Machine Learning Engineer (London, Hybrid)

About Us:
We are a **market-leading automotive analytics company** with a long history and deep domain expertise in the UK automotive industry. Our team is on a mission to **drive innovation** in automotive retail through data-driven solutions. We leverage large-scale vehicle and dealership service transaction datasets to build cutting-edge **predictive and prescriptive analytics** tools that help automotive retailers make smarter decisions.

The Role:
We're looking for a **Machine Learning Engineer (2–5 years experience)** to join our London-based team. In this hybrid role, you'll spend part of your week collaborating onsite in our London office and part working remotely. You will play a key role in developing and productionizing machine learning solutions end-to-end – from model development and deployment to ongoing monitoring and improvements. If you're passionate about MLOps, scalable ML systems, and applying data science to real-world industry problems, this role is for you.

What You'll Do

- **Design, Build & Deploy ML Systems:** Develop and maintain end-to-end machine learning pipelines – from data ingestion and feature engineering to model training, deployment, and ongoing management in production.

- **Model Deployment & Lifecycle Management:** Own the full lifecycle of machine learning models in production. You'll deploy models as scalable services (e.g. on AWS), automate retraining, perform A/B testing, and manage versioning/rollbacks as needed to keep models performing at their best.

- **Implement MLOps Best Practices:** Create robust CI/CD pipelines for ML (automated testing, containerization, continuous integration and delivery) to streamline model releases. Establish monitoring and observability for deployed models (tracking performance, data drift, and errors) and set up alerts/dashboarding to ensure system reliability.

- **Collaborate & Innovate:** Work closely with data scientists, data engineers, and product managers to turn prototypes into production-grade solutions. You'll incorporate domain knowledge from our automotive experts to improve model accuracy and business relevance, and help brainstorm new data-driven features and products.

- **Continuous Improvement:** Monitor model outcomes and system metrics in the wild, troubleshoot issues, and proactively improve performance. You'll apply updates or retrain models with fresh data to enhance accuracy, and implement enhancements to make our ML infrastructure more efficient, scalable, and secure.

Who You Are

- **Experienced ML Practitioner:** You have **2–5 years of hands-on experience** building, deploying, and maintaining machine learning models or similar data-driven products in a production environment.

- **AWS Proficient:** Strong experience with **AWS cloud services** (compute, storage, networking) and AWS ML tools. You're comfortable using services like EC2, S3, SageMaker, Lambda, etc., to build and host ML solutions at scale.

- **MLOps & Pipeline Savvy:** Demonstrated experience with **ML deployment pipelines** and MLOps best practices. You have worked with tools for continuous integration/delivery (CI/CD) and configuration-as-code, and you understand how to orchestrate workflows (e.g., using Docker, Kubernetes, Airflow, or similar). You also have hands-on knowledge of monitoring and observability frameworks to ensure model and system reliability (tracking drift, performance, uptime).

- **Strong Coding & ML Skills:** Proficiency in Python (and common ML/Data Science libraries such as pandas, scikit-learn, TensorFlow/PyTorch, etc.). You write clean, efficient code and are familiar with software engineering fundamentals like version control (Git), code reviews, and testing.

- **Analytical Problem-Solver:** Solid understanding of machine learning fundamentals and statistics. You can troubleshoot data and model issues, and optimize algorithms for performance. Experience working with large datasets (e.g., SQL or NoSQL databases, data warehouses) is a plus.

- **Team Player with Domain Interest:** Excellent communication skills and ability to work collaboratively in cross-functional teams. You can explain complex ML concepts to non-technical stakeholders. An interest in the **automotive industry** or experience with retail/transaction data is a bonus – you're excited about applying ML to solve real-world problems in this domain.

- **Education:** BSc/MSc in Computer Science, Data Science, Engineering, or a related field (or equivalent industry experience).

What We Offer

- **Industry Leadership:** Be part of a company with deep domain expertise in the UK automotive sector. You'll work on innovative projects at the intersection of predictive analytics and digital engagement, shaping the future of how businesses connect with their customers.

- **Impact & Growth:** A dynamic, collaborative work environment where your ideas are valued and can lead to immediate impact. You'll have the opportunity to grow your expertise, take on new challenges, and advance your career as our company continues to expand.

- **Flexible Work Environment:** Enjoy a hybrid work schedule – combine the energy and collaboration of our London office with the flexibility of remote work. We believe in work-life balance and adaptability.

- **Competitive Compensation:** A competitive salary and benefits package that rewards your experience and contributions. You'll be joining a forward-thinking team that celebrates success and nurtures talent.

How to Apply

Email careers@tekcor4.com with the job title in the subject line