Introducing AI

AI OVERVIEW · ML CONCEPTS · DEEP LEARNING & GENERATIVE AI · AWS AI SERVICES

STUDY & REVIEW

01 MUST KNOW

if you remember five things, remember these

01 **The nesting: AI ⊃ ML ⊃ deep learning ⊃ generative AI.** AI (a field since the 1950s) makes machines perform tasks that typically require human intelligence; ML is the broad set of AI techniques at its center; deep learning is ML on deeply layered neural networks; generative AI is deep learning built on foundation models.

02 **Train, then infer:** ML applies an algorithm to existing data to **train a model** that learns to recognize patterns; the model's output on **new** input data is an **inference**. Deployed models must be monitored — accuracy decays when real-world patterns change, and retraining is the fix.

03 **Specific vs. generalized: traditional ML** trains on a well-defined dataset to make one specific prediction (is this transaction fraud?). **Generative AI** uses **foundation models (FMs)** pretrained on vast data — they generate *new content* and answer generalized requests, even tasks never defined in advance.

04 **Good data is the fuel:** the quality and quantity of training data directly set model performance. Check **volume, quality, bias/representation, and timeliness** — plus PII protection, security, and governance. Data preparation typically consumes a very large share of an ML project's timeline.

05 **Responsible AI:** AI's insights and automation come with risks — bias, error, opacity, unclear accountability. Systems must be **accurate/fair/safe/secure**, **explainable/transparent**, and **governed/controlled** — and both builders *and* users own that duty.

02 KEY TERMS

TERM	DEFINITION
AI	Field, evolving since the 1950s, of making machines perform tasks that typically require human intelligence
Machine learning (ML)	AI techniques that use existing data to train a model to recognize patterns and predict on new input
ML model	The end result of training — an algorithm applied to a dataset that has learned to make predictions or decisions
Inference	The output a trained model produces by applying what it learned to new input data
Deep learning	ML that uses deeply layered neural networks trained on very large amounts of data; designed for a specific domain
Neural network	Model architecture of an input layer, one or more hidden layers (the complex processing), and an output layer
Generative AI	A type of deep learning that generates new content in response to requests and responds across broad domains
Foundation model (FM)	Very large deep learning model pretrained on massive data; handles generalized requests beyond its training
Labeled data	Training data that identifies the target of the prediction — an explicit tag or a target column
Structured data	Rows and columns with a well-defined schema (relational tables, spreadsheets); simple to query, not flexible
Semistructured data	Self-describing elements and attributes marked by tags (CSV, JSON, XML)
Unstructured data	No predefined structure (images, movies, clickstream); harder to analyze, most flexible with the right tools

03 ONE FIELD, NESTED

each layer is a type of the one before it

01 **AI** — machines doing human-intelligence tasks → 02 **ML** — data trains a model to predict → 03 **Deep learning** — deeply layered neural networks → 04 **Generative AI** — FMs that create new content

04 DATA ANALYTICS OR ML?

pick the approach that fits the problem

TRADITIONAL DATA ANALYTICS

when the problem is rule-shaped

- Business rules can be hardcoded
- Variables are limited
- Statistical modeling can solve it
- Analytics answers quickly enough

MACHINE LEARNING

only if analytics isn't a good fit — then check:

- Enough relevant, high-quality training data
- ML meets required accuracy, transparency, speed
- Skills and resources to build **and sustain** it
- Cost of ML matches the impact of the problem

05 THE THREE ML PARADIGMS

match the paradigm to the data you have

PARADIGM	TRAINING DATA	PROBLEM TYPES	APPLICATION
Supervised	Labeled	Classification, regression	Decide if a patient has a disease from diagnostic test results
Unsupervised	Unlabeled	Clustering, anomaly detection	Analyze spending patterns to spot possible fraud
Reinforcement	Rewards/penalties from the environment	Robotics, inventory control	Teach a robot to navigate a room

