Module 11

Workshop 1: Detecting Data Drift in Production

Icebreaker

Mind the shift

What might cause a model's input data to change between training and production?

Share your answers in the chat.

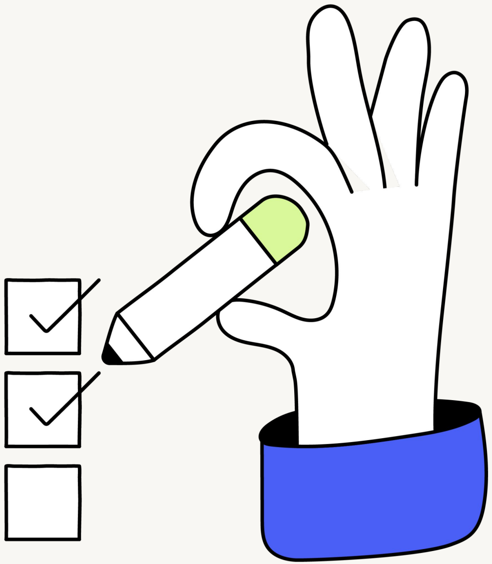

Module 11

Workshop 1: Detecting Data Drift in Production

Agenda

- **Review:** Recap key concepts from async units 1 and 2.
- **Practice:** Building a data drift alert system.
- **Closing:** Key takeaways and next steps.

Learning objectives

- **Reinforce the importance** of continuous monitoring and its role in mitigating deployment risks.
- **Implement a simple Python script** to detect data drift using statistical methods.
- **Simulate a production monitoring scenario** to visualise how changes in input data distribution can be automatically detected.
- **Connect technical monitoring solutions** to model governance and operational resilience.

Async review: Units 1 & 2

Go to My Multiverse, Module 11
Workshop 1

↓

Navigate to the **"Async review"** page

↓

Review the summary of Units 1 & 2

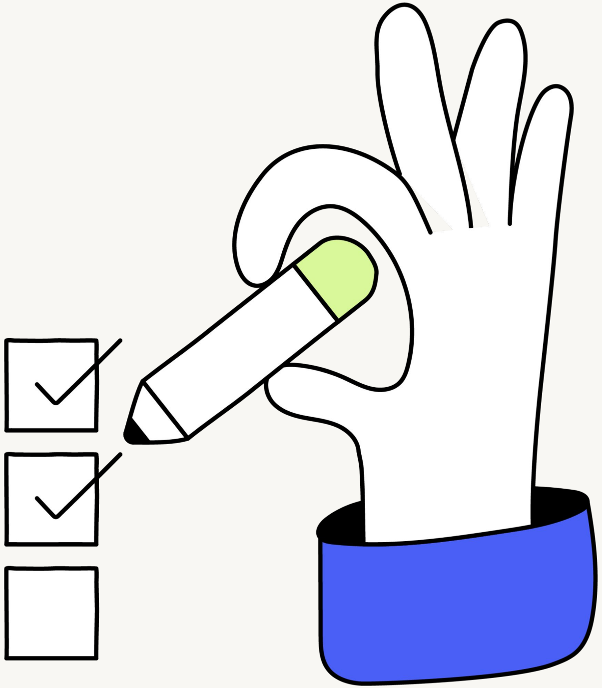

Unit 1: Fundamentals of ML model deployment

- **What ML deployment really means:** Making trained models available for real-world use and understanding how this differs from traditional software deployment.
- **Core MLOps concepts:** Practices that combine ML, DevOps, and automation to ensure models can be reliably deployed, monitored, and maintained.
- **Containerisation:** Using tools like Docker to package ML models and their dependencies for consistent and reproducible deployments.

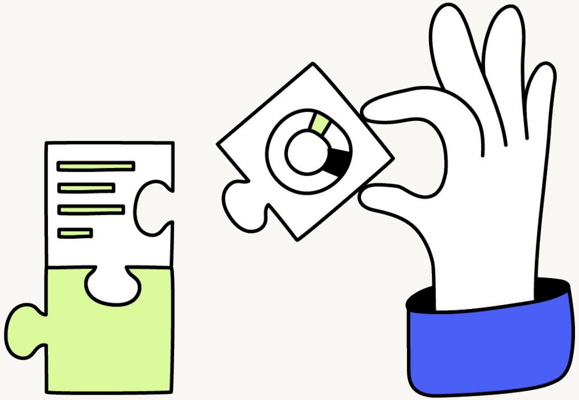

Unit 1: Fundamentals of ML model deployment

- **Deployment workflows:** Designing robust CI/CD pipelines, versioning models and code, and using strategies like blue/green and canary releases to manage change safely.
- **Monitoring and logging:** Setting up systems to track performance, detect drift, log anomalies, and trigger alerts—ensuring models remain accurate and useful over time.
- **Model governance:** Applying traceability, documentation, audit logs, and ethical safeguards to ensure responsible, reproducible, and compliant ML operations.

