

Training Course on Fine-Tuning and Customizing Pre-trained LLMs:

Professional Development Training Course Outline

Category	Data Science	Duration	5 Days
Base Fee	\$1,500 USD	Delivery Mode	Online / On-site Hybrid

Course Introduction

Training Course on Fine-Tuning & Customizing Pre-trained LLMs delves into the advanced techniques of **Fine-Tuning** and **Customizing Pre-Trained Large Language Models (LLMs)**. Participants will gain practical expertise in **adapting foundational models** to excel in specific **tasks and niche domains**, moving beyond generic applications. We will explore the methodologies for **optimizing LLM performance**, ensuring **domain-specific accuracy**, and unlocking the full potential of **generative AI** for specialized business needs. This training is crucial for professionals seeking to build **high-performing, bespoke AI solutions** and gain a significant **competitive advantage** in the rapidly evolving landscape of artificial intelligence.

In today's data-driven world, the ability to tailor powerful LLMs for precise applications is a critical skill. This program focuses on equipping participants with the knowledge and hands-on experience to transform general-purpose models into highly effective, **domain-aware AI assistants**. Through practical exercises, **real-world case studies**, and exploration of **cutting-edge parameter-efficient fine-tuning (PEFT) techniques**, attendees will master the art of **LLM adaptation**, addressing challenges like **data scarcity**, **hallucination mitigation**, and **cost optimization**. This course is designed to empower individuals and organizations to deploy **intelligent, specialized LLM solutions** that drive innovation and efficiency.

Programme Curriculum

Training Course on Fine-Tuning & Customizing Pre-trained LLMs

Introduction

Training Course on Fine-Tuning & Customizing Pre-trained LLMs delves into the advanced techniques of **Fine-Tuning** and **Customizing Pre-Trained Large Language Models (LLMs)**. Participants will gain practical expertise in **adapting foundational models** to excel in specific

tasks and niche domains, moving beyond generic applications. We will explore the methodologies for **optimizing LLM performance**, ensuring **domain-specific accuracy**, and unlocking the full potential of **generative AI** for specialized business needs. This training is crucial for professionals seeking to build **high-performing, bespoke AI solutions** and gain a significant **competitive advantage** in the rapidly evolving landscape of artificial intelligence.

In today's data-driven world, the ability to tailor powerful LLMs for precise applications is a critical skill. This program focuses on equipping participants with the knowledge and hands-on experience to transform general-purpose models into highly effective, **domain-aware AI assistants**. Through practical exercises, **real-world case studies**, and exploration of **cutting-edge parameter-efficient fine-tuning (PEFT) techniques**, attendees will master the art of **LLM adaptation**, addressing challenges like **data scarcity**, **hallucination mitigation**, and **cost optimization**. This course is designed to empower individuals and organizations to deploy **intelligent, specialized LLM solutions** that drive innovation and efficiency.

Course Duration

5 days

Course Objectives

1. Understand and implement various fine-tuning approaches, including **full fine-tuning**, **LoRA (Low-Rank Adaptation)**, and **QLoRA** for efficient model adaptation.
2. Acquire proficiency in **high-quality data collection**, **data cleaning**, and **feature engineering** for **domain-specific datasets**.
3. Develop advanced **prompt optimization strategies** to complement fine-tuning for **task-specific performance**.
4. Learn to employ **robust evaluation metrics** and **benchmarking techniques** for assessing **model accuracy**, **relevance**, and **bias mitigation**.
5. Implement **PEFT methods** to significantly reduce computational costs and training time for large models.
6. Grasp the fundamentals of **RLHF** and its role in aligning LLM behavior with desired outcomes.
7. Gain the ability to tailor LLMs for **industry-specific applications** in sectors like **healthcare**, **finance**, and **legal tech**.
8. Learn strategies and techniques to reduce the generation of **inaccurate or irrelevant content** in fine-tuned models.
9. Understand the considerations and best practices for **deploying customized LLMs** in production environments.
10. Address biases and promote fairness in datasets and model outputs.
11. Explore techniques for improving the **latency and throughput** of fine-tuned models.
12. Identify and utilize powerful **open-source foundational models** for cost-effective fine-tuning.
13. Develop comprehensive workflows for data preparation, fine-tuning, evaluation, and deployment.

Organizational Benefits

- Develop LLMs that deliver highly precise and contextually relevant responses tailored to specific business needs, reducing generic or irrelevant outputs.
- Automate complex, domain-specific tasks, leading to significant time and resource savings across various departments.

