Desmond
Junior Machine Learning and AI Engineer

SUMMARY

- Machine Learning and AI Engineer with 5 years in LLM-based Document AI, RAG pipelines, and computer vision systems using Python, PyTorch
- Expertise in MLOps with Docker, Kubernetes, MLflow, DVC, and cloud platforms (AWS, Azure, GCP)
- Fine-tuning enterprise LLMs (LoRA/QLoRA), building vector search with FAISS, and integrating OCR (Tesseract, PaddleOCR) for multimodal document understanding.
- Strong background in data engineering and analytics using PySpark, SQL, Airflow, and visualization tools (Power BI, Tableau) to support ML workflows.
- Experience in trustworthy AI evaluation including robustness, hallucination analysis, latency benchmarking, and failure-mode diagnostics for reliable AI systems.

TECHNICAL SKILLS

Main Technical Skills	Python, SQL, Machine Learning, Deep Learning, Docker
Programming Languages	C++, Python, R
AI & Machine Learning	AWS SageMaker, AWS SageMaker (Amazon SageMaker), CNN, Deep Learning, LangGraph, LSTM, Machine Learning, Mistral AI, ResNet, Scikit-learn, Xgboost
Python Libraries and Tools	Matplotlib, Plotly, PySpark, Scikit-learn, Seaborn
Data Analysis and Visualization Technologies	Apache Airflow, Databricks, Jupyter Notebook, Power BI, Random Forest, Tableau
Databases & Management Systems / ORM	dbt, Microsoft SQL Server, MongoDB, MySQL, PostgreSQL, SQL
Cloud Platforms, Services & Computing	GCP
Amazon Web Services	AWS Bedrock, AWS EC2, AWS S3, AWS SageMaker, AWS SageMaker (Amazon SageMaker)
Azure Cloud Services	Databricks
Google Cloud Platform	Google Data Studio
Deployment, CI/CD & Administration	CI/CD
Virtualization, Containers and Orchestration	Docker

Version Control	Git, GitHub, Github Actions
Third Party Tools / IDEs / SDK / Services	Microsoft Visual Studio Code
Other Technical Skills	LoRa, MLP

WORK EXPERIENCE

Research Assistant, Document Intelligence Systems, Czestochowa University of Technology

Duration: 2025 - 2026

Summary:

- Designed and evaluated large language model-based document understanding systems for extracting structured information from complex, layout-rich PDF documents
- Developed retrieval-augmented pipelines integrating transformer models, OCR, vision-language models, embeddings, and vector search for robust document question answering
- This project focused on improving document AI reliability and interpretability through advanced AI techniques

Responsibilities:

- Designed and evaluated document understanding systems using LLMs and OCR technologies.
- Developed retrieval-augmented generation pipelines combining multiple AI models and vector search.
- Implemented reproducible experiments with various performance and reliability metrics.
- Analyzed hallucination patterns and grounding failures to enhance system robustness.

Technologies: Large Language Models, Transformers, OCR (Tesseract, PaddleOCR), Vision-Language Models, Embeddings (Sentence-BERT), FAISS vector search, Python, ANLS, Exact Match, F1 metrics

Team Participant, Trustworthy AI and Cybersecurity Challenges - CISPA, Warsaw, Poland

Duration: 2026

Summary:

- Participated in the European AI and Cybersecurity Hackathon focusing on trustworthy AI and cybersecurity challenges
- The project involved developing methods for dataset inference, LLM memorization extraction, watermark detection, calibration, and robustness analysis under hidden benchmark conditions

Responsibilities:

- Worked on dataset inference and LLM memorization extraction techniques.
- Contributed to watermark detection and calibration methods.

- Performed robustness analysis and failure-mode diagnostics.
- Engaged in architecture-aware statistical inference and contrastive decoding.

