

Technology Division

Application Management Department

Position: Machine Learning Engineer

Job Description:

- Design, build, and maintain infrastructure for training, deploying, and monitoring ML models at scale.
- Collaborate with Data Scientists to productionize models with reliability and performance.
- Implement MLOps practices (CI/CD, automated retraining, model versioning, monitoring).
- Work within the Microsoft Azure ecosystem (Azure ML, Databricks, Synapse, Data Lake, Fabric).
- Optimize models for scalability and latency in production.
- Build and manage feature stores and model registries.
- Integrate AI models with enterprise systems (APIs, Power BI, Microsoft 365 apps).
- Develop dashboards to track model drift, accuracy, and performance.
- Partner with Data Engineers to ensure robust pipelines.
- Document processes and ensure compliance with governance and privacy standards.

Requirement:

- Proven experience as an ML Engineer, Data Engineer (ML focus), or MLOps Specialist.
- Hands-on expertise with Azure ML, Data Factory, Databricks, Synapse, Data Lake.
- Strong Python skills (PyTorch, TensorFlow, MLflow, scikit-learn).
- Proficiency in SQL and databases (PostgreSQL, Oracle, MSSQL).
- Experience with Pentaho Data Integration (a plus).
- Solid understanding of MLOps pipelines (CI/CD, Docker/Kubernetes, GitHub Actions/Azure DevOps).
- Knowledge of cloud deployment strategies (REST APIs, microservices, serverless).
- Experience with monitoring tools for model performance and data drift.
- Familiarity with data governance and responsible AI frameworks.
- Strong collaboration and problem-solving skills.
- Growth mindset and adaptability in fast-paced environments.