


AI Made Easy: Building Your Future with Amazon GenAI Tools

What is Artificial Intelligence

The definition of artificial intelligence has not been clearly defined in the first place. So, while a lot of people are talking about here, there is no agreement on how to define or measure this intelligence. Some of the definition can be given as below:

Definition 1

Science that empowers computers to mimic human intelligence such as decision making, text processing, and visual perception. The ability of a machine to replicate natural intelligence.



Definition 2

The theory and development of computer systems able to perform tasks normally requiring human intelligence, such as visual perception, speech recognition, decision making, and translation between languages etc.

Definition 3

Artificial Intelligence is the ability of a man-made system to replicate human thoughts.

e.g.

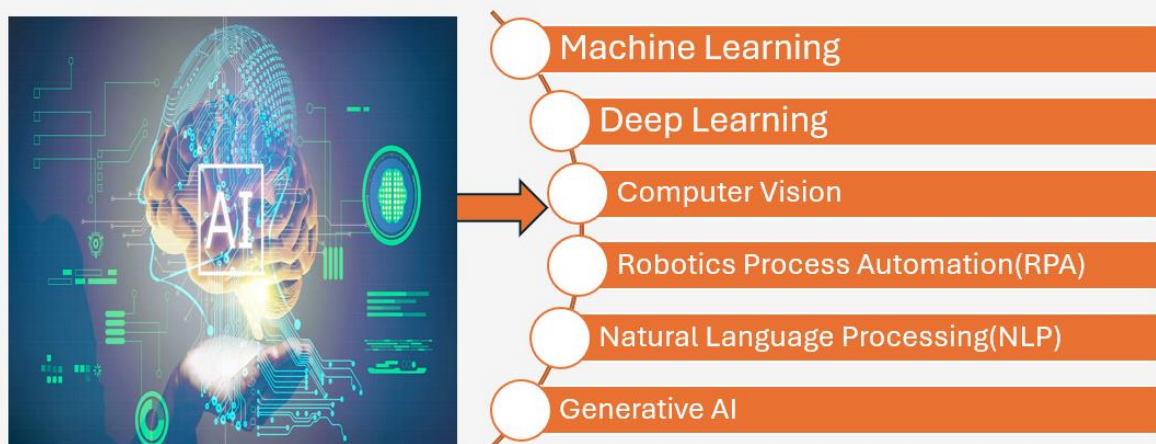
- Predicting whether an airplane arrives early, on-time, or late.
- Determining the likely repair costs for an accident involving a vehicle
- Stock Prices, Weather Condition, Electricity Consumption in a household, Total sales in a store.

Common Artificial Intelligence workloads

The common artificial intelligence (AI) workloads are used to enable machines to perform tasks that typically require human-level intelligence. These workloads are designed to help machines interpret and understand complex data, learn from past experiences, make predictions, and perform tasks autonomously.

By achieving these objectives, AI workloads can improve efficiency, accuracy, and speed of various processes across different industries.

Common AI Workloads



There are many common artificial intelligences (AI) workload that are used in various applications, which includes-

Machine Learning

This is often the foundation for an AI system and is the way we “teach” a computer model to make prediction and draw conclusions from data.

Machine learning (ML) has emerged as one of the most revolutionary technological developments. The use of machine learning (ML) is enabling businesses to accelerate digital transformation and enter the age of automation in the fiercely competitive corporate world. Specific would even



contend that in some industries, like digital payments, fraud detection in banking, or product suggestions, AI/ML is necessary to remain relevant.


Examples:

- Determine whether a customer is likely to switch to a competitor;
- Detect credit card fraud;
- When to perform preventive maintenance on a production robot.

Deep Learning

It is subfield of Machine Learning used for complex algorithm and deep neural net to train the model.

It's a field within machine learning that uses artificial neural networks with many layers, which is why it's called "deep." These layers are like the interconnected neurons in our brains, and they allow deep learning models to identify complex patterns in data.



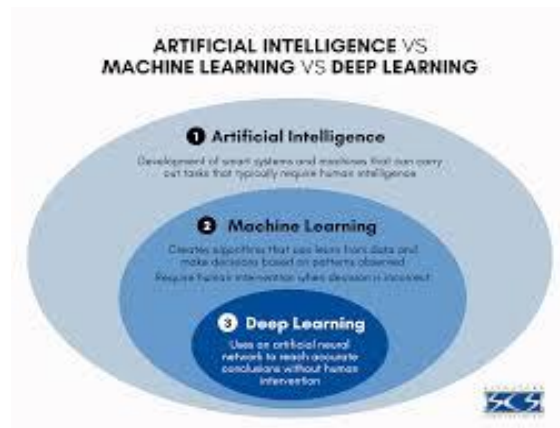
Here's a breakdown of deep learning:

- Inspired by the brain: Deep learning models are based on artificial neural networks, which are designed to mimic the structure and function of the human brain.
- Multiple layers: The "deep" in deep learning refers to the use of many layers of interconnected nodes in the neural network. Each layer learns to recognize different features in the data.
- Learns from data: Deep learning models can learn from vast amounts of data, including images, text, and sound. They can identify patterns and relationships in the data that humans might miss.
- Used for complex tasks: Deep learning is used for a variety of complex tasks, such as image recognition, speech recognition, natural language processing, and self-driving cars.

