UK AI & ML Engineering Salary Guide 2026

MLOps • AI Engineering • LLM Engineering • Data Engineering

The definitive guide to compensation for AI and machine learning engineering roles in the UK. Based on market data from Q4 2025 and January 2026, covering permanent salaries and contractor day rates across all experience levels.

KEY MARKET INSIGHTS

Median MLOps Engineer salary: £70,000 – £85,000

Median AI Engineer salary: £75,000 – £95,000

ML Engineer contractor day rate: £600 – £800

Senior/Lead roles commanding £100,000 – £150,000+

Contents

1. Executive Summary

The UK AI and ML engineering market remains highly competitive in 2026. Demand for production-focused engineers—those who can deploy, scale, and maintain ML systems—continues to outstrip supply.

Key Findings:

- **MLOps remains scarce:** Engineers with production experience command premium salaries, with median figures of £75,000–£85,000 in London.

- **LLM skills at premium:** Engineers with RAG, fine-tuning, and LLM deployment experience see 15–25% salary premiums over general ML roles.

- **Contract market stable:** Day rates for senior ML engineers remain strong at £600–£800/day, with specialists reaching £900–£1,100.

- **Leadership demand surging:** Head of MLOps and Director of AI roles are increasingly common, with salaries of £120,000–£180,000+.

- **Remote premiums narrowing:** London premiums remain at 10–20% over regional roles, but remote-first companies are compressing this gap.

2. MLOps Engineer Salaries

MLOps engineers bridge the gap between ML model development and production deployment. They build and maintain the infrastructure that enables data scientists to deploy models reliably at scale. Key skills include Kubernetes, Docker, CI/CD for ML, monitoring, and cloud platforms.

Permanent Salary Ranges (Annual, £)

Level	UK National	London	Remote UK
Junior (0–2 years)	£45,000 – £55,000	£50,000 – £62,000	£45,000 – £55,000
Mid-Level (2–4 years)	£55,000 – £75,000	£65,000 – £85,000	£58,000 – £75,000
Senior (4–7 years)	£75,000 – £95,000	£85,000 – £110,000	£75,000 – £95,000
Staff/Principal (7+ years)	£95,000 – £130,000	£110,000 – £150,000	£95,000 – £125,000

Skills Commanding Premium

• Kubernetes orchestration & Helm charts (+5–10%)

• MLflow, Kubeflow, or Weights & Biases (+5–8%)

• Terraform/Pulumi infrastructure-as-code (+5%)

• Ray or Spark for distributed training (+8–12%)

• GPU cluster management (+10–15%)

3. AI Engineer Salaries

AI Engineers build production AI systems, bridging ML research and software engineering. They integrate models into applications, optimise inference pipelines, and ensure AI systems work reliably at scale. This is distinct from research roles—AI Engineers focus on applied, production-grade work.

Permanent Salary Ranges (Annual, £)

Level	UK National	London	Remote UK
Junior (0–2 years)	£48,000 – £60,000	£55,000 – £68,000	£48,000 – £58,000
Mid-Level (2–4 years)	£60,000 – £80,000	£70,000 – £95,000	£62,000 – £80,000
Senior (4–7 years)	£80,000 – £105,000	£95,000 – £125,000	£82,000 – £105,000
Staff/Principal (7+ years)	£105,000 – £140,000	£125,000 – £170,000	£105,000 – £135,000

Most In-Demand Skills

- PyTorch production deployment (+8–12%)
- Model optimisation (quantisation, pruning) (+10–15%)
- Inference serving (Triton, TensorRT, vLLM) (+12–18%)
- Computer Vision systems (+8–12%)
- NLP/NLU production systems (+10–15%)

4. LLM Engineer Salaries

LLM Engineering is the fastest-growing specialisation in AI. These engineers build systems around large language models—not just prompts, but complete architectures including RAG pipelines, fine-tuning workflows, guardrails, and production inference at scale.

Permanent Salary Ranges (Annual, £)

Level	UK National	London	Remote UK
Junior (0–2 years)	£50,000 – £65,000	£58,000 – £75,000	£50,000 – £62,000
Mid-Level (2–4 years)	£65,000 – £90,000	£80,000 – £110,000	£68,000 – £90,000
Senior (4+ years)	£90,000 – £130,000	£110,000 – £150,000	£90,000 – £125,000
Lead/Principal	£130,000 – £180,000	£150,000 – £200,000+	£125,000 – £165,000

Premium Skills in LLM Engineering

- RAG architecture design & optimisation (+15–20%)
- Vector databases (Pinecone, Weaviate, Milvus) (+10–15%)
- LLM fine-tuning (LoRA, QLoRA, RLHF) (+15–25%)
- Inference optimisation (vLLM, TensorRT-LLM) (+12–18%)
- Agent frameworks (LangChain, LlamaIndex) (+8–12%)
- LLM evaluation & safety systems (+10–15%)

5. Data Engineer Salaries

Data Engineers build the pipelines and infrastructure that power ML systems. For AI-focused organisations, this means real-time feature stores, streaming pipelines, and ML-ready data architectures. The distinction between analytics and AI data engineering is increasingly important.

