

Lesson 12 — AI & ML Foundations

The world before LLMs — what machine learning is, what models do, training vs inference.

IT Landscape

Module 4

The words “AI” and “ML”

“AI is the dream. ML is the technique that’s actually working.”

- **AI** (Artificial Intelligence) — the broad goal: machines that act intelligently.
 - **ML** (Machine Learning) — one family of techniques where the machine *learns from data* instead of being explicitly programmed.
 - Most of what news outlets call “AI” in 2026 is actually ML.
-

Programming vs learning

Programming

A human writes the rules, step by step.

"If temperature > 30, send alert."

Machine learning

A human gives the machine thousands of past examples (temperature + did-something-bad-happen). The machine works out the rule itself.

ML shines where rules are too complex for humans to write.

What is a “model”?

“A model is a math function shaped by data.”

- **Input** goes in — a photo, a sentence, a row of features.
 - **Output** comes out — a label, a number, a probability.
 - The “shape” of that function — the **parameters** — was learned from training data.
 - Same architecture, different data → **different model**.
-

Training vs inference — the two phases

Training

Slow, expensive, done once (or periodically).

The model sees millions of labelled examples and adjusts its parameters until it predicts well.

Inference

Fast, cheap, done every time a user asks.

The trained model takes new input and produces output.

“Training is months of GPU bills; inference is milliseconds per request.”

Supervised, unsupervised, reinforcement

- **Supervised** — given input + correct answer, learn the mapping (spam classifier, photo tagger, house-price predictor).
 - **Unsupervised** — given input only, find structure (clustering customers, detecting anomalies).
 - **Reinforcement** — an agent acts in an environment, gets rewards, learns a strategy (game-playing AI, robotics, ad bidding).
-

The classical ML toolkit

Techniques that ran the field for ~30 years — and still ship today:

- **Linear regression** — predict a number.
- **Logistic regression** — predict a category.
- **Decision trees, random forests, gradient-boosted trees** — still the workhorses for tabular data.
- **k-means** — clustering.
- **SVMs** — classification.

In 2026, the boring tabular-data winner is still gradient-boosted trees, not deep learning.

Neural networks — what they actually are

- A stack of math operations (multiplications + a nonlinear squash) arranged in **layers**.
 - Each layer transforms its input; the final layer produces the answer.
 - Training nudges the millions of internal numbers (**weights**) so outputs match the training data.
 - The “deep” in deep learning just means **many layers**.
-

Why deep learning took off (~2012)

Three things had to line up:

- **Big labelled datasets** — ImageNet, the web.
- **GPUs** — parallel math, originally for video games, repurposed for neural nets.
- **A few algorithmic tricks** — better optimizers, better activations, better regularization.

Once they did, deep learning crushed classical techniques on **images**, then **speech**, then **language**.

Frameworks people use (as of 2026)

Framework	Owner	Where it fits
PyTorch	Meta	Research + production winner
TensorFlow / Keras	Google	Still huge, especially in industry
JAX	Google	Research darling
scikit-learn	Community	The classical-ML standard
Hugging Face Transformers	Hugging Face	The LLM bridge

Most ML engineers use 2 or 3 of these daily.

Where ML actually ships (none of these are LLMs)

- **Fraud detection** — banks, every credit-card swipe.
- **Recommender systems** — Netflix, Spotify, TikTok feed.
- **Search ranking** — Google, Amazon.
- **Medical imaging** — cancer detection in scans.
- **Speech recognition** — Siri, Alexa, captions.
- **Forecasting** — electricity grid, ride-hailing demand.
- **Autonomous-driving perception.**

The world has been “running on ML” for a decade.

Recap

- **ML = learning from data**, not from rules.
- The workflow is **train once, infer many times**.
- **Deep learning** is one corner of ML — very good at perception.

Coming up: Lesson 13 — what changed in 2022 when LLMs arrived, and where agents are taking us next.

