kubrick

The Big Squeeze

The consequences of a constrictive labor market for AI adoption

New research from Kubrick shines a light on the constricting data and AI labor market and what it could mean for AI strategy.

Mirela Gyurova, Talent Product Owner

Published 10th March 2026

AI and the slowing job market: correlation or causation?

The labor market is stalling, facing macroeconomic headwinds that prioritize productivity over headcount. LinkedIn data reveals a steady decline in global job postings since the pandemic (-20%), with large enterprises slowing hiring by almost 40% compared to pre-pandemic levels. And no one is feeling the squeeze more tightly than the entry-level workforce, where the average number of roles has declined by 29% since 2024.

> Misconception points to AI as the culprit for automating out junior-level roles.

In reality, low AI adoption rates mean it could not yet achieve this at any noticeable scale. Headlines continuously remind us that AI POCs are still largely stalling before production, with as many as 95% failing. Instead, researchers at the World Economic Forum point to economic uncertainty, monetary policy shifts, and post-pandemic hiring rebalancing as the main drivers of a slowing job market[1]. Although AI isn't replacing roles at scale, is it is certainly challenging them. In organizations where AI and automation are boosting productivity, namely Technology, Financial Services, and Healthcare, revenue per employee is increasing[2], meaning less investment in headcount to drive revenue.

If AI is causing market uncertainty elsewhere, what about roles in AI itself?

Unsurprisingly, this corner of the market is bucking the trend. LinkedIn reports 1.3 million new data and AI roles were created 2025, designed to fuel this technology boom. However, beneath the surface, the fast pace of change in AI is reshaping what those roles look like and the demand for skills and experience.

Mirela Gyurova, Talent Services Product Owner at Kubrick, investigates the current data and AI talent market and what it foreshadows about the success of AI adoption.

Methodology

We analyzed almost **1 million** job postings for data an AI roles on LinkedIn, across the US and UK markets. We tracked the influx and reduction of role postings across Q3 and Q4 2025 to create a snapshot of overall demand.

To understand demand levels for different role-types in data and AI, we grouped the job advertisements by skills and requirements into 6 key roles*:

01 **Data Engineering**

02 **Platform Engineering**

03 **Data and AI Governance**

04 **Data Product Management**

05 **AI and Machine Learning**

06 **Applied Decision Intelligence**

*See Appendix for role descriptions and skills requirements

To understand skill-level and experience, we used the self-reported experience level on a LinkedIn job advertisement:

Entry-level	Mid-to-Senior	Director	Executive

Due to the self-reporting nature, there may be some discrepancies in the skill level required at each experience level, but it is intended to capture market sentiment.

Contents

PART 1:

Constriction across the board

AI isn't replacing junior roles; experience is.

Our analysis found demand for mid-senior roles outweighs entry-level roles in all facets of data and AI, creating a constrictive labor market in the US and UK: inflated demand for more experience is constricting the flow of junior talent into a business.

The imbalance is most acute in the US, where there are more than 3.5x more mid-senior roles on the market than junior. In the UK, the demand is still 2.5x more.

US

Jobs Market Share - Seniority Breakdown

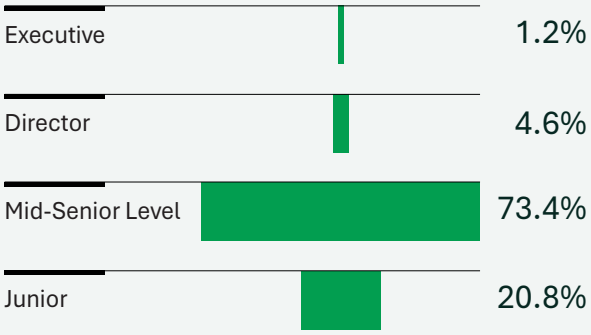

UK

Jobs Market Share - Seniority Breakdown

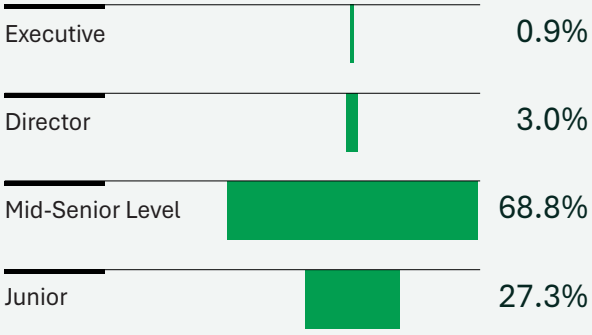

When analyzing this demand, there are a couple of concepts to keep in mind:

