

Introduction to Machine Learning and AI Engineering

This AI and Machine Learning course helps you build reliable AI systems and deploy usable AI/ML workflows in any job. You'll understand how models learn, build machine learning models, and master rigorous evaluation and pipeline design. The course lays out ML fundamentals like regression and tensors, then builds toward real LLM applications, focusing on hallucination reduction and retrieval-based apps. You'll integrate latency awareness to ensure testable, grounded outputs, learning how modern AI scales in production using tools like PyTorch, TensorFlow, and LangChain. By the end, you'll have a foundation in AI workflow thinking, equipping you with skills like ML fundamentals, RAG, evaluation, and AI security to consistently build reliable applications.

[CBT Nuggets course material](#) →

 STUDY PLAN

WEEK 1

What is AI, ML, DL, and GenAI?

Defining Artificial Intelligence

Machine Learning and Deep Learning Foundations

Transformers and Generative AI Basics

Challenge ☒

Compare Machine Learning and AI Engineering

Introduction

Train Validation Test Split

Baseline Training Curves

Improved LSTM and Overfitting

Model Comparison and Error Analysis

Challenge ☒

Introduction

From AI History to AI Engineering

AI Engineering and Foundation Models

Model Selection, Prompting, and RAG

RAG, Vector Stores, and AI Engineering Tradeoffs

Challenge

Explore Data Science & ML Development Environments

Introduction

What Is an AI Development Environment?

Conda and Anaconda for Environment Management

Jupyter Notebook and Google Colab

Vibe Coding

Cursor Agentic Coding

Challenge

Solution

Think in Machine Learning with Pseudocode

Introduction

Writing Psuedocode

Supervised and Unsupervised Learning

Linear Regression

Logistic Regression

Neural Networks in Pseudocode

Challenge ☒

Model Linear Regression As Learning A Line

Introduction

What Linear Regression Solves

Plotting the Data First

Interpreting Weight and Bias

Predictions and Residuals

Challenge ☒

Classify with Logistic Regression Decision Lines

Introduction

Classification with Logistic Regression

One Input Binary Classification

Concept Chains in Classification

Decision Lines in Two Dimensions

Challenge

Build a One-Neuron Model With the Perceptron

Introduction

Logistic Regression to Perceptron

What is a Perceptron?

What is a Weighted Average or Weighted Sum?

Start to Finish

Challenge ☒

Explain Scalars, Vectors, Matrices and Tensors

Introduction

Scalars & Vectors

Matrices & Tensors

Shapes & Dot Product

Matrix Multiplication

Challenge

Break Down the ML Pipeline Step by Step

Introduction

The Machine Learning Pipeline Matters

Data Ingestion and Data Preparation in Practice

Inspecting the Data with Pandas

Cleaning and Validating the Data

From Clean Data to the Full ML Pipeline

Challenge ☒

Build Your First PyTorch Machine Learning Model

Introduction

Understanding Computational Graphs

Loading Data into PyTorch

Building a PyTorch Model

Challenge ☒

Implement Linear Regression with PyTorch

Introduction

Implement Linear Regression with PyTorch & Pseudocode

Loading the CSV and Selecting the Feature and Target

Splitting Data and Converting to Tensors

Building and Training the Linear Regression Model

Challenge

Implement Logistic Regression with PyTorch

Introduction

Pseudocode and Data Setup for Logistic Regression

Feature Engineering for the Target Label

Splitting the Data and Preparing Tensors

Building and Training the Logistic Regression Model

Challenge ☒

Implement a Non-Linear Neural Network in PyTorch

Introduction

Create A Combined Dataset & Load Helper Utility

What Is Non-Linear Data?

Normalize Data Before Training the Neural Network

What is ReLU?

Challenge ☒

Analyze Tensors To Prevent PyTorch Errors

Introduction

Tensor Shapes and Dimensions

Tensor Shapes in Code

Missing Batch Dimensions

Tensor Data Types

Reshaping and Inspecting Tensors

Challenge

Apply Data Pipelines to Train PyTorch Models

Introduction

Why Batching Matters

Transforms and Dataset

Inspecting DataLoader Batches

Model Construction and Training

Challenge ☒

Explain Loss, Gradients, and Optimizers

Introduction

Mean Squared Error for Regression

Loss and Metrics in Classification

Gradients and Backward Pass

Optimizer Step and Learning Rate

Challenge ☒

Build a PyTorch Neural Network Image Classifier

Introduction

Device Management and Image Transforms

Loading and Inspecting MNIST Data

Flattening Images for Classification

Training and Evaluation Loops

Challenge ☒

Apply TensorFlow to Build Your First NLP Pipeline

Introduction

TensorFlow and Keras Basics

Text Data for NLP

Text Vectorization and Batching

Building a TensorFlow Text Classifier

Challenge ☒

Build an RNN Text Classifier with TensorFlow

Introduction

Word IDs and Embedding Vectors

Vocabulary IDs and Padding

Embedding Shapes and IMDb Movie Reviews

Baseline Embeddings and RNN Classifiers

Challenge ☒

Compare RNNs, LSTMs, and GRUs with TensorFlow

Introduction

RNN Memory Limits and Data Setup

Comparing Sequence Models

Challenge

Evaluate a TensorFlow Movie Review Text Classifier

Introduction

Train Validation Test Split

Baseline Training Curves

Improved LSTM and Overfitting

Model Comparison and Error Analysis

Challenge ☒

Introduction

From AI History to AI Engineering

AI Engineering and Foundation Models

Model Selection, Prompting, and RAG

RAG, Vector Stores, and AI Engineering Tradeoffs

Challenge

Explore LLM APIs for AI Engineering

Introduction

OpenAI API Setup + Documentation RAG Helper ☒

First OpenAI API Request

Inspect the OpenAI Response Object

Challenge

Analyze LLM Responses and Token Costs

Introduction

API Setup and Response Helpers

Inspecting Complete and Incomplete Responses

Word Splits and LLM Tokens

Delimiters, Prompt Injection, and Context Windows

Challenge

Program LLMs Using Prompt Engineering

Introduction

Weak and Strong Prompts

Context, Role, Expectation, and Shots

Prompt Context, RAG, and Fine Tuning

Delimiters, JSON, and Prompt Injection

Challenge

Build a RAG Application with Chroma & LangChain

Introduction

RAG Setup and Knowledge Base Basics

Chunking and Embeddings for RAG

Chroma Retrieval and Grounded RAG Answers

Challenge Solution

Evaluate AI System Outputs

Introduction

Evaluation Metrics and Human Review

RAG Evaluation Setup

Testable RAG Functions and Eval Data

Code-Based RAG Evaluators

Not Found Behavior and RAG Score Signals

CHALLENGE ☒

Automate AI Evaluation Workflows

Introduction

Evaluation Review and Notebook Setup

Rebuilding the RAG Evaluation Loop

Source Lines and Overall Eval Scores

Failure Inspection, LLM-as-Judge, and RAG Setting

LangChain, LangSmith, and Evaluation Workflows

CHALLENGE ☒

Integrate Tools Into AI Systems

Introduction

Plain LLMs and Tool-Enabled AI Systems

Defining Python Tools and Tool Schemas

Tool Loops and Tool Libraries

Reusable Tool Calling and Agent Foundations

CHALLENGE ☒

Build AI Agentic Workflows

Introduction

Agent Workflows and Stop Conditions

Agent Tools, Instructions, and Tracing

Agent Evaluation and Course Wrap-Up

CHALLENGE ☒

WEEK 2