CAROLINE GAKII

DATA SCIENTIST | AI TRAINER | LLM EVALUATOR | MACHINE LEARNING RESEARCHER

Meru, Kenya | caratechconsultancy@gmail.com | LinkedIn | GitHub | Mindful Data Science | Portfolio

PROFESSIONAL SUMMARY

Data Scientist and AI professional with hands-on experience evaluating large language models, developing machine learning solutions, and applying AI to real-world problems. Currently pursuing a **Master's degree in Data Science** while combining research experience with practical industry work evaluating frontier language models at Turing. Master's research focuses on AI-powered computer vision for medical diagnosis, training and evaluating CNN architectures including **ResNet50V2, EfficientNetV2, MobileNetV2, DenseNet121, and CoAtNet** for histopathology image classification. Skilled in Python, SQL, TensorFlow, machine learning, deep learning, computer vision, and AI quality assurance. Communicates technical findings through published research articles reaching a global audience of engineers and researchers.

CORE COMPETENCIES

AI EVALUATION
LLM Evaluation
AI Model Assessment
Prompt Evaluation
RLHF Concepts
Supervised Fine-Tuning
Structured Reasoning
AI Quality Assurance
Response Ranking
Prompt Engineering

MACHINE LEARNING
Machine Learning
Deep Learning
Computer Vision
CNN Development
TensorFlow · Keras
Model Validation
Hyperparameter Tuning
Transfer Learning

DATA SCIENCE
Python · SQL
Pandas · NumPy
Data Analysis
Data Cleaning
Feature Engineering
Data Visualization
Statistical Analysis
Matplotlib

EVALUATION
Root Cause Analysis
Fact Verification
Error Analysis
Analytical Reasoning
Rubric-based Eval
QA Alignment
Documentation
Technical Writing

PROFESSIONAL EXPERIENCE

AI Trainer

May 2026 — Present

Turing

- Evaluate AI-generated responses across complex reasoning and instruction-following tasks using structured evaluation rubrics.
- Compare multiple model outputs to identify the highest quality response based on accuracy, reasoning, and instruction adherence.
- Identify **hallucinations, factual inaccuracies, logical inconsistencies, and reasoning failures** in LLM outputs.
- Write detailed, evidence-based justifications supporting evaluation decisions with clear analytical reasoning.
- Work on evaluation projects involving **NVIDIA Nemotron and Qwen** large language models.
- Design prompts that challenge model reasoning and expose failure cases across diverse task types.
- Maintain high annotation quality while adhering to strict evaluation standards and consistency guidelines.
- Contribute human feedback used to improve large language models via RLHF pipelines.

Business Analyst

March 2026 — May 2026

Turing

- Conducted side-by-side comparisons of AI-generated responses applying structured evaluation criteria.
- Validated factual correctness of AI outputs using reliable, verifiable evidence sources.
- Identified reasoning errors, unsupported assumptions, and logical inconsistencies across evaluation tasks.
- Produced clear, evidence-based explanations for evaluation decisions to support quality improvement.

Founder & Editor — Mindful Data Science

2025 — Present

Medium Publication

- Founded publication making AI, ML, and data science accessible through research-driven practical tutorials.
- Published technical articles on computer vision, domain adaptation, preprocessing pipelines, and AI ethics.
- Translate complex ML research into practical guides for engineers and researchers globally.

MACHINE LEARNING RESEARCH PROJECTS

AI-Powered Digital Microscope for Cancer & Malaria Diagnosis

Master's Research

- Developed an AI-assisted digital pathology workflow using computer vision for histopathology image classification.
- Trained CNN architectures: **ResNet50V2, EfficientNetV2, MobileNetV2, DenseNet121, and CoAtNet**.
- Performed dataset preprocessing, augmentation, class balancing, and model optimization.
- Evaluated models using Accuracy, Precision, Recall, F1-score, ROC-AUC, Sensitivity, and Specificity.
- Investigated **domain shift and model generalization** across multiple medical imaging datasets.
- Integrated AI pipeline with an **OpenFlexure digital microscope** for low-cost diagnostic support.

TensorFlow | Keras | ResNet50V2 | EfficientNetV2 | CoAtNet | Computer Vision | Healthcare AI | Python

BreakHis Breast Cancer Histopathology Classification

Research Project

- Built CNN-based classification system for breast cancer histopathology images using the BreakHis dataset.
- Trained and evaluated ResNet50V2, DenseNet121, and MobileNetV2 architectures with clinical performance metrics.
- Focused on sensitivity, specificity, and ROC-AUC to assess clinical reliability of each model.

TensorFlow | DenseNet121 | MobileNetV2 | Pathology AI | ROC-AUC

Gastric Cancer AI Classification with Domain Adaptation

Research Project

- Developed deep learning model for gastric cancer histopathology classification using CoAtNet architecture.
- Investigated domain shift and model generalization across multiple medical imaging datasets.
- Designed preprocessing pipeline for heterogeneous data to improve real-world deployment robustness.

CoAtNet | Domain Adaptation | Medical Imaging | TensorFlow

TECHNICAL SKILLS

Programming & Data	ML & Deep Learning	Tools & Platforms
Python · SQL	TensorFlow · Keras	Git · GitHub
Pandas · NumPy	CNN Architectures	MySQL
Matplotlib · Seaborn	Transfer Learning	Google Colab
Scikit-learn	Model Validation	VS Code
Jupyter Notebook	Hyperparameter Tuning	LLM Evaluation Platforms

SELECTED TECHNICAL PUBLICATIONS

I Trained ResNet, CoAtNet, and DenseNet on Real Cancer Images. Here's What Surprised Me — Mindful Data Science

What I Learned When My 96% Accurate Gastric Cancer Classifier Met the Real World — Mindful Data Science

What Domain Shift Taught Me About Building Machine Learning Models for Production — Mindful Data Science

Why AI Requires African Voices: Creating Inclusive Models for a Diverse World — Mindful Data Science

The Nightmare of Heterogeneous Data: Building an Invariant Preprocessing Pipeline — Mindful Data Science

SELECTED ACHIEVEMENTS

- Evaluated **frontier language models** (NVIDIA Nemotron, Qwen) through structured AI quality assessment at Turing.
- Applied analytical reasoning to identify logical errors and improve AI-generated outputs across diverse task types.
- Built **multiple computer vision models** for healthcare applications using state-of-the-art CNN architectures.
- Conducted research on **domain shift and ML model generalization** across heterogeneous medical imaging datasets.
- Published **5+ technical articles** on ML research, computer vision, and AI ethics in Mindful Data Science.
- Integrated AI pipeline with OpenFlexure digital microscope for low-cost diagnostic support in resource-limited settings.

EDUCATION

Master of Science in Data Science ***(Ongoing)*** — 2023 — Present

Meru University of Science and Technology

Machine Learning · Deep Learning · Data Mining · Time Series Analysis · Big Data Analytics · Optimization

Bachelor of Computer Technology — Completed

Meru University of Science and Technology

PROFESSIONAL STRENGTHS

Critical Thinking · Analytical Reasoning · AI Evaluation · Technical Communication · Research · Problem Solving · Attention to Detail · Continuous Learning

PORTFOLIO & LINKS

Portfolio: carolinegakii.github.io GitHub: CarolineGakii Mindful Data Science: medium.com/mindful-data-science

LinkedIn: caroline-gakii-75775415b

Made with IBM Bob