01

Market Churn:

This is the sum of both new roles being created and replacing existing roles. This is a hot market for enterprises who are trying to grow AI capability, as well as for the professionals who are cashing in on their experience.

02

Market Sentiment:

The seniorities of these roles are self-reported, so the experience of a 'mid-senior' role can vary from 3 years to 5 years to perhaps even 10, depending on the organization's definition. The upper and lower bounds are much wider than the typical 0-2 years of entry-level, so it will inherently encapsulate more roles. While self-reporting means we cannot exactly weigh up roles by number of years of in industry or expertise in a particular technology, we can use it to understand the sentiment of the market around the value of experience. In a still new and emerging market like data and AI, experience doesn't always equate to capability when certain tools and approaches are just a couple of years old – or less.

So, what's the problem with a constrictive labor market?

At a macro-level, constrictive labor markets risk sustainable growth and future workforce strategy. Instead of building a healthy workforce pyramid, which is essential for encouraging innovation and resilience, organizations are demanding more mid-to-senior level roles than entry-level, creating a diamond-shaped organization.

This impedes succession and leadership pipelines while limiting growth capacity, leading to employee churn. Meanwhile, expensive resources are stuck doing repetitive, functional tasks instead of focusing on strategic vision.

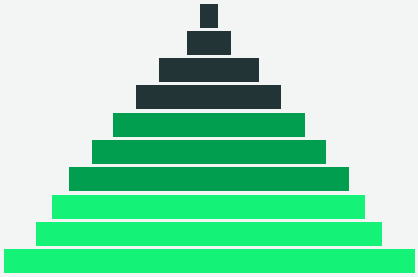

Expansive Workforce Model

- Cost-effective talent mix
- Innovation from emerging talent
- Succession planning
- Resilience to market shifts
- Higher productivity and ROI

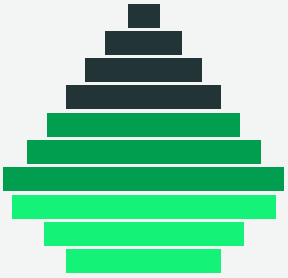

Constrictive Workforce Model

- High operational costs
- Leadership pipeline gaps
- Limited growth capacity
- Reduced organisational agility
- Ineffective use of expert capability

The immediate impact is also profound. Inflated demand for mid-to-senior talent is inflating wages in an already high-value market. Roles with AI skills are, on average, 23% higher than comparable roles without, leaping beyond the 13% wage premium of a master's degree. This is fueling an AI ROI crisis, whereby the collective cost of talent and technology (i.e., cloud costs) exceeds the return on AI, especially as POCs continue to stall. In short, businesses are lacking the right skills mix to enable AI that actually generates value.

> This ROI crisis is self-perpetuating, as organizations fight for the top talent to drive AI outcomes but continue to inflate wages further as a result.

Wages for roles with AI skills are:

Higher than comparable roles

Higher than roles requiring a Master's degree

PART 2:

Data Engineering is feeling the pinch

When we examined the sub-verticals of data and AI roles, there was one field feeling constriction most acutely: Data Engineering.

In the US market, the demand for mid-senior talent most drastically outweighs junior by 5:1.

