

Jeffrey Kwan

Senior Data Scientist, Applied Intelligence

London

PRACTICE

Machine learning and artificial intelligence; large language models, transformers and embeddings; natural language processing; production ML deployment and MLOps.

SECTORS

Legal services; insurance and financial services; consumer goods.

“I like to start from first principles. The best machine-learning solutions come from understanding the data, the maths and the problem deeply, then choosing the right tool, not from reaching for the most fashionable one.”

Brief

Jeffrey Kwan is a senior data scientist at Linklaters, where he works across the firm’s flagship AI-driven products and initiatives, including **Drafting Risk Factors**, **Contract Analysis** and **Live Matters**. He has been instrumental in developing the core machine-learning logic behind Drafting Risk Factors, integrates new use-cases into Contract Analysis, and supports Live Matters work with rapid, interactive proof-of-concept demos delivered to tight deadlines. Known inside the team as a versatile, end-to-end data scientist who can plug into any project, Jeffrey works from first principles and applies cutting-edge techniques in model fine-tuning, including embedding and transformer models. He brings more than **5 years** of experience applying machine learning in production. Before joining Linklaters, he spent over five years at **Hastings Direct**, one of the United Kingdom’s largest general insurers, progressing from data analyst to senior data scientist. There he delivered multiple ML systems each generating **£1 million+** in commercial benefit, built end-to-end NLP and behavioural models, and supervised junior data scientists.

Selected AI Solutions

Plays a central role in **Drafting Risk Factors** at Linklaters, where he developed the core machine-learning logic that powers the product.

Integrates new use-cases into **Contract Analysis** and supports **Live Matters** work with rapid, interactive proof-of-concept demos delivered under short deadlines.

Works from first principles and applies cutting-edge techniques in model fine-tuning, including **embedding** and **transformer** models, balancing model sophistication with the practical demands of legal-domain accuracy.

At **Hastings Direct**, delivered **multiple ML models with £1 million+** commercial benefit each across pricing, operations and commercial functions, spanning regression (predicting repair costs and salvage values to optimise steerage), classification (large-injury propensity, liability) and behavioural (contractual breach and soft-fraud detection).

Managed a feature-engineering project at Hastings that generated **£1.6 million** in additional premium from historic quote data, working end-to-end with a junior data scientist he mentored.

Built NLP systems to extract value from more than **1 million** operations and medical notes, including a **Bidirectional Transformer** to predict deterioration from early-stage medical summaries, prompt engineering with LLMs to identify dangerous driving behaviour, and Word2Vec text embeddings reused across other downstream language models.

Productionised ML in **Databricks** and **Azure ML**, working with engineering to package models as **Docker** containers served via **Kubernetes** for live inference, and built offline batch-scoring pipelines using Azure ML Pipelines and LogicApp.

Education and Languages

Jeffrey holds a **BSc in Psychology** from the **University of Bristol**. The cross-disciplinary background, combining the empirical rigour of psychology with applied machine learning, informs his focus on understanding the underlying problem before choosing the model.

JEFFREY KWAN

Senior Data Scientist, Machine Learning & AI · London