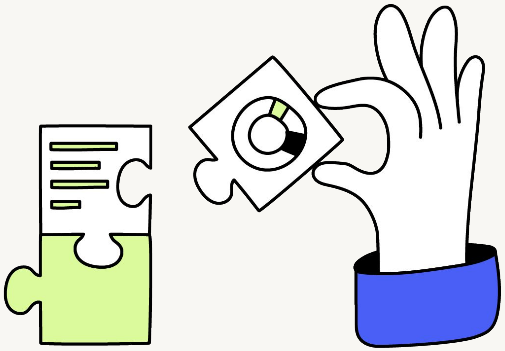

Unit 2: Risk management in model deployment

- **Understanding deployment risks** across technical, operational, business, and ethical dimension including drift, security gaps, and regulatory exposure.
- **Evaluating deployment approaches** such as batch, real-time, and edge, each with distinct risk profiles, especially in automated pipelines.
- **Applying safe transition strategies** like dark launches, A/B tests, and canary or blue/green deployments to minimise disruption and enable early issue detection.

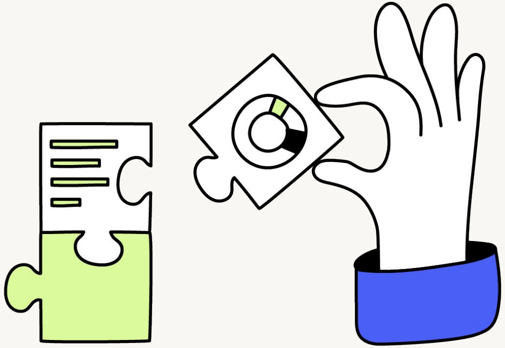

Unit 2: Risk management in model deployment

- **Developing mitigation plans** that include robust monitoring, rollback mechanisms, explainability tools, and compliance safeguards.
- **Analysing real-world scenarios** to identify and address risks related to latency, security, and regulations with targeted, practical solutions.

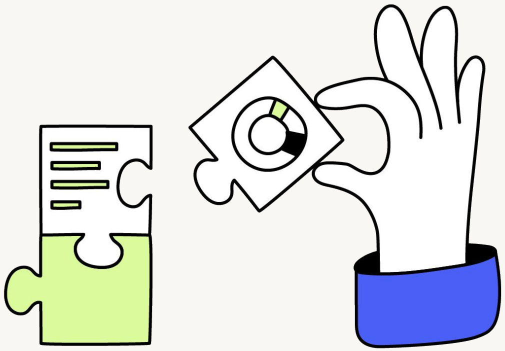

Why monitoring matters in ML production

- Monitoring is not just about system uptime—it's about model health.
- Key metrics include:
  - Model performance (e.g., accuracy, F1 score).
  - Prediction latency and failure rates.
  - Input data characteristics—e.g., missing values, feature distribution shifts.
- Effective monitoring helps detect issues like model drift, silent failures, and unexpected input formats early.

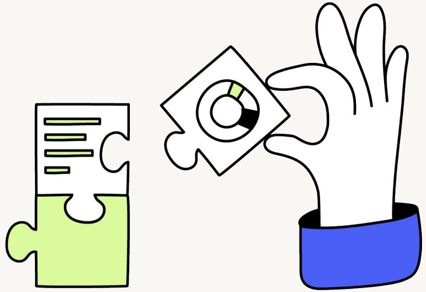

Monitoring as a risk mitigation tool

- Deployed models face **technical, operational, and business risks.**
- **Data drift** is a major technical risk—where real-world input data changes over time, reducing model accuracy.
- This can lead to poor decisions, customer dissatisfaction, or regulatory issues.
- Monitoring is the first line of defense against drift and performance degradation.

