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

Updated on 02/26/2026

Register

Fine-tuning Expert Training: QLoRA, Axolotl & MLOps

3 days (21 hours)

Overview

Fine-tuning Expert: QLoRA, Axolotl & MLOps teaches you how to adapt an LLM to your business data with controlled GPU costs. You will leave with a reproducible chain for training, evaluating, and deploying a specialized model (customer support, extraction, classification, RAG).

The training aims to improve your operational skills in QLoRA (quantization + LoRA), training orchestration via Axolotl, and MLOps best practices to industrialize your iterations. We cover dataset preparation (format, cleaning, prompts), hyperparameter selection, memory management, and evaluation strategies (benchmarks, business tests).

The approach is centered on guided workshops and reproducible demos: launching a run, tracking metrics, comparing checkpoints, exporting, and packaging. Deliverables: Axolotl scripts, data prep pipeline, evaluation report, LoRA/merged model, and deployment procedure.

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

Objectives

- Prepare a tuning instruction dataset and define a prompt strategy.
- Configure and execute QLoRA fine-tuning with Axolotl.
- Optimize GPU consumption (quantization, batch, gradient checkpointing).
- Evaluate and compare runs (metrics, business tests, regression).
- Industrialize with MLOps practices (versioning, traceability, deployment).

Target audience

- ML Engineers / Data Scientists
- ML Engineers/MLOps Engineers
- Python developers working on LLMs
- AI technical managers

Prerequisites

- Good level of Python (packages, environments, scripts)
- Solid understanding of deep learning (transformers, training)
- Practical experience with Git and the terminal
- Understanding of data formats (JSON/JSONL) and cleaning

Technical prerequisites

- Linux/macOS/Windows machine with WSL2 recommended
- Minimum 16 GB RAM (32 GB recommended)
- NVIDIA GPU recommended: minimum 12 GB VRAM (24 GB recommended), CUDA installed
- Python 3.10+, conda/venv, PyTorch, Hugging Face, Axolotl, bitsandbytes

Fine-tuning Expert training program: QLoRA, Axolotl & MLOps

[Day 1 - Morning]

LLM fine-tuning fundamentals and adaptation strategy (SFT, DPO, LoRA)

- Clarifying objectives: assistant, extraction, classification, controlled generation
- Choosing the approach: SFT vs. DPO vs. RLHF (when and why)
- Understanding LoRA: rank, alpha, target modules, impact on quality and cost
- Preparing the GPU environment: CUDA, drivers, VRAM, bfloat16/fp16, multi-GPU constraints
- Hands-on workshop: Diagnosing a use case and choosing a fine-tuning strategy.

[Day 1 - Afternoon]

QLoRA in practice: Quantization, memory, and training stability

- QLoRA principle: 4-bit (NF4), double quantization, compute dtype
- Key settings: batch size, gradient accumulation, sequence length, packing
- Stability: gradient checkpointing, clipping, warmup, scheduler, seed, and reproducibility
- Measurements and monitoring: loss, perplexity, overfitting, generation samples

- Hands-on workshop: Launching initial QLoRA fine-tuning on an open-source model and analyzing metrics.

[Day 2 - Morning]

Axolotl: configuration, datasets, and SFT pipelines

- Structure of an Axolotl project: YAML, models, datasets, outputs, and checkpoints
- Data formats: instruction-following, chat templates, JSONL, system/user/assistant fields
- Pre-processing: cleaning, deduplication, filtering, balancing, train/val split
- Optimizations: packing, flash-attn, paged optimizers, VRAM/throughput management
- Hands-on workshop: Building an SFT dataset and executing a reproducible Axolotl run.

[Day 2 - Afternoon]

Alignment and quality: evaluation, DPO, and robustness testing

- Defining an evaluation protocol: test sets, fixed prompts, business criteria
- Measures: win rate, exact match, toxicity, hallucinations, latency, and cost
- DPO introduction: preference data, chosen/rejected pairs, common pitfalls
- Robustness testing: jailbreaks, out-of-distribution data, functional regressions
- Hands-on workshop: Set up automatic evaluation + a DPO run on a small preference set.

[Day 3 - Morning]

Fine-tuning MLOps: traceability, versioning, and governance

- Traceability of experiments: hyperparameters, seeds, artifacts, logs, and run comparisons
- Versioning: data, prompts, models, LoRA adapters, Axolotl configuration, and dependencies
- Secret management and compliance: PII, dataset/model licenses, auditability
- Promotion strategies: baseline? candidate? staging? production with exit criteria
- Hands-on workshop: Setting up a versioning workflow (data + model + configuration) and a release checklist.

[Day 3 - Afternoon]

Packaging, deployment, and monitoring: from LoRA adaptation to inference

- Choosing the delivery mode: LoRA adapters vs. merge, inference quantization, formats (safetensors)
- Serving the model: inference parameters (temperature, top_p), context limits, batching
- CI/CD: non-regression testing, pre-merge evaluation, artifact build and internal publication
- Monitoring: drift, perceived quality, user feedback, alerts on costs/latency/errors
- Hands-on workshop: Build a release pipeline with tests + deploy an inference endpoint and define tracking metrics.

Target companies

This training is intended for both individuals and companies, large or small, wishing to train their teams in a new advanced IT technology or to acquire specific business knowledge or modern methods.

Positioning at the start of training

The placement test at the start of the training course complies with Qualiopi quality criteria. Once they have finalized their registration, learners receive a self-assessment questionnaire that allows us to gauge their estimated level of proficiency in different types of technologies, as well as their expectations and personal goals for the upcoming training course, 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 be problematic for the monitoring and smooth running of the training session.

Teaching methods

Practical training: 60% practical, 40% theory. Training materials distributed in digital format to all participants.

Organization

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

Assessment

At the end of the session, a multiple-choice questionnaire is used to verify that the skills have been correctly acquired.

Certification

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