Proprietary Data and LLMs

Course Overview

This advanced course is designed for senior machine learning engineers, tech/product leads, and hands-on engineering managers aiming to leverage Large Language Models (LLMs) for enterprise applications, especially those with domain-specific requirements and sensitive data. Participants will dive deep into the art and science of fine-tuning pre-trained models to excel in tasks ranging from instruction following to reinforcement learning.

Learning Outcomes

By the end of this course, you will:

1. Gain a comprehensive understanding of fine-tuning techniques, from continued base model pre-training to reinforcement learning with human feedback (RLHF)
2. Develop practical skills in fine-tuning LLMs through extensive hands-on coding exercises using state-of-the-art tools and frameworks
3. Learn to adapt pre-trained models to specific industries and domains, ensuring they excel in specialized tasks
4. Master strategies to fine-tune models while preserving sensitive data which is critical for many enterprise use cases
5. Explore methodologies for effectively evaluating fine-tuned models and strategies for deploying them in production environments

Recommended Prerequisites

- Proficiency in Python programming
- Solid understanding of machine learning fundamentals
- Familiarity with deep learning concepts
- Intermediate knowledge of Natural Language Processing (NLP) is recommended

Detailed Curriculum Schedule

Topic	Topics
Introduction to LLMs and Fine-Tuning (2 hr)	• Overview of Large Language Models (LLMs) and the current tech/tools landscape for building applications • To tune, or not to tune ➜ guidance for when fine-tuning is appropriate with case studies • Understanding different fine-tuning concepts • Strategies for preserving data privacy in fine-tuning (i.e. federated p-tuning, differential privacy, etc.)
Hands-on with Different Fine-Tuning Methods (2 hrs)	• Few-shot learning (FSL) for domain-specific tasks • Retrieval augmented generation best practices • Instruction fine-tuning for task-specific applications • Meta AI Humpback method to automate the creation of instruction-tuning datasets for any corpus
Human vs. AI Feedback & Continuous Improvement (2 hrs)	• Reinforcement learning with human feedback (RLHF) • Reinforcement learning with AI feedback (RLAIF) • Evaluation and continuous improvement approaches

Register Now!

About FourthBrain

FourthBrain trains engineers, developers, data scientists, and leaders to make an impact in the Artificial Intelligence field, with our flexible, accessible education programs. We are training a new generation of engineers and leaders with more than just technical ability; they have an awareness and mindset of what is needed to succeed with AI. We are part of the AI Fund, founded by Andrew Ng.