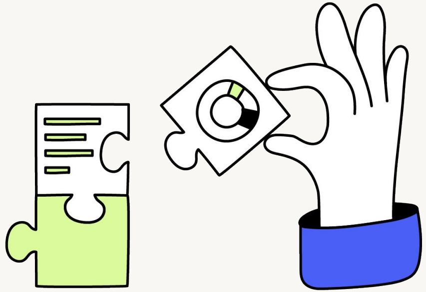

Skills application: Building a data drift alert system

Skills application introduction

- **Objective:** In this skills application, you'll analyse input data distributions over time to detect data drift in a deployed ML system. Starting with a reference dataset, you'll compare it against simulated production data, apply a statistical test (Kolmogorov–Smirnov), and build a basic monitoring loop that triggers alerts when drift is detected.
- **You'll learn how to:**
  - Ensure your deployed models have robust monitoring systems in place to detect data degradation early.
  - How to address data drift issues to ensure your models are stable and are aware of potential risks.

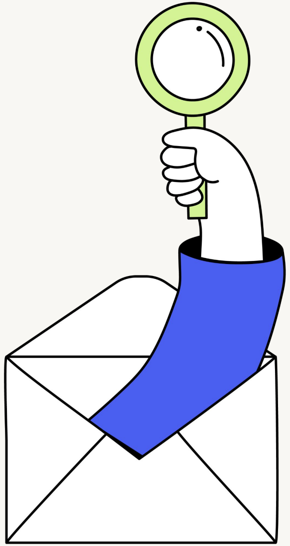

Skills application

- Work in breakout rooms in small groups.
- Use the provided datasets.
- The challenge is designed to be solved individually, but collaboration is encouraged.
- Find all instructions and resources on Ariel.

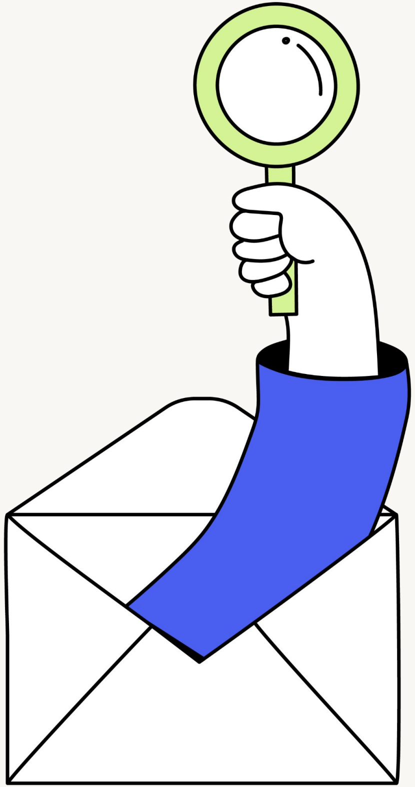

Access the skills application

Go to My Multiverse, Module 11
Workshop 1

Go to the Skills application page and follow the instructions provided

Be prepared to discuss your learnings when we come back together

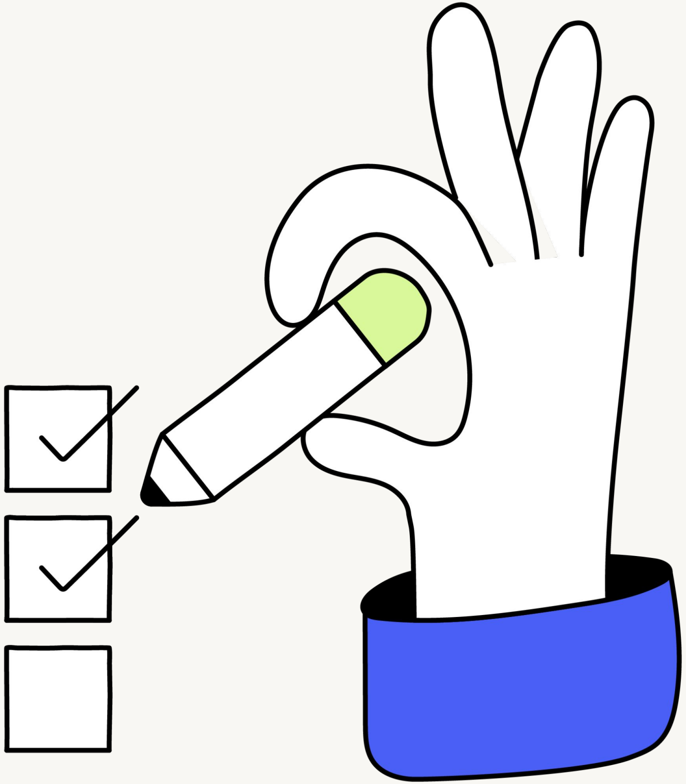

Skills application debrief

Skills application debrief

How does the statistical test used in this activity support real-time monitoring and logging in a deployed ML system?