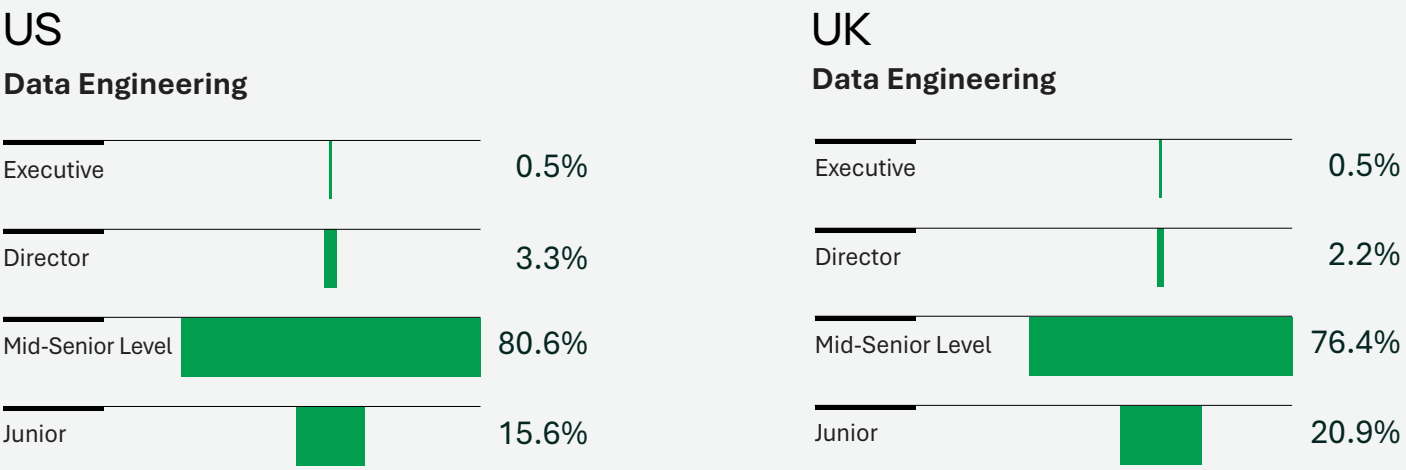

There are a few factors at play.

01

Offshoring

Along with software engineering and cyber security, data engineering is increasingly offshored to reduce headcount and labor costs in the US and UK. The global IT outsourcing market continues to boom, estimated to breach $1billion before 2030, as Brazil, Poland and other markets bring competitive pricing against India's market stronghold.

02

Platform Modernization

Modern data platforms are also reducing foundational and repetitive data engineering tasks with automation. Market-leading platforms like Databricks and Snowflake have a series of features for automating loading new data, pipeline orchestration, and CI/CD integration, as well as Generative AI assistance for code monitoring and updates.

Yet, even with these two mechanisms of accelerating data engineering, most organizations still struggle with their data foundations.

Many organizations have found themselves operating in a complex patchwork of platforms and systems that don't integrate sufficiently to scale AI.

These siloes fragment data, so it becomes outdated, incomplete, and more difficult to use, meaning AI POCs that function perfectly in a controlled environment simply don't have supporting architecture and data to work for real users.

Without sufficient data engineering capability at an entry-level to execute against architecture and migration strategy, organizations are slow to decommission these legacy platforms. Elongated migration periods diminish ROI; more money is sunk into maintaining legacy systems while organizations underutilize committed spend on modern cloud platforms.

In a constrictive data engineering labor market, organizations are further fueling the ROI crisis of their data engineering by adding wage inflation to the mix.

> When over-experienced and costly talent meets slow legacy decommissioning, the total cost of data and AI outweighs its value.

PART 3:

The AI opportunity

The 1.3 million jobs created by the AI investment upswing is a rising the tide for junior AI talent more than any other area of data and AI roles.

The market is still notably constricted, with demand skewed to mid-senior roles by 2-3x. However, this snapshot captures a market as it just begins to surge.