Relationship Between AI, ML and DL

The given diagram illustrates the relationship between artificial intelligence (AI), machine learning (ML), and deep learning (DL). As per diagram, we can understand that

- **Artificial Intelligence (AI):** This is represented at the top of the diagram. AI is the broadest concept of the three and encompasses both machine learning and deep learning.
- **Machine Learning (ML):** This is a subset of AI and is shown branching out from underneath AI. Machine learning algorithms allow computers to learn and improve without being explicitly programmed.
- **Deep Learning (DL):** This is a subset of machine learning, shown as a smaller branch forking out from underneath machine learning.



Computer Vision

This is used to take images or videos and then provide information such as object detection and Image classification.

Most computer vision solutions are based on machine learning models that can be applied to visual input from cameras, videos, or images. The following describes common applications of computer vision.

- Image classification
- Object detection
- Semantic segmentation
- Faces Services
- Optical Character recognition (OCR)



Natural Language Processing

This is used to analyze text in documents or other text sources. It can interpret speech and generate speech. It also has the ability to translate spoken or written phrases between different languages.

NLP enables you to create software that can:

- Analyze text documents to extract key phrases and recognize entities (such as places, dates, or people).
- Perform *sentiment analysis* to determine how positive or negative the language used in a document is.
- Interpret spoken language, and synthesize speech responses.
- Automatically translate spoken or written phrases between languages



Robotics Process Automation (RPA)

Robotic Process Automation (RPA) is essentially using software robots to automate repetitive tasks on your computer. Imagine a tireless, digital assistant that can mimic what you do on various programs.

Here's how RPA works:

- **Software Robots:** RPA utilizes software programs called bots that can interact with digital systems and applications just like a human would.
- **Repetitive Tasks:** These bots are designed to automate repetitive tasks that follow a defined set of rules. This can include things like data entry, copying and pasting information between applications, generating reports, and moving files.
- **Increased Efficiency:** By automating these tasks, RPA frees up employees to focus on more complex, strategic work. It can also improve accuracy and reduce errors since bots don't get tired or make mistakes.



Overall, RPA is a type of business process automation (BPA) and it deals with automating well-defined procedures rather than using complex learning algorithms.

Generative AI

Generative AI is a powerful technology that enables machines to generate content, such as text, images, and videos, based on patterns learned from existing data. It has applications in various industries, including content creation, product design optimization, and operational efficiency.

Generative artificial intelligence (generative AI or GenAI) is artificial intelligence capable of generating text, images, or other media, using generative models.



Generative AI models learn the patterns and structure of their input training data and then generate new data that has similar characteristics.

In simple terms, Generative AI promises to take a bunch of data and use it to gain hidden insights and put the data to work in various ways. It does this in a way that is very similar to how our brains work. It looks at all the data in all the different places, then says something based on what it saw.

These models can create new data points that have similar statistical properties to the training data they were exposed to. Generative AI models often use techniques like neural networks, particularly generative adversarial networks (GANs), variational autoencoders (VAEs), and autoregressive models. These models learn patterns and structures from a dataset and then generate new samples by either sampling from a learned probability distribution or by learning to produce outputs that are like the training data.

Generative AI has numerous applications as below:

Image Generation: Generating realistic images of people, animals, scenery, etc.

Text Generation: Producing human-like text, such as stories, poems, or even code snippets.

Music Generation: Creating new music tracks or generating melodies.

Video Generation: Generating video sequences, deepfakes, or synthetic training data for computer vision tasks.


Data Augmentation: Generating synthetic data to augment existing datasets for training machine learning models.

Relationship between AI, ML, DL & GenAI

Generative AI falls primarily within the domain of machine learning and, more specifically, within deep learning.

Generative AI techniques involve training models to generate new data samples that resemble a given training dataset. This can include generating images, text, audio, or even videos. Generative models, such as Generative Adversarial Networks (GANs), Variational Autoencoders (VAEs), and autoregressive models, are examples of

deep learning techniques used in generative AI. Generative AI has applications in various domains, including image synthesis, text generation, music composition, and data augmentation.