- Create proprietary AI capabilities that differentiate the organization in the market by addressing unique industry challenges with specialized LLM solutions.
- Reduce reliance on expensive general-purpose API calls by deploying more efficient, task-specific models trained on targeted data.
- Maintain greater control over sensitive data by fine-tuning models in-house, aligning with compliance and regulatory requirements.
- Minimize the generation of erroneous or misleading information, leading to more reliable and trustworthy AI outputs.
- Accelerate the development and deployment of new AI-powered products and services, fostering a culture of innovation.
- Embed specialized knowledge and jargon directly into the model's architecture, enabling it to understand and respond with industry-specific nuance.

Target Audience

1. Machine Learning Engineers
2. Data Scientists
3. AI Developers.
4. NLP Researchers
5. MLOps Engineers
6. Product Managers (AI-focused).
7. Solutions Architects
8. Technical Leads / Team Leads

Course Outline

Module 1: Foundations of LLM Fine-Tuning & Customization

- Introduction to Pre-trained LLMs: Architectures (Transformer, GPT, BERT), strengths, and limitations for specific tasks.
- The Need for Fine-Tuning: Why generic LLMs fall short in specialized domains; comparing fine-tuning with prompt engineering and RAG.
- Types of Fine-Tuning: Full fine-tuning vs. parameter-efficient methods (PEFT).
- Key Concepts: Transfer learning, catastrophic forgetting, overfitting, and generalization in LLMs.
- Ethical Considerations in LLM Adaptation: Bias, fairness, and responsible AI development.
- **Case Study:** Analyzing how a general-purpose LLM struggles with medical jargon and patient privacy concerns, highlighting the need for domain-specific adaptation.

Module 2: Data Preparation & Curation for Domain Adaptation

- High-Quality Dataset Acquisition: Sourcing and collecting relevant, domain-specific data (text, code, structured data).
- Data Cleaning & Preprocessing: Tokenization, normalization, handling noisy data, and anonymization of sensitive information.
- Data Labeling and Annotation Strategies: Supervised learning paradigms, active learning for efficient labeling.
- Building Custom Datasets: Best practices for creating balanced, representative, and diverse datasets for fine-tuning.
- Dataset Versioning & Management: Tools and techniques for tracking and managing datasets.

- **Case Study:** Curating a legal document dataset for contract analysis, focusing on identifying clauses and redacting sensitive client information for compliance.

Module 3: Parameter-Efficient Fine-Tuning (PEFT) Techniques

- Introduction to PEFT: Why PEFT is crucial for large models; benefits in terms of computational resources and time.
- LoRA (Low-Rank Adaptation): Theory, implementation, and practical application for efficient fine-tuning.
- QLoRA: Quantized LoRA for further memory and computational savings.
- Other PEFT Methods: Exploring techniques like Prefix-Tuning, Prompt-Tuning, and Adapter layers.
- Choosing the Right PEFT Method: Decision criteria based on dataset size, computational resources, and desired performance.
- **Case Study:** Fine-tuning a large financial LLM (e.g., Llama-2) on a limited dataset of earnings call transcripts using LoRA to achieve high accuracy on sentiment analysis without full model retraining.

Module 4: Supervised Fine-Tuning (SFT) Implementation

- SFT Workflow: Step-by-step process from data preparation to model deployment.
- Training Frameworks & Libraries: Hands-on with Hugging Face Transformers, PyTorch, and TensorFlow for SFT.
- Hyperparameter Optimization: Learning rate scheduling, batch size, epoch selection, and regularization techniques.
- Model Checkpointing & Resumption: Strategies for saving and loading model states during training.
- Distributed Training for Large Models: Scaling fine-tuning across multiple GPUs or machines.
- **Case Study:** Implementing an SFT pipeline to adapt a pre-trained LLM for customer service chatbot responses, focusing on improving intent recognition and factual accuracy for product-specific queries.

Module 5: Evaluation & Performance Optimization

- Quantitative Evaluation Metrics: Perplexity, ROUGE, BLEU, F1-score, accuracy, and domain-specific metrics.
- Qualitative Evaluation & Human-in-the-Loop: Expert review, user feedback, and A/B testing for fine-tuned models.
- Debugging Fine-Tuning Issues: Identifying and resolving common problems like overfitting, underfitting, and catastrophic forgetting.
- Model Drift & Monitoring: Strategies for detecting and addressing performance degradation in deployed models.
- Benchmarking Custom LLMs: Comparing performance against baseline models and industry standards.
- **Case Study:** Evaluating a fine-tuned LLM for legal document summarization, comparing its ROUGE scores against human-generated summaries and identifying areas for improvement through qualitative review.

Module 6: Reinforcement Learning from Human Feedback (RLHF) Concepts

- Introduction to RLHF: Aligning LLM behavior with human preferences and values.

