

SAMIR GENAI ACADEMY

PROFESSIONAL STUDY GUIDE SERIES

Phase 01 · GenAI Foundations

01

AI, Machine Learning and Generative AI

How each one grew out of the last — and where the LLM sits

Module 01 of 25 · **1 hour**

Instructor: **Samir Kumar** · Edition 2026

STUDY GUIDE

01 Overview

How each one grew out of the last — and where the LLM sits. This guide gives you the core theory, the tools that matter, a hands-on build, working code, and the pitfalls to avoid — everything you need to learn the topic to a professional standard and apply it in real projects.

1 hour

LESSON LENGTH

5 + 4

CONCEPTS + BUILD
STEPS

5

FRAMEWORKS & TOOLS

01/25

PHASE 01

02 Learning Objectives

- ▶ Tell artificial intelligence, machine learning, deep learning and generative AI apart, and say how they nest.
- ▶ Explain what “the model” is, and why it is built from data rather than written by hand.
- ▶ Describe a neural network from one neuron upwards, and say what “deep” means.
- ▶ Decide whether a given task needs prediction or generation.

03 Core Concepts

AI is a discipline, not a product

Artificial intelligence is the field of building machines that do work we normally need a brain for: learning, problem solving and decision making. The idea dates from 1950, and for its first thirty years it meant a person writing rules by hand and a machine following them. That approach works for a purchase-approval threshold and collapses on a photograph, because nobody can write down every rule for reading one.

Machine learning: show examples, not rules

Machine learning is a subfield of AI. Instead of writing the rules, you supply data — a million emails already marked spam or not — and the algorithm finds what the spam ones have in common. It produces the rule itself. Your spam folder, recommendation systems and card-fraud alerts are all machine learning, each trained for one specific task.

The model is the maths, and the data makes it

The maths that turns an input into an output is called the model: something goes in, a result comes out. In practice it is a large structure of parameters rather than one equation. Crucially, nobody writes it — it is derived from the training data, so different data produces a different model with different answers. When a model is wrong, look at the data before the code.

Neural networks and deep learning

One artificial neuron takes several numbers, multiplies each by a weight, adds them up and passes the result on; training means adjusting those weights. Arrange neurons into layers, stack the layers, and you have a neural network: input layer, hidden layers, output layer. “Deep” simply means many layers. Each layer learns something more abstract than the last — edges, then shapes, then faces — which is what finally made images, speech and video tractable. Deep learning split into NLP for text and speech, and computer vision for images and video.

Predictive versus generative

Everything up to this point analysed something and predicted an outcome: spam or not, churn or not, cat or not — one answer chosen from a set that already existed. Generative AI produces new content instead: text, images, audio, video, code. The test is one question: does it pick an answer, or write one? A generative model does not store its training examples; it learns their underlying patterns and structures, which is why it needs an enormous amount of material. A large language model is generative AI working in text, and sits at the centre of the whole picture.

04 Visual Diagram




Figure 1 — The four-layer AI stack. Value increasingly accrues at the upper orchestration and application layers, where this course focuses.

05 Key Frameworks & Tools

Anthropic Claude API

OpenAI API

Google Vertex AI

Hugging Face Hub

scikit-learn (classical ML)

06 Hands-On Build

✂ Sort the tools you use, then make one API call

1. List twelve AI products you or your team actually use.
2. Place each in the smallest circle that contains it: rules-based AI, machine learning, deep learning or generative AI.
3. For each one, say whether it predicts an answer or generates new content.
4. Then obtain an API key, store it in an environment variable, and make one authenticated call to a model.

07 Code Walkthrough

Your first frontier-model API call (Anthropic Claude)

PYTHON

```
import os
from anthropic import Anthropic

client = Anthropic(api_key=os.environ["ANTHROPIC_API_KEY"])

resp = client.messages.create(
    model="claude-opus-5",
    max_tokens=512,
    messages=[{
        "role": "user",
        "content": "Summarise the difference between machine learning "
                  "and generative AI in 3 bullets.",
    }],
)
print(resp.content[0].text)
```

08 Lab Assistant

Lab 1 • Make your first frontier-model call

Prerequisites: Python 3.10+ · pip install anthropic · An API key in ANTHROPIC_API_KEY

1. Export your key: `export ANTHROPIC_API_KEY=sk-...`

Expected: `echo $ANTHROPIC_API_KEY` prints your key

2. Run the Module-1 code snippet

Expected: A 3-bullet summary prints to the console

3. Change the model id and re-run

Expected: Output style/latency differs — note it in your rubric

Troubleshooting

- 401 Unauthorized → key not exported in this shell.
- `ModuleNotFoundError` → run `pip install anthropic` in the active venv.

09 Pitfalls & Key Takeaways

⚠ COMMON PITFALLS

- Treating AI, machine learning and generative AI as competing technologies rather than nested ones.
- Blaming the code when a model behaves badly — the training data is the usual cause.
- Assuming a generative model can recite its sources; it learned patterns, not copies.
- Hard-coding API keys in source — always read them from the environment.

✓ KEY TAKEAWAYS

- AI contains machine learning, which contains deep learning, which contains generative AI, at whose centre sits the LLM.
- The model is the maths, and it is built from the data — different data, different model.
- Deep learning is neural networks with many layers, and it is what made messy data work.
- Predictive AI picks an answer that already exists; generative AI writes a new one.

 **Companion video:** Watch the module 01 lesson in your Student Dashboard, then use this guide to revise, copy the code, and complete the build. Download both for offline study.

10 Further Resources

Anthropic API documentation • Google Cloud: introduction to generative AI • “Attention Is All You Need” (2017) — the transformer paper, covered in Module 2