+33 (0)1 48 10 09 46
formation@ambient-it.net
140 Avenue Jean Lolive 93500 Pantin - France

Updated on 10/14/2026

Sign Up

Roboflow Training with Docker

2 days (14 hours)

Overview

Roboflow is a computer vision platform that allows you to prepare data, train models, and deploy vision applications. With Roboflow Inference and Docker, you can run your models on your own infrastructure, expose an API, and build reproducible workflows tailored to your business needs.

Our Roboflow with Docker training will teach you how to link image preparation to containerized deployment, distinguishing between the platform's services and locally executed components.

You'll learn how to organize a versioned dataset for detection, verify annotations, and interpret model results. You'll then set up an inference server in Docker and retrieve its predictions from a Python application.

By the end of the training, you'll be able to structure a deployment using Docker Compose to manage services, volumes, and configuration. You'll know how to measure latency, diagnose errors, and secure sensitive access.

This training also covers workflows for linking detection with business logic, as well as best practices for versioning, testing, and maintenance to ensure your application's reliability in real-world conditions.

Like all our training courses, this one will introduce you to **the latest stable version** of the technology and its new features.

Objectives

- Prepare and version a detection dataset in Roboflow to train a usable model
- Deploy a Roboflow Inference server in a container optimized for the available hardware
- Consume predictions using the Inference SDK from a client application
- Build a workflow that integrates detection, filtering, and business results
- Ensure the reliability of a Docker Compose architecture through security, testing, and diagnostics

Target Audience

- Python Developer
- Machine Learning Engineer
- Data Scientist
- DevOps Engineer
- Computer Vision Engineer

Prerequisites

- Experience with Python for writing scripts and manipulating data
- Knowledge of Docker to launch a container and understand images, ports, and volumes
- Experience using HTTP and the JSON format in application communication
- Basic understanding of datasets, annotations, and evaluation in machine learning

Technical Prerequisites

- Linux workstation with administrative privileges and a 64-bit processor; 16 GB of RAM recommended and 20 GB of available storage
- A working installation of Docker Engine with the Docker Compose plugin to run Linux containers
- Python environment with a package manager and Visual Studio Code editor available
- Roboflow account with an API key and permissions appropriate for the exercises, plus an Internet connection that allows for downloading models and container images
- Optional graphics card for the CPU track; GPU acceleration requires a and the NVIDIA Container Toolkit properly configured

Our Roboflow with Docker training program

[Day 1 - Morning]

Preparing data and the model

- Roboflow architecture: cloud platform, data management, training, and self-hosted inference server
- Organizing workspaces and projects, using object detection as a common thread

- Importing images, verifying annotations, and separating training, validation, and test sets
- Creating a dataset version, preprocessing, and data augmentation without data leakage
- Running a training session, reading accuracy, recall, and mAP, using a provided pre-trained model to ensure the exercises proceed
- Hands-on workshop: preparing a small annotated dataset, creating a version, and analyzing the results of an
  available model.

[Day 1 - Afternoon]

Running Roboflow Inference with Docker

- Role of the Roboflow Inference Server; difference between local execution and hosted API
- Selecting a compatible CPU or GPU image, starting with the Inference CLI, and reading the equivalent Docker command
- Configuring ports, volumes, and variables; exposing port 9001 locally; and ensuring model cache
- Identifying the project and model version, API key, and initial download of weights
- First request with InferenceHTTPClient, reading predictions, and diagnosing connection or access errors
- Hands-on workshop: Start the containerized server, mount a persistent cache, and query a model from a Python script.

[Day 2 - Morning]

Integrating predictions and workflows

- Using the Inference SDK on local images and batches; distinguishing between the HTTP client and the runtime engine
- Interpreting classes, scores, and coordinates; visualizing detections; and adjusting the Confidence
- Building a workflow combining image input, detection, filtering, and business-specific counting
- Running the workflow on the local server, validating inputs, and utilizing structured outputs
- Processing a short video, measuring latency after warm-up, and balancing trade-offs between resolution, frame rate, and accuracy
- Hands-on workshop: Create a counting workflow, run it locally, and compare its results on images and a short video.

[Day 2 - Afternoon]

Ensuring reliable deployment and operation

- Setting up the server and client with Docker Compose, internal network, volumes, and restart policy
- Reproducibility through pinned image versions, model versions, and externalized configuration

- Protection of API keys and network access, client-read secrets, and server exposure restrictions
- Diagnostics with logs and resource metrics, error tracking, memory monitoring, and response times
- Non-regression testing on a test set, controlled updates, and rollback procedures
- Hands-on workshop: deploying a reproducible Compose application, testing its disaster recovery, and documenting its operation.

Target Audience

This training is intended for both individuals and companies—large or small—that wish to train their teams in a new, advanced IT technology or to acquire specific domain knowledge or modern methodologies.

Pre-training Assessment

The pre-training assessment complies with Qualiopi quality standards. Upon final registration, the learner receives a self-assessment questionnaire that allows us to evaluate their estimated proficiency with various types of technologies, as well as their expectations and personal goals for the upcoming training, within the limits imposed by the selected format. This questionnaire also allows us to anticipate certain connection or internal security issues within the company (intra-company or virtual classroom) that could pose challenges for monitoring and ensuring the smooth running of the training session.

Teaching Methods

Hands-On Training: 60% Practical, 40% Theoretical. Training materials distributed in digital format to all participants.

Organization

The course alternates between theoretical input from the instructor—supported by examples and reflection sessions—and group work.

Assessment

At the end of the session, a multiple-choice quiz is administered to verify that participants have successfully acquired the skills.

Certification

A certificate will be issued to each trainee who has completed the entire training program.

Training organization registered under number 11 75 54743 75. This registration does not constitute state accreditation.
© Ambient IT 2015-2026. All rights reserved. Paris, France - Switzerland - Belgium - Luxembourg