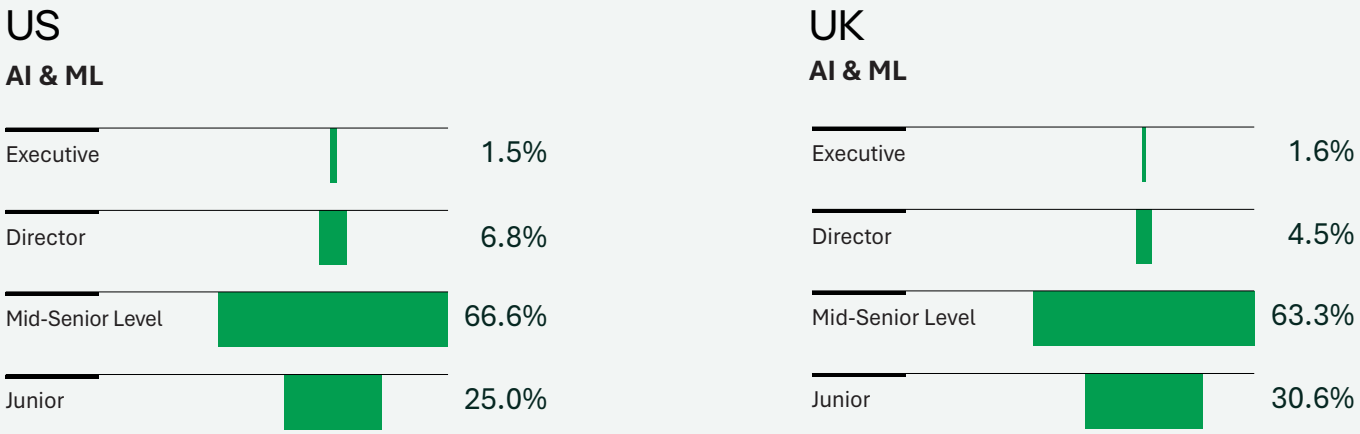

AI Engineering is now the fastest growing role on the market across all industries

According to LinkedIn's Jobs on the Rise 2026, demand for the role has risen 146% over the last year, underscoring the criticality of skills to productionize AI for the two thirds of organizations that are still stuck in the pilot phase of AI.

In the AI value chain, from problem identification to value realization, this role is designed to put technology into action: to take scoped use-cases and develop them for deployment and operationalism by prototyping data, building robustness, and systematizing the monitoring and governance of an AI solution once in production.

The AI Value Chain

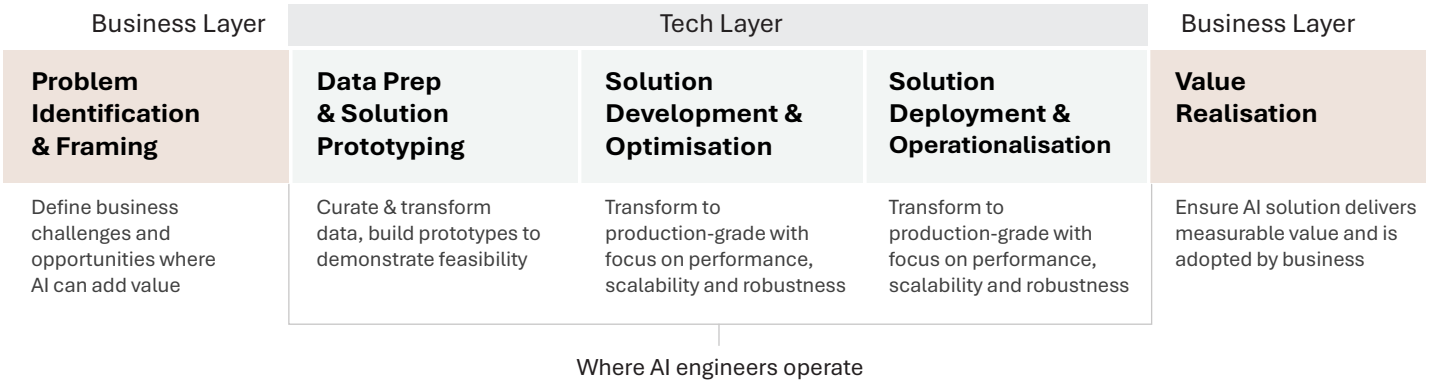

In practice, AI Engineers help create solutions which drive value across different business lines and domains. We classify these types of solutions in four key areas of impact:

Solution	Knowledge discovery and Intelligent automation	Operational forecasting and resource optimization	Risk intelligence and anomaly detection	Customer intelligence and next-best-action systems
Function	Surface enterprise knowledge and automate workflows, combining information retrieval, workflow execution, and governance Eliminate time spent searching for information and automating document-driven workflows	Predictive forecasting for demand and operational volumes to optimize staffing, inventory, maintenance and asset utilization Improve service reliability and performance by aligning resource with predicted demand	AI systems for dynamic pricing, risk detection, or identifying fraud, anomalies or suspicious activity in transactions and operations Improve risk accuracy, leading to better marking and lower risk exposure, reducing manual review	Data-driven systems that model customer behavior to predict churn, segment audiences, personalize experience and recommend next-best-actions Improve targeting and efficacy of marketing and sales motions
Outcome	Increase productivity	Reduce costs	Mitigate risks	Increase revenue

This speed of growth creates a risk for a talent shortage.

> As we've seen in previous tides of technology talent in the last decade, namely in data analytics and software engineering, traditional education models are slow to catch the wave.

Advising on this emerging competition for AI engineering talent, the World Economic Forum stresses the importance of **"verifying real capability"** instead of relying on titles or degree diplomas, where the nuances of AI Engineering may not have been covered by prior work experience or studies. In this still emerging segment, searching for the right talent will require an expansive approach beyond university pools, leveraging other models for training that bring in new talent and enable people with some domain experience to pivot into this role.

The economics – and ROI – at stake

The risk of a constrictive labor market is exacerbated by our current social economics, where the retirement of Gen X-ers will leave a gap of approximately 6 million roles in the next 5-7 years that the smaller Gen Z and Gen Alpha populations cannot fill. Beneath this macro socioeconomic challenge is the reality for each business: low succession rates and a risk to long-term talent development. These gaps can only be filled with costly reliance on consultancies and contractors, meaning little development of inhouse competitive IP.

Game theory offers a clear warning for the economics of this market.

Herding fails; what wins is divergence.

When all agents converge on the same strategy in a competitive environment, that strategy becomes predictable and exploitable. A labor market, where all companies hire the same level of experience, is not in equilibrium.

Any firm that changes tact, by investing counter-cyclically in entry-level talent, gains fitness in the evolutionary sense: lower marginal costs, faster delivery capacity and resilience against talent shocks. Organizations that fixate on 'experience' over 'potential' or 'capability' may risk falling victim to the economic theory of 'sunspots': when sentiment, psychology, or rumours direct market outcomes instead of fact.

The confronting consequence that almost all businesses are feeling is a failure to embrace AI. Board dissatisfaction is rising, as the tension between expectation and reality grows: 81% of Executives believe AI will give them a competitive advantage, yet only 12% have the data architecture mature enough to support it.

As the "total cost of AI" (inclusive of talent and cloud costs) still outweighs the efficiency gains, organizations need to break the cycle. It's less about asking 'why is AI failing?' and instead 'how do we make it work?'. Without the right talent mix, execution is inefficient and ineffective.

So how do we reconcile the immediate needs with long-term strategic vision?

Euan Scott, Chief Delivery Officer at Kubrick, reflected on the consequences of a constrictive labor market for delivery capability.

"In moments of economic uncertainty, junior headcount is the first thing to pause, as businesses put pressure on their experienced workforce to deliver. I've seen these reactive tactics play out over and over again in my career", Scott said. "The future consequences are inevitable: low succession, missing skills, and a greater risk to strategic transformation. However, it is difficult to prioritize these long-term needs when the immediate economics are so startling."

"The key is to make short-term tactics align with strategy; to demonstrate the skill gaps that mean projects with a direct impact on the balance sheet won't get done on time", he continued.

