

Introduction to Machine Learning

This Introduction to Machine Learning training covers how to turn vast datasets into actionable insights, predictions, and trends. Learn to build models with scikit-learn, PyTorch, TensorFlow, and explore LLM development with OpenAI, LangChain, and HuggingFace. Perfect for junior data scientists, IT teams, or as a comprehensive reference for machine learning and AI best practices.

[CBT Nuggets course material](#) →

WEEK 1

Explore How AI Agents Navigate Driving Directions

154 min.

What is Artificial Intelligence?	11
Grand Search Auto	16
Explore the Frontier	6
Depth-First Search	8
Breadth-First Search	6
Greedy-Best First and A* Search	7
Challenge	1

Apply Probability to Real-World AI Problems

Introduction	2
Probability of Rolling One 6-sided Die	11
Die Roll Simulation	1
Die Roll Probabilities	1
Probability of Rolling Two 6-sided Dice	8
Probability Distribution of Rolling Two 6-sided Dice	1
Conditional Probability	6
Bayesian Networks	9
Baye's Rule	1
Recapitulation	2
Challenge	1

Define What is Machine Learning?

Introduction	8
What is Machine Learning?	9
Supervised	14

Unsupervised	4
Build an Image Classifier	8
Predicting Lumber Prices with Linear Regression	6
Challenge	1

Setup a Machine Learning Development Environment

Introduction	6
--------------	---

WEEK 2

153 min.

Anaconda, Conda and Jupyter	16
Google Colab	8
Cloud Services: AWS, GCP, and Azure	7
Vast.ai: the market leader in low-cost cloud GPU rental	4
Challenge	1

Explore Data Pipelines and Linear Regression

Introduction	5
What is a Machine Learning Model?	11
Predicting Lumber Prices: Data Collection	7
Predicting Lumber Prices: Data Cleaning & Preprocessing	6
Predicting Lumber Prices: Feature Extraction	18
Challenge	1

Apply Regression Concepts for Supervised Learning

Introduction	5
--------------	---

A Brief and Bizarre History of Linear Regression	5
Explore Linear Relationships: Ordinary Least Squares	16
Seaborn Line of Best Fit	7
Ordinary Least Squares with Matlab's PolyFit	11
Challenge	5

Examine Cost Functions and Parameter Tuning

Introduction	4
Mean Absolute Error	4
Mean Squared Error	2
Root Mean Squared Error	3
Cost Functions	7

WEEK 3

152 min.

Calculate Your Model's Performance	21
Challenge	1

Implement Gradient Descent for Linear Regression

Introduction	4
Exploring Gradient Descent Concepts	6
Exploring The Gradient Descent Algorithm	5
Gradient Descent Behind the Scenes	14
Implementing The Gradient Descent Algorithm	12
Challenge	1

Vectorize Operations for Multiple Regression

Introduction	4
Multiple Linear Regression	9
Vectorization	8
Non-Vectorized Operations	7
Interpreting the Weights	8
Vectorized Operations	3
Congrats! 🎉	1
Challenge	1

Explore Feature Engineering and Data Preparation

Introduction	7
What is Feature Engineering?	5
Handling Missing Data	11
Handling Outliers	7
One Hot Encoding	6
Define, Split and Scale Features	9
Measuring Survival Accuracy	2

Challenge	1
-----------	---

Identify Key Classification Algorithms

Introduction	5
From Regression to Classification	9
Binary Classification Visualization	1
Multiclass Classification Visualization	1
Logistic Regression	7

Decision Trees	7
Random Forests	5
Support Vector Machines	3
Perceptrons	4
Challenge	1

Implement Logistic Regression with Python

Introduction	6
What is Logistic Regression?	5
The Sigmoid Formula and Function	3
Logistic Regression in 4 lines of Code	8
Implement Logistic Regression	26
Challenge	1

Build a Python Decision Tree Classification Model

Introduction	6
What is a Decision Tree?	1
Concepts Video	9
Entropy, Information Gain, and Gini Impurity	5
Import Libraries, Feature Engineering and One-Hot Encoding	14
Train, Test, Predict, and Measure Model Performance	14
Challenge	1

Build a Python Random Forest Classification Model

Introduction	4
What is a Random Forest	7

WEEK 4

154 min.

WEEK 5

157 min.

Random Forest Concepts	6
Import Libraries, Feature Engineering and One-Hot Encoding	10
Train, Test, Predict, and Measure Model Performance	9
Bonus: Hyperparameter Tuning Video	3
Challenge	1

Apply Regularization to Overcome Overfitting

Introduction	5
What is Overfitting?	8
Three Options for Handling Overfitting	6
Overfitting for Classification	6
Comparing Cost Functions	8
Perform Logistic Regression with Regularization	7
Challenge	1

Build a Support Vector Machine Classifier

Introduction	4
What is a Support Vector Machine?	7
Optimal Hyperplanes and the Margin	8
Data Loading and PreProcessing	13
Build and Evaluate the Model	6
Breast Cancer Wisconsin (Diagnostic) Dataset	3
Challenge	1

Build a K-Nearest Neighbors Classifier

Introduction	8
What is K-Nearest Neighbors?	9
KNN vs. Other Classifiers	6
What is Imbalanced Data?	6
Data Loading and EDA	6
Data PreProcessing	10

WEEK 6

153 min.