Technologies: Large Language Models, Statistical Inference, Contrastive Decoding, Robustness Analysis, Watermark Detection

Instructor - Landmark Certification Centre ML Training Programs

Summary:

- Delivered practical instruction in machine learning concepts including regression, classification, model evaluation, and Python-based workflows
- Guided learners through implementing models and understanding end-to-end training and evaluation processes

Responsibilities:

- Delivered practical machine learning instruction covering key concepts and workflows.
- Guided learners in model implementation and result interpretation.
- Provided end-to-end training and evaluation practice support.

Technologies: Python, Machine Learning (Regression, Classification), Model Evaluation

Enterprise LLM fine-tuning and MLOps

Summary: Implemented parameter-efficient fine-tuning workflows for large language models using LoRA/QLoRA, Hugging Face, MLflow, DVC, and cloud-based experimentation to enable reproducible model development.

Responsibilities:

- Developed fine-tuning workflows for enterprise LLMs.
- Managed reproducible experimentation pipelines using MLOps tools.

Technologies: LoRA, QLoRA, Hugging Face, MLflow, DVC, Cloud Platforms

Conversational Document Intelligence/RAG System

Summary: Developed a document-grounded conversational pipeline using sentence embeddings, vector retrieval, LangChain, and context-aware prompting to answer questions over unstructured content.

Responsibilities:

- Built conversational QA system grounded in document data.
- Implemented retrieval-augmented generation with embeddings and vector search.
- Applied context-aware prompting techniques.

Technologies: Sentence Embeddings, Vector Retrieval, LangChain, Prompt Engineering, RAG

Churn and Retention Analysis

Summary: Built a classical machine learning workflow for churn and retention analysis involving data cleaning, feature engineering, class-imbalance handling, Random Forest modeling, performance evaluation, and feature-importance analysis.

Responsibilities:

- Performed data cleaning and feature engineering.

- Handled class imbalance in datasets.
- Developed Random Forest models and evaluated performance.
- Conducted feature importance analysis.

Technologies: Python, Random Forest, Data Cleaning, Feature Engineering, Class-Imbalance Handling

Computer Vision Projects

Summary:

- Applied computer vision techniques including image classification and object detection using OpenCV, CNN/ResNet architectures, and YOLOv8
- Included transfer learning and model evaluation workflows

Responsibilities:

- Developed image classification and object detection models.
- Applied transfer learning techniques.
- Evaluated model performance and accuracy.

Technologies: OpenCV, CNN, ResNet, YOLOv8, Transfer Learning

EDUCATION

- **Master of Science in Artificial Intelligence and Data Science**
  Degree completed in July 2026. Academic average: 4.60/5.00; thesis grade: 4.50/5.00; diploma examination: 4.88/5.00; overall final result: 4.68/5.00. Relevant coursework includes Neural Networks and Machine Learning, Algorithms, Computer Vision, Pattern Recognition and Image Retrieval, Big Data and Data Mining, Autonomous Systems, Intelligent Signal Processing, Probabilistic Systems, Artificial Intelligence in Medicine, Computer Forensic Analysis.
  Feb 2025 - Jul 2026
- **Bachelor of Technology in Electronics and Telecommunications Engineering**
  Foundation in electronics, telecommunications, signal processing, networked systems, embedded-system concepts, and multidisciplinary engineering problem-solving.
  2016 - 2020

CERTIFICATIONS

- **Machine Learning Specialization - DeepLearning.AI / Stanford University, Coursera**
  Aug 2024
- **Python and MySQL Mastery - Code with Mosh**
  Mar-Apr 2023
- **Applied Machine Learning - University of Michigan, Coursera**
  Aug 2024
- **Business Analysis and Process Management - Coursera**
  Apr 2024
- **Data Analytics Fundamentals - IBM, Coursera**
  Feb 2024
- **Prepare Data for Exploration - Google, Coursera**
  Mar 2024
- **Data Analytics and Visualization Internship - BCG X**
  2024

- **Data Analytics and Visualization Internship - Accenture**
  2024
- **Data Analytics and Databases on AWS - AWS**
  Dec 2024
- **Data Science Internship - Quantum**
  2024
- **LangChain Certification - LLM Application Development**
  2025