TD SYNNEX

Hewlett Packard Enterprise

AMD

HPE-AMD

An Introduction to Artificial Intelligence (AI)

From improving healthcare outcomes to elevating the retail shopping experience, outsmarting financial services fraud, and optimizing production lines, AI, together with HPE ProLiant powered with AMD EPYC™ 4th GEN, is reshaping industries and lifestyles around the world.

"AI is any task performed by a machine that, if a human carried out the same activity, we would say the human had to apply intelligence to accomplish the task."

McCarthy and Minsky (1956)

What is AI?

Artificial Intelligence (AI) refers to the development of computer systems and algorithms that can perform tasks that typically require human intelligence. These tasks include reasoning, learning, problem solving, perception, and natural language understanding.

Today, AI is the main tool for Business Analytics (BA).

But what is Business Analytics? BA is the practice of iterative, methodical exploration of an organization's data, and is used by companies committed to data-driven decision making.

Data is the new oil. Data alone is... useless!

BA uses AI to deliver a role for data: transforming data into knowledge to generate actionable insights. And what is new? New algorithms? No! They are concepts from the 1980s (or older). **Data Availability** and **Computation Capacity** are new! With AI, we can apply new ideas to old concepts to drive unprecedented growth.

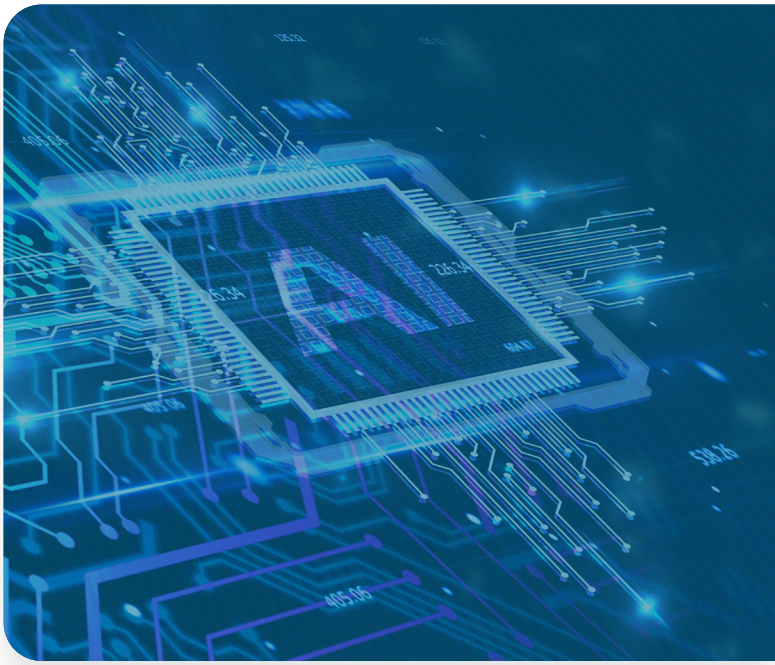

AI can be applied to any system, from small to large AI models. Here are some use cases and HPE-AMD Enterprise AI solutions:

