

Welcome to Episode 8 of AI Practitioner Exam Bites.

Let's start by looking at the review question from the last episode about what type of AI would be most suitable to automatically categorize and respond to customer emails based on their content.

The correct answer is D Natural Language Processing - as the AI needs to be able to understand and interpret free-form text written by a human, plus even possibly formulate a response. A bit later in this series you'll see how Natural Language Processing is the backbone of Large Language Models used in generative AI.

The other options are incorrect as they don't relate to the processing of human-written text, but instead focus on other use-cases of AI.

In today's episode we're diving into AWS managed AI/ML services, a key topic for the AWS Certified AI Practitioner exam, looking at the exam objective: *Explain the capabilities of AWS managed AI/ML services (for example, SageMaker, Amazon Transcribe, Amazon Translate, Amazon Comprehend, Amazon Lex, Amazon Polly).*

AWS offers a range of powerful, ready-to-use AI and ML services. Probably the most well-known of these is Amazon SageMaker, AWS's flagship machine learning platform. It covers the entire ML workflow from data preparation, through to model training, evaluation, and deployment.

Do you have to use SageMaker to do AI and ML? No, you could fire-up any computer (in cloud, or on premises) and install open-source machine learning frameworks like TensorFlow, PyTorch, or scikit-learn to build and train your own models. However, SageMaker makes the process much easier through managing all the infrastructure, training workflows, algorithms, etc. for you. In other words, it's a one-stop-shop.

But what if you just want to do something like converting speech to text, or detect objects in images, and don't want to have to train your own model, then deploy and manage that model? Well, AWS offers a bunch of high-level AI services where you can just feed them input (like text or images) and they use their own pre-trained models to provide you with a response. Let's take a look at a few of these:

Need to convert speech to text? That's Amazon Transcribe's specialty. It's great for generating subtitles, transcribing customer service calls, or creating searchable archives of audio and video content. **Here you can see it converting what I am saying now into text.**

Next, Amazon Translate. As the name suggests, this service provides real-time and batch language translation. It supports a wide range of languages, making it ideal for localizing content or enabling multilingual customer support.

Amazon Comprehend is a natural language processing (NLP) service that helps you extract insights from text. It can identify the language, recognize entities, detect key phrases, and analyze sentiment in your text data. Here you can see some text of a product review and it successfully identifies the various entities within the review.

If you're looking to build conversational interfaces using voice and text, Amazon Lex is your go-to. It's the technology behind Alexa and can be used to create chatbots and other conversational interfaces.

Amazon Polly turns text into lifelike speech. It's great for creating voiced applications, accessibility features, or even generating podcast content.

...and I want to do one more, which is one of my favorites...Amazon Rekognition. This provides powerful computer vision capabilities. It can analyze images and videos to detect objects, faces, text, and activities.

Rekognition is great for tasks like facial analysis, content moderation, or identifying products in images. Here is a photo of yours truly with an adorable puppy called Mabel...

These services are designed to be easy to integrate into your applications, often requiring just a few API calls to get started. They're scalable, cost-effective, and continuously improved by AWS, so you're always working with state-of-the-art AI capabilities.

Remember, for the exam, you should understand the basic purpose and capabilities of each of these services. You don't need to know how to implement them in detail, but you should be able to identify which service is appropriate for different AI/ML tasks.

...so let's test your knowledge with a review question!

A company wants to build a multilingual customer support chatbot that can understand customer inquiries in various languages, provide responses in the customer's preferred language, and transcribe voice messages for later analysis. Which combination of AWS services would be most appropriate for this task? Select THREE.

- A) Amazon Lex**
- B) Amazon Polly
- C) Amazon Translate**
- D) Amazon Transcribe**
- E) Amazon SageMaker
- F) Amazon Comprehend

That's all for this bite! Next time, we'll be moving into the final task statement for this first domain - looking at the ML development lifecycle in depth.