06 TRADITIONAL ML VS. DEEP LEARNING VS. GENERATIVE AI

know which trains on what, and how broad the output is

APPROACH	TRAINS ON	SCOPE OF OUTPUT	EXAMPLE USES
Traditional ML	A well-defined, specific dataset	One specific prediction per problem	Fraud yes/no, ticket categorization
Deep learning	Very large volumes of varied, unstructured data	Interrelated complex tasks **within one domain**	Self-driving cars, weather prediction, recommendations
Generative AI	Massive data (pretrained FMs)	**New content**; generalized requests across broad domains	Chat apps, coding assistants, image generators

07 AWS AI/ML SERVICES AT A GLANCE

one-line job of each — exam gold

SERVICE	WHAT IT DOES	REMEMBER
EC2 AI instances	Compute purpose-built for training performance and, separately, for serving inferences	Infrastructure layer
SageMaker AI	Integrated environment to build, train, and deploy models; managed infrastructure tools + workflows	Full ML lifecycle
Amazon Bedrock	Access to popular foundation models through an API	Gen AI apps, no infrastructure
Amazon Q	Generative AI assistant (covered in later modules)	Pre-trained service
Amazon Transcribe	Audio files → text output	Speech-to-text
Amazon Translate	Text from one language → another	Translation
Amazon Polly	Text → synthesized speech	Text-to-speech

SERVICE	WHAT IT DOES	REMEMBER
Amazon Comprehend	Derives insights from text documents	Text analysis

08 EXAM TRAPS

where the knowledge check gets you

- × “Computer vision and smart devices are alternatives to ML.” — They are AI **specializations that incorporate ML techniques**; ML sits at the center of AI.
- × “Generative AI sits outside machine learning.” — Generative AI is a type of **deep learning**, which is itself a subset of ML.
- × “More data always fixes the model.” — Biased or unrepresentative data yields biased models; **quality, representation, and timeliness** matter as much as volume.
- × “A model that tests well stays accurate forever.” — Production patterns drift (fraudsters change tactics); **monitor inferences and retrain**.
- × “Every data problem calls for ML.” — With limited variables and hardcodable rules, **traditional data analytics** is the better fit; even then, weigh data availability, accuracy/transparency/speed requirements, skills, and cost vs. impact.
- × “Unlabeled data can't train a model.” — **Unsupervised learning** trains on unlabeled data (clustering, anomaly detection); it generally needs more data and more complex models.
- × “Deep learning models are general-purpose.” — Deep learning models are designed for a **specific domain**; broad, generalized responses are the signature of **foundation models**.
- × “Generative AI is always the best modern answer.” — For categorizing existing content (like support tickets), **traditional ML** is more suitable and more cost-effective.

09 SELF-QUIZ

answers are upside-down below

Q1 Arrange from broadest to most specific: deep learning, AI, generative AI, ML.

Q2 What term names the output a deployed model produces when it applies its training to new input data?

Q3 Which ML paradigm requires labeled training data, and what are its two classic problem types?

Q4 Which ML paradigm learns from rewards and penalties derived from an environment?

Q5 Generative AI uses very large deep learning models pretrained on massive amounts of data — what are they called?

Q6 Name the three general types of data an ML solution might use.

Q7 Which AWS service provides an integrated environment for building, training, and deploying ML models?

Q8 Which AWS service lets developers reach a variety of popular foundation models through an API?

Q9 Which three AWS AI services, in order, turn an English audio file into Spanish speech?

Q10 Sample exam: a customer service department wants to automatically categorize incoming support tickets — traditional software, traditional ML, generative AI, or manual sorting?

ANSWER KEY (rotate the page)

Q1 AI → ML → deep learning → generative AI **Q2** an inference **Q3** supervised learning; classification and regression **Q4** reinforcement learning **Q5** foundation models (FMs) **Q6** structured, semistructured, unstructured **Q7** Amazon SageMaker AI **Q8** Amazon Bedrock **Q9** Transcribe → Translate → Polly **Q10** traditional ML — it learns classification from historical ticket data