

Welcome to episode 5 of the AI Practitioner Exam Bites!

Let's start by looking at the review question from the last episode, asking us to identify the most appropriate type of machine learning.

In this case, the answer is C unsupervised learning. This is because there are no predefined labels or categories, and the goal is to discover inherent patterns or groupings in the data - which is a common unsupervised learning technique.

Supervised learning (A) requires labeled data, which is not available in this scenario.

Reinforcement learning (C) is used for decision-making processes in interactive environments, which doesn't apply here.

Federated learning (D) is a machine learning technique where you train algorithms across multiple decentralized devices holding local data samples, which is not relevant to this clustering task.

Today we are going to move into the second Task Statement 1.2 in the Fundamentals of AI and ML domain - *looking at practical use cases of AI*.

In particular, in this episode we are covering the objectives:

Recognize applications where AI/ML can provide value (for example, assist human decision making, solution scalability, automation); and

Determine when AI/ML solutions are not appropriate (for example, cost-benefit analyses, situations when a specific outcome is needed instead of a prediction).

First up, let's talk about some examples of applications where AI/ML shines:

Assisting human decision-making: AI can analyze vast amounts of data and provide insights faster than humans, supporting better-informed decisions in fields like healthcare diagnostics or financial risk assessment.

Enhancing solution scalability: AI systems can handle increasing workloads efficiently - think of chatbots that can simultaneously assist thousands of customers.

Automation: AI excels at automating repetitive tasks - for example, in manufacturing, AI-powered robots can perform precise, repetitive actions tirelessly with consistency.

Pattern recognition: AI can identify complex patterns in data, making it valuable in fraud detection or predictive maintenance.

Now, when might AI/ML not be the best choice?

Cost-benefit considerations: If the cost of implementing and maintaining an AI system outweighs the potential benefits, it's not a good fit. Always do a thorough cost-benefit analysis.

When specific outcomes are needed: If you need a definite, explainable result rather than a prediction or probability, AI might not be suitable. For instance, in critical safety systems where every decision must be 100% certain and traceable.

Limited data availability: AI systems need large amounts of quality data to train effectively. If you don't have sufficient data, AI might not be the right solution.

Ethical concerns: In situations where AI decisions could lead to bias or unfair treatment, human judgment might be more appropriate.

Regulatory restrictions: Some industries have strict regulations that may limit or prohibit the use of AI for certain applications.

Remember, the key is to understand both the potential and limitations of AI/ML. As an AI Practitioner, you need to recognize where AI can add value and where traditional methods might be more suitable.

Let's do a review question before we wrap-up this episode.

A large multinational corporation is looking to improve its operations across various departments. In which of the following scenarios would an AI/ML solution likely provide the MOST value?

- A) Calculating employees' exact salaries based on hours worked and pay rate*
- B) Maintaining an alphabetical list of employee names and contact information*
- C) Assigning parking spaces to employees based on a predefined seniority system*
- D) Detecting potential fraudulent transactions in real-time across millions of daily financial operations***

We'll review this question in our next episode, see you then!