	Industry				Organizational					
	Automotive	FinTech	Healthcare	Manufacturing & Retail	Customer Service	Finance	Sales	IT Dept & Network	HR & Marketing	Operations & Administration
Use-case	• Driving Assistant • Vehicle Cyber-security • Vision Systems • Self-driving cars	• Fraud Detection • Financial Analytics Platform • Insurance Risk Analysis • Expense Management • Credit Lending & Scoring • Regulatory Compliance • Customer Service Automation	• Patient Data Analytics • Personalized Medication & Care • Real-time Triage • Early Diagnosis • Healthcare Chatbot • Device & Drug Comparative Effectiveness	• Quality Control Automation • Supply Chain Optimization • Safety & Compliance Monitoring • Cashierless Checkout • Personalized Recommendations • Inventory Management • Customer Loyalty Program	• Call Analytics & Classification • Call Intent Discovery • Chatbot for Customer Service • Chatbot Analytics • Intelligent Call Routing • Voice Authentication • Survey & Review Analytics	• Invoicing • Billing Reminder	• Sales Forecasting • Lead Generation • Predictive Sales • Sales Call Analytics • Sales Attribution	• Secure Communication • Smart Security • Predictive Intelligence for Security • AI-driven Issue identification • Endpoint Analytics	• Employee Monitoring • Marketing Analytics • Personalized Marketing • Context-Aware Marketing	• Cognitive & Intelligent Automation • Robotic Process Automation • Predictive maintenance • Building Management • Digital Assistant
Data Use-cases	• Analytics Platform • Automated Machine Learning • Conversational Analytics • E-Commerce Analytics				• Image Recognition & Visual Analytics • Data Management & Monitoring • Data Visualization • Data Labelling & Transformation					
HPE-AMD solutions	AI Inferencing HPE ProLiant DL: • 325 Gen 10 Plus V2 + Gen 11* • 345 Gen 10 Plus + Gen 11* • 365 Gen 10 Plus + Gen 11* *powered with AMD EPYC™ CPUs and/or AMD Instinct™ GPUs				• AI Training and tuning • HPC for AI (HPC Infrastructure & services)		HPE ProLiant DL • 385 Gen 10 Plus V2 +Gen 11* HPE Cray • XD 665* • XD 675* *powered with AMD EPYC™ CPUs and/or AMD Instinct™ GPUs			

For further details please visit:

➔ HPE and AMD Server Systems

How has AI developed?

AI can be broadly categorized into several key components and subfields:

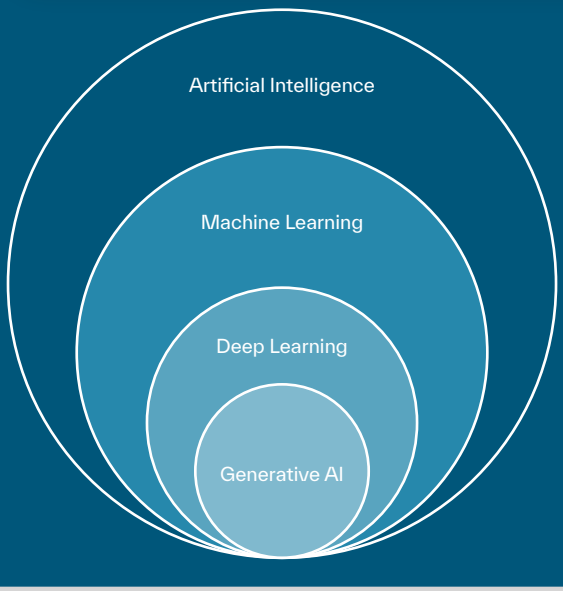

Hewlett Packard Enterprise

AMD

1. Machine Learning (ML)

ML is a subset of AI focused on the development of algorithms that allow computers to learn from, and make predictions or decisions based on, data. ML algorithms detect patterns in large data sets, allowing machines to learn and adapt.

