



AI Professional: Master Machine Learning, NLP, and AI Governance Training Course

5 – 9 July 2027
Sharm El-Sheikh

AI Professional: Master Machine Learning, NLP, and AI Governance Training Course

Ref: 103600324_100076 **Date:** 5 – 9 July 2027 **Location:** Sharm El-Sheikh **Fees:** 4100 Euro

Overview

This five-day AI professional programme in machine learning, NLP and AI governance is a practical tour of how AI models are actually built, from a spreadsheet of raw data to a working text classifier. Participants work in guided Python notebooks using scikit-learn, PyTorch and Hugging Face libraries, and carry one realistic customer-feedback dataset through the whole week: cleaning it, predicting churn with classic machine learning, training a small neural network, fine-tuning a transformer to tag complaints by topic, and finally writing a model card that explains what the model can and cannot be trusted to do. No prior coding experience beyond basic scripting is assumed; every notebook is pre-built so participants can focus on understanding each step. This course is delivered by Agile Leaders Training Center.

Who Should Attend

- Analysts who use Excel or SQL and want to understand how predictive models are trained
- Developers adding a text or image feature to an existing application
- Product owners who brief data teams and need to judge what is feasible
- Researchers and engineers moving from statistics into neural networks and language models
- Team leads preparing an internal AI pilot and its documentation

Departments and Industries

Relevant to any team sitting on text, tabular or image data it wants to use.

- Customer service and contact centre analytics
- Marketing and e-commerce personalisation teams
- Clinical documentation and medical imaging groups
- Quality inspection on manufacturing lines
- Banking teams working with transaction and document data

Learning Objectives

By the end of this course, participants will be able to:

- Use pandas to prepare a messy dataset: missing values, encoding, scaling and train/test splits.
- Compare logistic regression, random forest and gradient boosting models trained in scikit-learn.
- Evaluate a confusion matrix, precision, recall and ROC curve and choose a sensible threshold.
- Build and train a small feed-forward and convolutional network in PyTorch.
- Apply Hugging Face to fine-tune a pre-trained transformer for text classification.
- Build a model card and a short data sheet describing intended use, limits and known biases.



Course Agenda

Day 1: From Raw Data to a Clean Table

- What machine learning is, and when a simple rule works better
- Notebook setup: Python, pandas and plotting in Jupyter
- Exploring the course dataset: distributions, outliers and leakage traps
- Handling missing values, categorical encoding and scaling
- Train, validation and test splits done correctly, ending with a clean, model-ready table

Day 2: Classic Machine Learning

- Linear and logistic regression explained with one feature, then many
- Decision trees, random forests and gradient boosting, used to build a churn model that beats the baseline
- K-means clustering to segment customers
- Cross-validation and grid search for hyperparameters
- Reading feature importance without over-interpreting it

Day 3: Neural Networks

- Neurons, layers, loss functions and gradient descent in plain terms
- Writing a training loop in PyTorch
- Overfitting, dropout and early stopping seen on live curves
- Convolutional networks for images: filters and pooling
- Transfer learning with a pre-trained image model to classify product photos as defective or acceptable

Day 4: Working with Language

- From words to numbers: tokenisation, bag-of-words and embeddings
- How transformers and attention work, at a whiteboard level
- Fine-tuning a pre-trained model to tag customer complaints by topic and urgency
- Named entity extraction and sentiment scoring
- Prompting a large language model versus fine-tuning a small one

Day 5: Responsible Use and Putting a Model to Work

- Checking a model for unequal error rates across customer groups
- Personal data in training sets: minimisation and anonymisation
- Writing a model card and a data sheet
- Wrapping a model in a simple API and a demo interface
- Planning an internal pilot: success measure, users and fallback, with a closing capstone review



Practical Exercises

The following suggested activities carry the course dataset from raw table to documented model.

- Suggested activity: build a churn prediction notebook and compare the model against a simple baseline.
- Suggested activity: train a defect image classifier using transfer learning.
- Suggested activity: fine-tune a pre-trained transformer as a complaint topic tagger.
- Suggested activity: complete a model card for your own final model.

FAQs

How much programming do I need?

Comfort with basic scripting or spreadsheet formulas is enough. Notebooks are provided pre-written; participants edit and run cells rather than write code from scratch.

How many hours does the course run?

Four to five hours a day across five days, around 20 to 25 hours in total, most of it hands-on lab time.

Do I need a powerful laptop?

No. Labs run in a hosted notebook environment, including the neural network and transformer exercises.

Conclusion

Participants leave having built three working models themselves, a churn predictor, an image classifier and a complaint tagger, with the vocabulary to discuss machine learning, NLP and AI governance confidently and a documented model card they can use as a template for the next AI project back at work.