeds, viewing historical records, and performing partial plate searches with date filters.
- Added an object tracking module using YOLOv8 with ByteTrack algorithm to calculate average vehicle exit times, providing analytics on congestion points for traffic optimization.

CVPRO + - ATS Resume Builder & Analyzer

- Developed an AI-powered resume parsing system using Google Gemini API and advanced prompt engineering, achieving 98%+ accuracy in extracting personal info, work experience, skills, and education into standardized JSON.
- Built an intelligent ATS analysis engine with NLP + ML scoring across five dimensions (contact info, formatting, content, keywords, professional impact), delivering actionable feedback and optimization suggestions.
- Implemented LLM-driven job-tailored resume generation that analyzes job descriptions, extracts relevant keywords, and dynamically customizes resumes to improve applicant success rates.
- Engineered a robust document processing pipeline with spaCy NLP, regex, and validation libraries to handle diverse formats, ensure structured data extraction, and provide fallback mechanisms for edge cases.

DALScooter – Serverless Bike Rental Platform

- Designed and implemented a serverless backend using AWS Lambda, API Gateway, and DynamoDB, building RESTful APIs for authentication, feedback, and analytics with robust error handling and validation.
- Developed a real-time analytics and visualization module leveraging Lambda, DynamoDB, and Hugging Face's sentiment analysis API, with Looker Studio dashboards for usage metrics and customer trends.
- Integrated AWS services (Cognito, DynamoDB, Lambda, API Gateway, CloudWatch) to deliver a scalable, secure, and cost-effective cloud-native architecture.

Distributed Maritime Data Management & Analytics Platform

- Designed and implemented a distributed maritime database with 20+ normalized tables for vessel tracking, port operations, and AIS data, using horizontal fragmentation across MySQL instances with a global data catalog.
- Built a data analytics pipeline in PySpark, MongoDB, and Python to process email datasets with sentiment analysis, DBSCAN clustering, and traffic visualization, enabling automated insights and reporting.

Core Skills

Programming Languages: Python, SQL, Java, C, C++, HTML, CSS, JavaScript

Libraries & Frameworks:

PyTorch, Pandas, PySpark, Matplotlib, Seaborn, OpenCV, Ultralytics, spaCy, Regex, CustomTkinter, LangChain

Databases & Vector Stores:

MySQL, PostgreSQL, MongoDB, DynamoDB, Amazon ElastiCache (Redis), Qdrant, Azure AI Search, S3

Cloud Platforms & Services:

Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform (GCP), Databricks

Development Tools & Version Control:

Git, Jira, Jupyter Notebook, VS Code, Google Colab, Anaconda, Notepad++, Microsoft Teams, Slack, Microsoft 365, Github/Gitlab

Operating Systems: Windows, macOS, Linux, Unix, Ubuntu

Analytical & Problem-Solving:

Critical thinking, Structured problem-solving, Data-driven decision-making, Detail-oriented

Collaboration & Communication:

Strong teamwork, Cross-functional collaboration, Technical documentation, Stakeholder communication

Adaptability & Innovation:

Quick learner, Adaptable to new technologies, Innovative mindset, Continuous improvement

Project & Time Management:

Agile methodology familiarity, Sprint planning, Requirements gathering, Task prioritization, Deadline management

Leadership & Initiative:

Proactive contribution, Mentoring peers, Ownership of deliverables, Ability to work independently and in teams

Education

Dalhousie University

Sep 2024 - Apr 2026

M.S. Applied Computer Science

Osmania University

Jun 2019 - Jun 2023

B.E. Information Technology