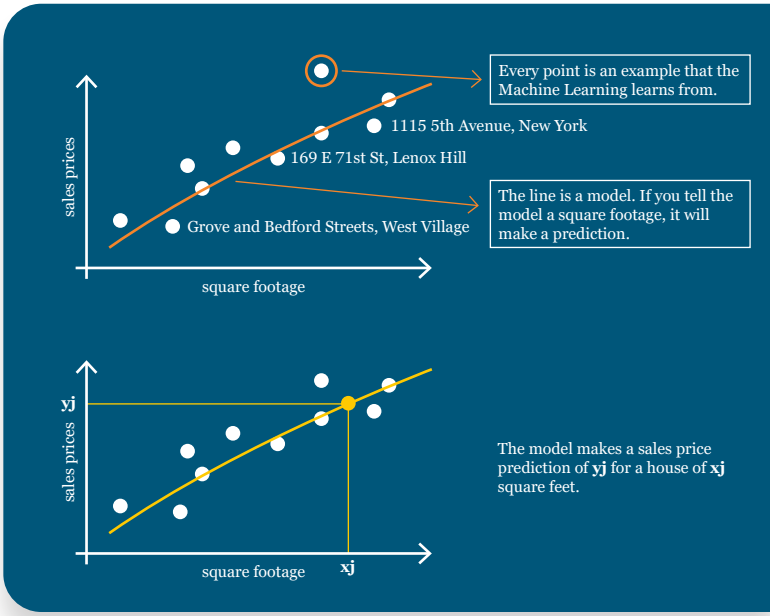

Types of ML:

Supervised Learning:

There are multiple ML problems that require a model to predict an output variable given a number of input variables. The models are trained on labelled data. These problems can be divided into classification and regression problems, depending on the data type of the objective field:

- Regression

  When the objective field is numeric. For these problems, an ML algorithm is used to build a model that predicts a continuous value. That is, given the fields that define a new instance, the model predicts a real number. For example, “the price of a house”, “the number of units sold for a product”, “the potential revenue of a lead”, “the number of hours until next system failure”, etc.

  Regression examples:

  - How much a store will sell?
  - How many products a customer will buy?
  - How much customers will spend?
  - What is your life expectancy?
  - How many customers will churn?

- Classification

  When the objective field is categorical. For these problems, an ML algorithm is used to build a model that predicts a category (label or class) for a new example (instance). That is, it classifies new instances into a given set of categories (or discrete values). For example, true or false, fraud or not fraud, high risk, low risk or medium risk, etc.

  Classification examples:

  - Will this loan be default?
  - What item should I recommend?
  - Does this patient have diabetes?
  - Will this customer churn?

Both classification and regression problems are called ‘supervised’ in the sense that the values of the output variable have either been provided by a human (for example, whether the patient had been diagnosed with diabetes or not) or by a deterministic automated process (for example, customers who did not pay their fees in the last three months are labelled as ‘delinquent’).

*Time Series (TS)

TS is a sequentially indexed representation of your historical data that can be used to forecast future values of numerical properties. It is used to analyze time-based data when historical patterns can explain the future behavior such as sales forecasting, website traffic, production and inventory analysis or weather forecasting, among other use cases.

Unsupervised Learning:

The model is trained on unlabeled data to find hidden patterns.

- Clustering
  There are problems that require datasets to be separated into subsets of instances bearing some similarities. Cluster analysis is an ML task that partitions a dataset and groups together those instances that are similar. It separates a set of instances into a number of groups so that instances in the same group, called 'cluster', are more similar to each other than to those in other groups.

- Association Discovery
  There are problems that require meaningful relationships among two or more values in large datasets across thousands of values, for example, discovering which products are bought together by customers (i.e., market basket analysis), finding interesting web usage patterns, or detecting software intrusion.

- Anomaly Detection (AD)
  There are problems that require identifying the instances within a dataset that do not conform to a regular pattern, for example, detecting any kind of fraud, or discovering errors in your operations. The AD is an unsupervised learning method that is capable of detecting anomalous instances in unlabeled datasets. This means that you do not need to collect a training dataset knowing in advance which instances are anomalous and which are normal. The algorithm can find suspicious patterns in your data given a set of input fields.

- Topic Modelling (TM)
  TM is an unsupervised learning task that helps you find the topics underlying a collection of documents. In other words, it is a form of text mining to identify the hidden thematic structure in a corpus. In order to understand topic models, it is important to consider that documents are generated using one or a combination of topics, and each topic is a group of co-occurring terms with different probabilities.

- Principal Component Analysis (PCA)
  Many datasets contain an extremely large number of fields, or highly correlated fields, resulting in suboptimal model performance. PCA is one technique that can be used to transform a dataset in order to yield uncorrelated variables, or as a first step in dimensionality reduction.