- **Reward Modeling:** Designing and training reward models from human preference data.
- **Proximal Policy Optimization (PPO):** Understanding the algorithm and its application in RLHF.
- **Direct Preference Optimization (DPO):** A simpler alternative to PPO for preference learning.
- **Challenges and Limitations of RLHF:** Data collection, ethical considerations, and computational expense.
- **Case Study:** Applying RLHF principles to a content generation LLM to refine its writing style and tone to match a specific brand's guidelines, reducing off-brand outputs.

Module 7: Advanced Customization & Deployment Strategies

- **Retrieval-Augmented Generation (RAG) with Fine-Tuned LLMs:** Integrating external knowledge bases for enhanced factual accuracy.
- **Continual Learning & Incremental Fine-Tuning:** Updating models with new data without retraining from scratch.
- **Model Quantization & Pruning:** Techniques for reducing model size and improving inference speed for deployment.
- **Serving Fine-Tuned LLMs:** Deployment on cloud platforms (AWS SageMaker, Google Cloud AI Platform, Azure ML) and edge devices.
- **API Integration & Scalability:** Building robust APIs for accessing fine-tuned models and handling high inference loads.
- **Case Study:** Deploying a fine-tuned LLM for real-time customer support in a call center, demonstrating low-latency responses and high accuracy on frequently asked questions, potentially integrated with a RAG system for dynamic information retrieval.

Module 8: Industry Applications & Future Trends

- **LLMs in Healthcare:** Clinical note summarization, diagnostic assistance, drug discovery.
- **LLMs in Finance:** Fraud detection, market analysis, personalized financial advice.
- **LLMs in Legal Tech:** Contract review, legal research, e-discovery.
- **Custom LLMs for Code Generation & Development:** Improving developer productivity.
- **Future of LLM Adaptation:** Multi-modal fine-tuning, smaller specialized models, ethical AI governance.
- **Case Study:** Developing a custom LLM for a specific manufacturing domain, capable of analyzing technical specifications, generating maintenance reports, and answering queries about machinery operations, leading to improved operational efficiency.

Training Methodology

This course employs a **blended learning approach**, combining theoretical foundations with extensive **hands-on practical sessions**. The methodology is designed to ensure participants gain both conceptual understanding and applied skills in fine-tuning and customizing LLMs.

- **Interactive Lectures & Discussions:** Engaging presentations of core concepts, current trends, and best practices.
- **Live Coding Demonstrations:** Step-by-step walkthroughs of fine-tuning pipelines using popular frameworks like Hugging Face.
- **Hands-on Labs & Exercises:** Practical assignments where participants implement fine-tuning techniques on real-world datasets.
- **Case Study Analysis:** Deep dives into industry-specific applications, examining challenges, solutions, and outcomes of LLM customization.

- **Q&A and Troubleshooting Sessions:** Dedicated time for addressing participant queries and debugging common issues.
- **Project-Based Learning:** Opportunities for participants to work on a mini-project, applying learned skills to a task of their choice.
- **Collaborative Learning:** Encouraging peer-to-peer learning and group problem-solving.
- **Access to Cloud Computing Resources:** Provision of necessary GPU-enabled environments for computationally intensive tasks.

Register as a group from 3 participants for a Discount

Send us an email: info@fineskilltrainingcenter.org or call +254769199797

Certification

Upon successful completion of this training, participants will be issued with a globally- recognized certificate.

Tailor-Made Course

We also offer tailor-made courses based on your needs.

Key Notes

- The participant must be conversant with English.
- Upon completion of training the participant will be issued with an Authorized Training Certificate
- Course duration is flexible and the contents can be modified to fit any number of days.
- The course fee includes facilitation training materials, 2 coffee breaks, buffet lunch and A Certificate upon successful completion of Training.
- One-year post-training support Consultation and Coaching provided after the course.
- Payment should be done at least a week before commence of the training, to Fineskill Training Center account, as indicated in the invoice so as to enable us prepare better for you.

Upcoming Scheduled Sessions

Start Date	End Date	Location	Price
21 Sep 2026	25 Sep 2026	Online	\$1,000
21 Sep 2026	25 Sep 2026	Physical	\$1,500

Start Date	End Date	Location	Price
21 Sep 2026	25 Sep 2026	Physical	\$0
21 Sep 2026	25 Sep 2026	Physical	\$0
21 Sep 2026	25 Sep 2026	Physical	\$0
21 Sep 2026	25 Sep 2026	Physical	\$0
21 Sep 2026	25 Sep 2026	Physical	\$0
21 Sep 2026	25 Sep 2026	Physical	\$0

Ready to enroll or request a customized program? Book online or contact us:

Register Online Now

Email: info@fineskilltrainingcenter.com | Website: www.fineskilltrainingcenter.com