Different AWS tools for Generative AI



Amazon Web Services (AWS) products that can help you use the power of generative AI are as below:

- ❖ Amazon SageMaker
- ❖ Amazon Bedrock
- ❖ Amazon CodeWhisperer
- ❖ Amazon Q

Amazon SageMaker

A wide range of tools are combined in one fully managed service called Amazon SageMaker to provide high-performance, affordable machine learning (ML) for any use case.

Using tools like notebooks, debuggers, profilers, pipelines, MLOps, and more, you can create, train, and implement machine learning models at large scale with SageMaker – all within a single integrated development environment (IDE).



SageMaker offers streamlined access control, transparency over your machine learning initiatives, and support for governance requirements. Furthermore, you can create your own FMs—vast models trained on enormous datasets—using specially designed tools for FM experimentation, deployment, and fine-tuning.

With only a few clicks, SageMaker provides access to hundreds of pretrained models, including FMs that are made available to the public.

Benefits

Choice of ML tools: IDEs for data scientists and no-code interfaces for business analysts are two examples of ML tools that can be chosen to allow more individuals to use ML for innovation.

Fully managed, scalable infrastructure: With integrated, purpose-built tools and affordable, high-performance infrastructure, you can build your own ML models, including FMs to drive generative AI applications.

Robust and conscientious machine learning workflows: To enhance auditability and transparency, automate and standardize MLOps procedures and governance throughout your company.

Amazon Bedrock

Amazon Bedrock is a fully managed service that offers a choice of high-performing foundation models (FMs) from leading AI companies like AI21 Labs, Anthropic, Cohere, Meta, Stability AI, and Amazon via a single API, along with a broad set of capabilities you need to build generative AI applications with security, privacy, and responsible AI.



Using Amazon Bedrock, you can easily experiment with and evaluate top FMs for your use case, privately customize them with your data using techniques such as fine-tuning and Retrieval Augmented Generation (RAG) and build agents that execute tasks using your enterprise systems and data sources.

Since Amazon Bedrock is serverless, you don't have to manage any infrastructure, and you can securely integrate and deploy generative AI capabilities into your applications using the AWS services you are already familiar with.

Amazon CodeWhisperer

In 2023, Amazon CodeWhisperer's capabilities increased dramatically. I've enjoyed following this AI-powered productivity tool's development over the past year, from a fantastic concept and implementation to a polished, feature-rich product. For that reason, I was thrilled to have the opportunity to speak with many of the tool's developers and learn more about how the application was put together. As explained by Doug Seven, General Manager of Amazon CodeWhisperer:



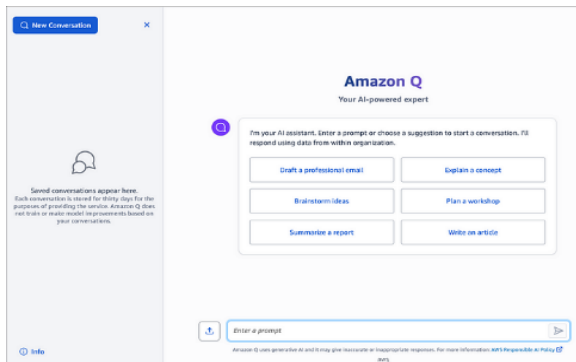
“Amazon CodeWhisperer is an AI-powered developer productivity tool for the IDE and command line. Developers can achieve significant productivity boosts with in-line code suggestions based on the code they are working in and their natural language comments, and now can leverage Amazon Q with

CodeWhisperer, to ask about or get help writing and improving their code and find and remediate security vulnerabilities with the included code scanning capability.”

It seamlessly works with all of your preferred processes, such as Amazon SageMaker Studio, JetBrains, AWS Cloud9, and VS Code. Doug outlined the following significant developments in CodeWhisperer since the platform's announcement of public availability in April 2023: "Since CodeWhisperer's general availability in April, we have been striving to enhance the user experience overall. This involves expanding the number of languages that receive function-block suggestions, adding new smart triggers to more accurately forecast when to offer a code proposal, and updating the model multiple times to enhance the quality of code suggestions. The average acceptance rate across all languages and application cases has increased from about 20% to 35%. With CodeWhisperer now including Amazon Q, developers can now inquire about their code and take advantage of Amazon Q's.

Amazon Q

Amazon Q is a new generative AI-powered personal assistant that can be customized for your company and is meant for work. Through connections to the code, data, enterprise systems, and information repositories of your organization, Amazon Q enables you to engage in dialogues, solve issues, produce content, acquire insights, and take action. Amazon Q gives workers quick access to pertinent information and guidance to help them complete activities more quickly, make decisions and solve problems more quickly, and encourage creativity and innovation at work.



With Amazon Q's user-based plans, you may customize the product's features, cost, and options to suit your needs. Based on your company's current identities, positions, and permissions, Amazon Q can customize its interactions for every single user. Amazon Q user content is never used by AWS to train the underlying models. Put another way, your company's information is kept secret and safe.