

# Machine Learning Foundations

Course Brochure



# MACHINE LEARNING FOUNDATIONS

Pilot-ready machine learning capability—structured, practical, and built for real delivery.

Machine learning education is everywhere. **Workplace-ready, deployable capability is not.** Most learners aren't short on information — they're short on a coherent pathway. Fragmented tutorials create isolated skills, thin project narratives, and uncertainty when it's time to justify choices to stakeholders.

**Machine Learning Foundations** is a **practitioner-led, outcomes-driven** course built to close that gap. You'll learn to translate a real workplace problem into a clearly scoped ML initiative, define defensible success criteria, and assemble a credible first deployment pathway — reaching Pilot-Level readiness with a project package you can stand behind.

---

## Who this course is for

- IT practitioners moving into applied ML who need a structured pathway to prove capability quickly with a workplace-credible deliverable.
- Software engineers / AI developers delivering an ML pilot, who want a scoping + evaluation + deployment-pathway framework to avoid building the wrong thing and get stakeholder approval faster.

## What you'll gain

- Build real ML models end-to-end using Python + scikit-learn.
- Evaluate models, interpret results, and explain outcomes to non-technical stakeholders.
- Turn a workplace problem into a guided, well-formulated ML capstone project — with clear success metrics, risk trade-offs, and a credible pilot pathway.

## What you'll walk away with

- ML Project Blueprint (use case → target variable → required data → success metrics → key risks)
- Evaluation Design (baseline + metric justification + error trade-offs + what conclusions are valid)
- Pilot-Level Deployment Pathway Outline

## What this course is

You'll follow a structured, build-led learning path: start with core ML concepts and the Python toolchain, then move quickly into hands-on model development—data preparation, regression/classification, and evaluation—so you can produce clear, explainable results. You'll then extend into modern applications (selected advanced/unsupervised methods, an end-to-end project), and gain guided exposure to deep learning, LLMs, and real-world deployment considerations through demonstrations.

## CURRICULUM-AT-A-GLANCE

### Block 1. Machine Learning Introduction

A crisp orientation to the ML landscape: data analysis foundations, core AI concepts, and how machine learning fits into an end-to-end lifecycle—grounded in the Python toolchain used in practice.

- Introduction to Data Analysis and Data Science
- Introduction to AI: history and basic concepts
- ML definition, concepts, and lifecycle
- Python libraries and tools for data science and ML

---

### Block 2. Hands-on Machine Learning

Build capability through application: prepare data, train supervised models, evaluate performance with the right techniques, and complete a full end-to-end project with exposure to high-signal use cases.

- Data preprocessing and visualization
- Supervised learning: regression and classification
- Evaluation of ML models
- Cross-validation and hyper-parameter tuning
- Advanced predictive models
- Unsupervised models
- Full, end-to-end ML project
- Salient applications: recommender systems

---

### Block 3. Advanced Machine Learning

A guided view of modern ML directions—deep learning, computer vision, and generative/LLM applications—plus what it means to deploy models in real settings, supported by demos.

- Deep learning foundations
- Computer vision techniques and tasks
- Generative models and LLMs; using pre-trained LLMs
- Deploying AI/ML models in real-world + example demos

## The Problem We Solve

### The Stakeholder Moment

The real test rarely happens while you're learning. It happens at the decision table — when an ML initiative is proposed, reviewed, or challenged under stakeholder scrutiny.

In that moment, you're expected to do more than “show a model.” You need to:

- Frame the problem and justify why ML is (or isn't) the right approach
- Define measurable success and own the trade-offs behind it
- Surface risks early — privacy, bias, error costs — before they become objections
- Outline a credible delivery pathway beyond a notebook prototype

### Why This Matters Now

In the current GenAI/LLM hype cycle, learning drift is almost designed in — people over-index on advanced topics while skipping what makes projects deployable: problem definition, success metrics, and decisions aligned to deployment reality.

### This course is for you if...

You're a software engineer, AI developer, or IT professional who needs a reliable way to turn a workplace problem into a scoped ML initiative — with defensible metrics, explicit risks, and stakeholder-ready reasoning, not “more ML theory.”