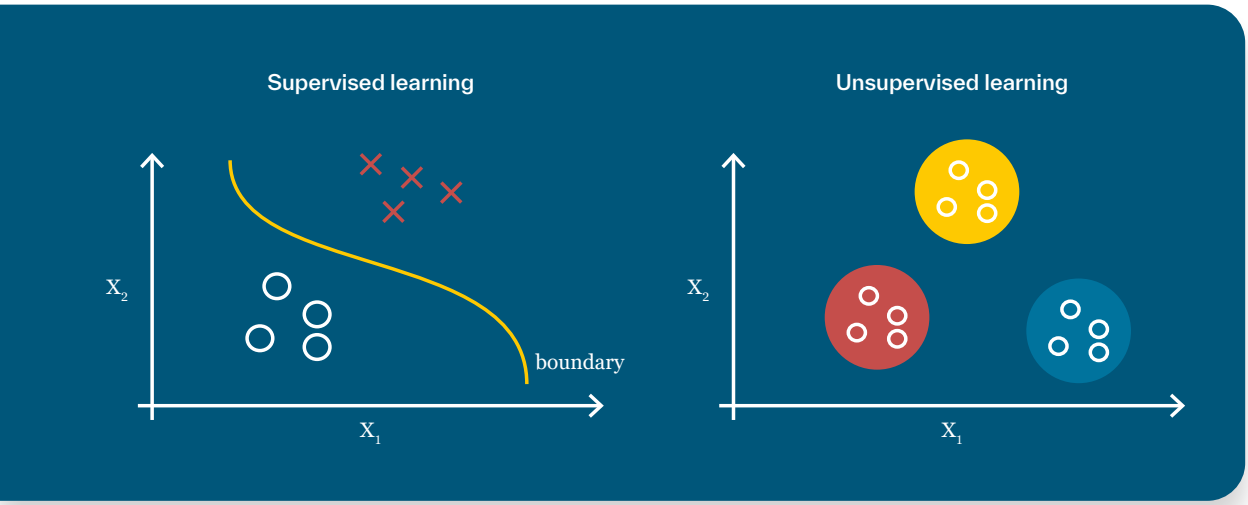

2. Deep Learning (DL)

DL is a subset of AI that teaches computers to learn by example, much like humans do. Inspired by the Brain, DL models are built to work similarly to the human brain, using structures called neural networks.

These networks consist of many layers of interconnected neurons that process information. Just as you learn to recognize objects by seeing many examples, deep learning models learn by analyzing large amounts of data. For instance, a model might be trained to recognize cats by being shown thousands of cat images.

The 'deep' in deep learning refers to the many layers in these neural networks. Each layer extracts more detailed features from the data. Early layers might detect simple shapes, while later layers identify complex patterns and objects.

DL is particularly good at understanding complex and high-dimensional data, such as images, speech, and text.

Unlike traditional machine learning, which often requires manual feature selection, DL automatically discovers the best features for a given task, DL automatically extracts features from large amounts of data.

3. Generative Artificial Intelligence (GEN AI)

GEN AI refers to a class of AI systems designed to generate new content, such as text, images, music, and more, by learning from existing data. These systems use advanced ML techniques, particularly neural networks, to produce outputs that mimic human-like creativity and reasoning.

CONCLUSION

All these AI components often interact and integrate to create sophisticated AI systems.

For example, an autonomous vehicle uses computer vision to navigate, ML to improve its decision making over time, and robotics to control its movements. Similarly, virtual assistants like Siri or Alexa combine natural language processing (NLP), speech recognition, and ML to understand and respond to user commands.

Overall, AI is a multifaceted and rapidly evolving field, with each part contributing to the development of intelligent systems capable of performing complex tasks in any kind of situation, business or organization.

Achieve the highest performance of your AI workloads with HPE ProLiant servers powered by AMD.

Now is not the time to fear the new. It is time to understand it and prepare to embrace it.

Hewlett Packard Enterprise

AMD