Permanent Salary Ranges (Annual, £)

Level	UK National	London	Remote UK
Junior (0–2 years)	£40,000 – £50,000	£45,000 – £58,000	£40,000 – £48,000
Mid-Level (2–4 years)	£50,000 – £70,000	£60,000 – £80,000	£52,000 – £68,000
Senior (4–7 years)	£70,000 – £90,000	£80,000 – £105,000	£70,000 – £88,000
Staff/Principal (7+ years)	£90,000 – £120,000	£100,000 – £140,000	£90,000 – £115,000

AI-Focused Data Engineering Premium Skills

- Feature store design (Feast, Tecton) (+12–18%)
- Real-time streaming (Kafka, Flink) (+10–15%)
- Databricks/Spark at scale (+8–12%)
- dbt + modern data stack (+5–8%)
- ML pipeline integration (+10–15%)

6. Leadership & Principal Roles

Senior technical leadership in AI/ML is highly sought after. These roles combine deep technical expertise with strategic thinking and team leadership. Compensation packages often include significant equity, especially at startups and scale-ups.

Leadership Salary Ranges (Annual, £)

Role	UK Range	London Range	Equity/Bonus
Head of MLOps	£120k – £160k	£140k – £180k	10–25%
Head of ML Engineering	£130k – £170k	£150k – £200k	15–30%
Director of AI	£150k – £200k	£170k – £250k	20–40%
VP of Engineering (AI)	£180k – £250k	£200k – £300k+	25–50%
Chief AI Officer	£200k – £350k+	£250k – £400k+	30–60%+

Note: Equity/Bonus figures represent typical percentage of base salary. Actual equity stakes vary significantly by company stage.

7. Contractor Day Rates

The contract market for AI/ML specialists remains strong. Day rates below are for contractors operating through their own limited company (outside IR35). Inside IR35 roles typically see 15–20% reduction in effective rate.

Day Rates by Role (£/day)

Role	Mid-Level	Senior	Lead/Principal
MLOps Engineer	£450 – £550	£550 – £700	£700 – £900
ML Engineer	£500 – £600	£600 – £750	£750 – £950
AI Engineer	£500 – £650	£650 – £800	£800 – £1,000
LLM Engineer	£550 – £700	£700 – £900	£900 – £1,200
Data Engineer	£400 – £500	£500 – £650	£650 – £850
ML Platform Engineer	£500 – £650	£650 – £800	£800 – £1,000

IR35 Considerations

The IR35 landscape continues to shape the contractor market. Key factors for outside IR35 status:

- Genuine project deliverables vs BAU support
- Right of substitution in practice
- Control over working methods
- Mutuality of obligation

8. London vs Regional Comparison

London maintains a salary premium for AI/ML roles, though the gap has narrowed with the rise of remote work. The premium varies by role and seniority.

London Premium Analysis (Senior Roles)

Location	vs London	Typical Range	Notes
London	Baseline	£75k – £120k	Highest salaries
Cambridge	–5% to –10%	£70k – £110k	AI research hub
Manchester	–10% to –15%	£65k – £100k	Growing AI hub
Edinburgh	–10% to –15%	£65k – £100k	Fintech & AI
Bristol	–12% to –18%	£62k – £95k	Growing tech
Remote UK	–8% to –15%	£68k – £105k	Varies by HQ

9. Startup vs Enterprise

Compensation structures differ significantly between startups and enterprises. Understanding these differences helps evaluate total compensation.

Compensation Structure Comparison

Factor	Early Startup	Scale-up (A–C)	Enterprise
Base Salary	70–85% market	90–100% market	100–110% market
Equity	0.1% – 1.0%	0.02% – 0.2%	RSUs: £10–50k/yr
Bonus	Rare	10–20% target	15–25% target
Benefits	Basic	Competitive	Comprehensive
Risk Profile	High	Medium	Low

10. Market Commentary

Key Trends for 2026

1. Production focus dominates

Companies prioritise engineers who can ship to production over those with purely research backgrounds.

2. LLM specialisation creates new tier

Engineers with RAG, fine-tuning, and inference optimisation experience command 15–25% premiums.

3. AI infrastructure is the bottleneck

GPU cluster management, distributed training, and cost optimisation skills are increasingly valuable.

4. Hybrid remains standard

Most companies offer 2–3 days in office. Fully remote roles available but often at slightly lower compensation.

5. Contract market stabilising

After IR35 disruption, the market has found equilibrium. Specialist contractors maintain strong rates.

11. Methodology

This salary guide is compiled from multiple data sources:

- • IT Jobs Watch salary data (Q4 2025 – January 2026)
- • Glassdoor UK salary submissions
- • Indeed UK salary data
- • Morgan McKinley 2025 Salary Guide
- • LinkedIn Talent Insights
- • PayScale UK data
- • Direct market intelligence from CodeRise Talent placements

Salary ranges represent the 25th to 75th percentile of market data. Actual salaries may fall outside these ranges based on company size, funding stage, candidate background, and negotiation. Data reflects the UK market as of January 2026.

NEED HELP HIRING AI & ML TALENT?

CodeRise Talent specialises in MLOps, AI Engineering,

LLM Engineering, and Data Engineering recruitment.

■ andrew@coderisetalent.io

■ www.coderisetalent.io