From a risk management perspective, how could automated drift detection like this help prevent major model failures?

If this monitoring system triggered an alert, what kind of rollback strategy would make sense in response?

Key takeaways

1. **Early detection starts with smart monitoring:**

   By comparing live input data to a reference set using statistical tests, you can spot drift before it impacts performance—keeping your model reliable and responsive.

2. **Drift alerts help keep models deployment-ready:**

   A simple monitoring loop that triggers alerts when data shifts ensures your system stays stable, risk-aware, and ready for real-world use.

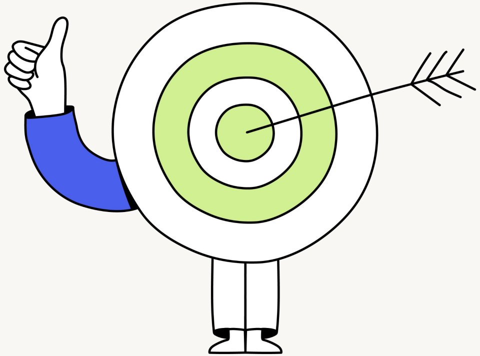

Applying key takeaways to your organisation

Applying key takeaways

1. How could integrating automated drift detection improve the reliability and accountability of the ML systems you work on?

2. Where in your current workflows would adding performance or data monitoring offer early warning signs of model degradation or risk exposure?

Share your plan to apply new skills

Go to My Multiverse, Module 11
Workshop 1

↓

On the Apply new skills to your role page, find the discussion activity

↓

Take 2 minutes to share how you will apply this skill in your role.

↓

Take time this week to read what others have written and get new ideas!

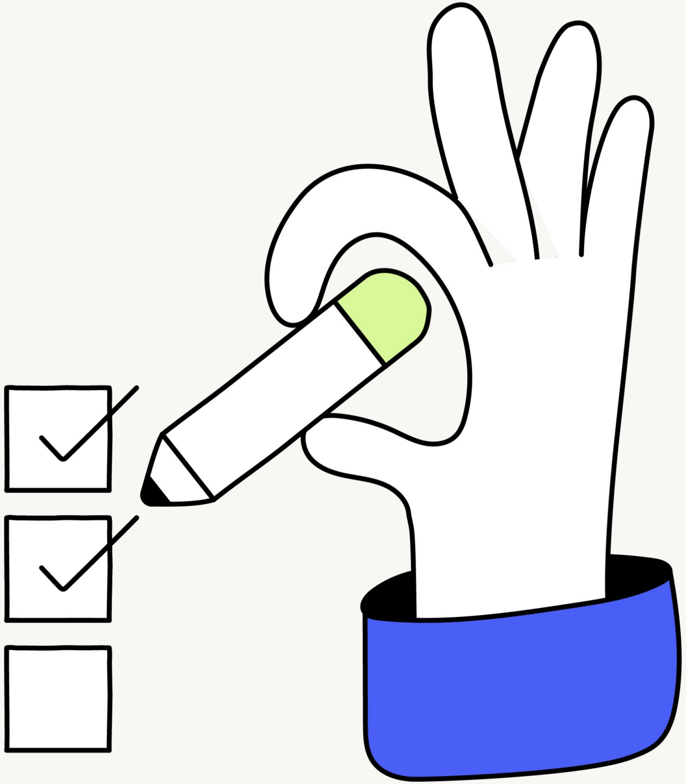

Closing out

- **Next workshop:** Module 11 Workshop 2
- **Have questions?**
  - Chat with Atlas
  - Reach out to your Cohort Coach
  - Sign up for Tutoring
- **Log your OTJ hours**
- **Take the Workshop Survey** on My Multiverse