

Mastering AI and Machine Learning

A Comprehensive Curriculum for Foundational and Advanced Topics

AI & ML Curriculum Overview

This comprehensive curriculum is designed to guide learners through the core foundations of AI and machine learning, advancing through specialized topics such as generative AI, large language models, retrieval-augmented generation systems, and multi-agent systems. Additionally, it covers critical aspects of MLOps, deployment, scaling, and AI governance. Below is a detailed breakdown of each phase.

Phase 1: Core AI & ML Foundations

Mathematics for AI

- **Topics:** Linear Algebra, Probability, Statistics, Optimization, Gradient Descent
- **Tools:** NumPy, SciPy

Python for Machine Learning

- **Skills:** Data manipulation, visualization, feature engineering
- **Libraries:** Pandas, Matplotlib, Scikit-Learn

Core ML Algorithms

- **Algorithms:** Regression, Classification, Clustering, Decision Trees, Random Forest, XGBoost
- **Tools:** Scikit-Learn

Deep Learning Essentials

- **Concepts:** Neural networks, CNN, RNN, LSTM, Transformers (intro)
- **Frameworks:** TensorFlow, PyTorch

Model Lifecycle Management

- **Processes:** Training, validation, testing, model evaluation metrics, hyperparameter tuning
- **Tools:** MLflow, DVC

Phase 2: Generative AI & Large Language Models

Foundation Model Concepts

- **Topics:** What are LLMs, Tokenization, Transformer Architecture, Self-Attention
- **Models:** GPT, LLaMA, Claude, Gemini

Prompt Engineering

- **Techniques:** Instruction prompting, few-shot learning, role-based prompting, chain-of-thought
- **Platforms:** OpenAI Playground, PromptLayer

Customizing LLMs

- **Techniques:** Fine-tuning, LoRA, prompt-tuning, embeddings, adapter layers
- **Frameworks:** HuggingFace, OpenAI

LangChain Basics

- **Concepts:** Introduction, Chains, Agents, Tools, Memory, Callbacks, Streaming responses
- **Tools:** LangChain

Vector Databases

- **Functions:** Embedding creation, similarity search, storage, and retrieval
- **Databases:** FAISS, Pinecone, Chroma

Phase 3: Retrieval-Augmented Generation (RAG) Systems

RAG Architecture Overview

- **Importance:** Why RAG is needed, how it improves accuracy and grounding
- **Conceptual Frameworks:** Knowledge Store Design

Knowledge Store Design

- **Processes:** Chunking, embedding, retrieval, re-ranking, indexing
- **Tools:** FAISS, Pinecone, Chroma

Integrating RAG with LLMs

- **Techniques:** Context injection, query optimization, hallucination control
- **Platforms:** LangChain, OpenAI API

Advanced RAG Use Cases

- **Applications:** Enterprise knowledge base search, document Q&A, conversational RAG
- **Tools:** Custom APIs

Enterprise Implementation

- **Considerations:** Secured RAG pipelines, compliance, latency optimization
- **Platforms:** AWS Bedrock, GCP Vertex AI

Phase 4: Agentic AI — Multi-Agent Systems & Orchestration

Introduction to Agentic AI

- **Overview:** Evolution from LLM apps to multi-agent ecosystems
- **Concepts:** AI Agent Architecture

AI Agent Architecture

- **Components:** Planner, Executor, Memory, Tool Interface, Communication Layer
- **Tools:** CrewAI, AutoGen, LangGraph

LangGraph Deep Dive

- **Topics:** Graph-based agent orchestration, task flows, message passing, state management
- **Frameworks:** LangGraph

CrewAI / AutoGen Overview

- **Concepts:** Multi-agent collaboration, task delegation, error handling
- **Platforms:** CrewAI, AutoGen

Model Context Protocol (MCP)

- **Functions:** Tool registration, context passing, execution orchestration
- **Frameworks:** MCP Frameworks

Orchestrator Implementation

- **Skills:** Multi-agent coordination across LLMs, APIs, voice & data layers
- **Tools:** Custom SDKs, LangGraph + MCP

Real-World Integration

- **Applications:** Genesys, AWS Connect, NovaSonic, Twilio, and enterprise connectors
- **Tools:** Intellectt Universal Connector

Phase 5: MLOps, Deployment, and Scaling

MLOps Fundamentals

- **Topics:** CI/CD for ML, model versioning, monitoring, A/B testing
- **Tools:** Docker, MLflow, Kubeflow

Cloud Deployment

- **Techniques:** Containerization, Fargate, EKS, Lambda orchestration
- **Platforms:** AWS, Azure, GCP

LLM & Agent Deployment

- **Skills:** Serving LLMs securely, latency control, API gateways
- **Tools:** FastAPI, LangServe

Enterprise-Grade DR & Failover

- **Design:** Designing fault-tolerant AI systems
- **Solutions:** AWS Multi-AZ

Phase 6: AI Governance, Ethics & Compliance

Responsible AI Practices

- **Topics:** Model interpretability and explainability
- **Tools:** SHAP, LIME, Fairlearn

Data Privacy & Compliance

- **Considerations:** Data privacy, IP ownership, and audit frameworks
- **Tools:** GDPR Tools, CCPA Tools

Security & Access Management

- **Skills:** Secure API key and access management for multi-agent setups
- **Tools:** Vault, AWS Secrets Manager

This curriculum provides a structured pathway to mastering diverse aspects of AI and ML, equipping learners with the necessary skills to innovate and implement AI solutions responsibly and effectively..

© 2025 TekOne. All rights reserved. No part of this content may be reproduced, distributed, or transmitted in any form without prior written permission from TekOne.