"That's where we focus our talent in Kubrick: creating the workforce who can enable businesses to decommission and migrate those costly legacy platforms which are leaking millions of dollars or utilize underspent cloud permissions to actually productionize AI POCs into value-generating products. It isn't about headcount; it's about the capability to deliver ROI."

Capability meets strategy

Case study: Retail Banking

We helped a major US retail bank build the inhouse capability to scale advanced analytics and AI with proven revenue growth.

The client had a strategic ambition to transform their legacy, retrospective marketing analytics into a real-time decision-making engine. Together with Kubrick, they built a custom talent pipeline to scale the inhouse capability to make it a reality.

> Combining data product management, data engineering, and AI, we aligned their product roadmap and technology stack with specially trained resources to accelerate delivery.

With hyper-personalized segmentation, messaging, and tactics, the client grew their digital channel revenue by 40%, attributed to their all-in-one analytics engine. In the process, they have retained Kubrick consultants as full-time team members, creating the inhouse capability to keep developing and scaling the tool.

Attributed revenue growth

Kubrick resources onboarded

Retention as full-time team members

About Kubrick

We build the capability to embrace AI; with a workforce we create ourselves.

At Kubrick, we to transform lives through data and AI. That means creating better outcomes and opportunities for our clients and our consultants.

Since 2016, we've built a workforce of data and AI professionals that keep ahead of the pace of technology by finding and training brilliant individuals ourselves.

By combining expert solution delivery with team augmentation services, we cover the full spectrum of the AI journey: from idea inception to adoption. We're hands-on, embedded in client teams. We deliver and we don't walk away. We create the capability and the culture that drives AI transformation.

Whether you need to inject capability into a project quickly, sustainably scale your team, or embark on a programmatic workforce strategy, we're here to help.

2500+

Data and AI rofessionals trained

150+

Organizations supported

2x

Industry average female representation in technology roles

Get in touch

Jay Lockwood, Senior Vice President
jaylockwood@kubrickgroup.com

Connect with me on LinkedIn

Appendix

Role	Primary functions	Common skills
Data Engineering	Building, maintaining, and optimizing data pipelines, storage systems, and processing frameworks that move, transform and deliver data at scale	Python, SQL, ETL/ELT pipelines, distributed data processing (Spark/Kafka), cloud data platforms (AWS/Azure/GCP) and Lakehouse/ Medallion-style architecture requirements
Platform Engineering	Designing, deploying, automating and operating cloud infrastructure - including compute, networking, storage, IaC, container orchestration and cloud security	Cloud platform expertise (AWS/Azure/GCP), infrastructure-as-code tools (Terraform/ CloudFormation/Bicep), Kubernetes and container orchestration, CI/CD automation and cloud networking and security skills
Data and AI Governance	Defining and enforcing policies, standards, metadata, access controls, regulatory compliance and data quality frameworks that ensure data is trustworthy, well-managed, secure and ethically used	Data quality, metadata management, data cataloguing, access control, privacy/compliance frameworks (GDPR/CCPA) and enterprise governance tooling such as Collibra or Alation
Data Product Management	Defining, delivering and scaling data products, analytics platforms or data-driven initiatives through product strategy, agile delivery, stakeholder alignment and program or project execution	Agile delivery, product road mapping, stakeholder management, requirements gathering, KPI/metric definition and decision-making using BI dashboards/tooling
AI and Machine Learning	Developing, training, fine-tuning, deploying or operationalizing ML, AI or generative models - as well as focusing on building the platforms, pipelines and infrastructure that support them	ML modelling frameworks (TensorFlow/PyTorch), LLM and generative AI techniques, feature engineering, model deployment/ monitoring pipelines, vector databases, MLOps tooling (MLflow, SageMaker, Airflow)
Applied Decision Intelligence	Generating insights through statistical analysis, predictive modelling, experimentation and advanced analytics to support decision-making	Python/R, statistical modelling, machine learning techniques, forecasting and optimization, experimentation (A/B testing) and data visualization tools such as Power BI, Tableau or Matplotlib.