---

## Trainer Profile

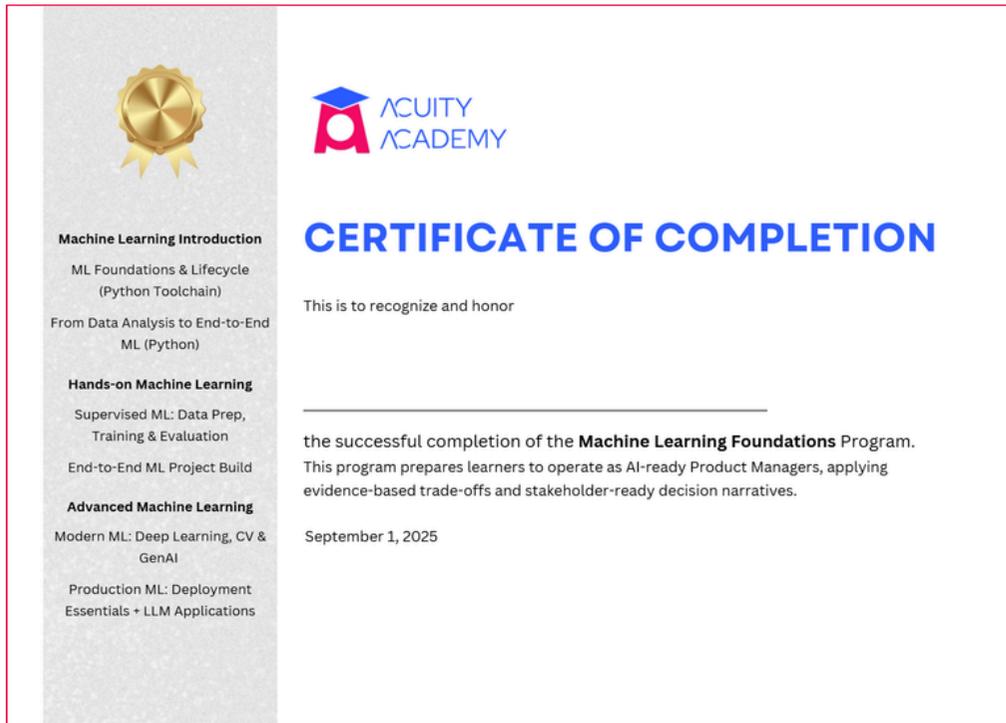


### Iván Palomares Carrascosa

**DataCamp** Instructor | Lead Content Writer — **KDnuggets & Machine Learning Mastery** | **AWS & Google Cloud** ML Certified

Iván is a teacher, trainer, and applied AI practitioner with 15+ years in the field—well before the generative AI wave. He specializes in data-centric machine learning, helping companies and professionals move beyond hype to build real solutions across Machine Learning, Deep Learning, and Data Science, grounded in rigorous understanding of the full data lifecycle.

## Completion Credential



## Earn your Machine Learning Foundations Certificate of Completion

—a standards-based completion credential that recognises demonstrated, workplace-relevant capability.

### Entry Requirement

- A basic IT / software development literacy
- No strict Python prerequisite, but you should be prepared to learn through hands-on coding in Python, using industry-standard libraries such as pandas, scikit-learn, and TensorFlow
- Comfortable working with tabular datasets (e.g., CSV/Excel-style data) and following structured steps for cleaning, analysis, and basic visualisation.
- Participants should be comfortable using digital tools and open to hands-on learning

### Delivery Model

- Online cohort (24 hours): a structured, cohort-based programme designed to convert effort into measurable progress—fast, coherent, and outcome-led.
- Guided workflow delivery: an end-to-end path is made explicit, anchored by early project decisions and a clear milestone sequence.
- Rubric-based feedback loops: instructor review at two checkpoints to keep quality standards clear and prevent “building the wrong thing.”
- Outcome-standard completion: progress is tied to a defined bar (Pilot-Level readiness), not just participation.

### Course Fee

<b>Course Fee:</b>	\$1500.00
<b>GST:</b>	\$135
<b>Total:</b>	<b>\$1635.00</b>

### Course Timing

**Monday & Friday : 7pm - 10pm**

*\*Subject to change*