Build and Evaluate the Model	6
Challenge	1

Explore Neural Network Basics With The Perceptron

Introduction	6
Neurons as the building blocks of neural networks	5
Perceptrons As Artificial Neurons	5
How Activation Functions Work	5
Why Linearly Separable Data Is Key	6
Build A Simple Binary Perceptron Classifier	8
Solution Code	1
Challenge: Complete The Perceptron Function ☒	4
Solution Video	7

Implement a Perceptron for Classification

Introduction	4
What is a Perceptron?	4
The Perceptron Rule and Neurons	9
Implement a Perceptron from Scratch	11

The Perceptron Challenge	2
Challenge Code	1
Solution Video	2
Solution Code	1
Bonus Resources	6

Explore PyTorch Fundamentals for Machine Learning

Introduction	4
What Is PyTorch and Why It Is Useful?	11
Set up a PyTorch Development Environment	6
Leverage Tensors Concepts	12
Leverage Tensors Programmatically	11
Challenge	2

Leverage PyTorch Tensor Attributes and Operators

Tensor attributes	13
-------------------	----

WEEK 7

152 min.

Tensor Math Operators	8
Matrix Multiplication	10
The PyTorch Double Challenge	5
Physics Lab Challenge	1
Matrix Multiplication Golden Rule	1
Breakdown for Solution 1: Physics Lab	1
Green Bottle Challenge	1
Solution 2: Green Bottle	1

Explore Fundamental PyTorch Tensor Operations

Review: Matrix Multiplication Errors	12
Min, Max, Mean, and Sum (Tensor Aggregation)	8
Navigating Positional Min Max Values	5
The Challenge	4
Solution Video	9

Apply PyTorch Tensor Manipulation and Indexing

Reshape, View, and Stack Tensors	13
Squeeze and Unsqueeze Tensors	11
Permute Tensors	5
Index Tensors	9
Challenge: Tensor Transformer	9

Explore Gradient Descent & Back Propagation

Introduction	6
Gradient Descent	5
Forward Propagation	6
Back Propagation	5
Training, Validation, and Test Datasets	5
Split The Train Test Datasets	12

WEEK 8

199 min.

Build a Linear Regression Model	9
Challenge: PyTorch Research	1

Predict Ice Cream Sales with PyTorch Regression

Introduction	4
Device Agnostic Conditions & Load Data	7
Pre-Processing	7
Model Building	5
Mini-Challenge: Model Training & Model Evaluation	8
Saving and Loading PyTorch Models	5
Challenge	6

Implement a Logistic Regression Model with PyTorch

Introduction	5
Review Sklearn Titanic Classification	5
Sklearn Python Code	1
Perform PyTorch Titanic Classification	22
Challenge	2
Completed PyTorch Code	8

Explore Neural Network Classification with PyTorch

Review Logistic Regression PyTorch Workflow	9
Load Make Moons Dataset & Pre-processing	12
Define Neural Network Architecture	12
Train and Evaluate Model	12
Visualize Decision Boundary with Probability	3
Challenge: PyTorch Workflow	3

Build a PyTorch Classifier with Non-Linearity

Review Neural Network Classification Without Non-Linearity	1
--	---

Build a Neural Network Classification With Non-Linearity	52
--	----

WEEK 9

158 min.

Challenge: PyTorch Workflows	3
------------------------------	---

Explore Multi-class Classification with PyTorch

Review of "Binary Classification with PyTorch"	11
Step 1: Setup and Prepare Data	12
Step 2: Visualize Data (EDA)	8
Step 3: Define Neural Network Architecture	9
Challenge	1
Solution Videos	16
Solution Notebook and url	1

Tune Hyperparameters and Analyze Fit with PyTorch

Review: Explore Multi-class Classification with PyTorch	5
Create, Preprocess, and Visualize the Spiral Dataset	9
Define Neural Network Architecture	6
Explore Hyperparameter Tuning	13
Explore Underfitting and Overfitting	9
Challenge	1
Solution Video	3

Discover What's New with PyTorch 2.0

Universal Device Setup in PyTorch 2.0	6
---------------------------------------	---

Key Features of PyTorch 2.0	6
Traditional PyTorch 1.0 Vs PyTorch 2.0: torch.compile()	17
Challenge ☒	5
Completed Solution	1

Explore TensorFlow Machine Learning Foundations

Introduction to TensorFlow Tensors	16
------------------------------------	----

WEEK 10

154 min.

Create Tensors with TensorFlow	4
Create Random Tensors with Numpy	8
Challenge ☒	9

Explore TensorFlow Aggregation and Manipulation

Why Shuffle Tensors?	4
TensorFlow Seeds	5
Tensor Attributes	6
Tensor Indexing	3
Changing Tensor Data Types & Tensor Aggregation	8
Tensor Positional Methods	4
Challenge ☒	6

Implement Matrix Multiplication with TensorFlow

Basic Tensor Operation	4
TensorFlow Math Functions	3
Matrix Multiplication Foundations	11

Perform Matrix Multiplication	13
Challenge	10

Reshape, Transpose, and Alter TensorFlow Tensors

Review: Matrix Multiplication	8
Altering Tensors	8
Transpose & Reshape Tensors	5
Tensor Expansion	7
Challenge ☒	17

Squeeze, Encode, and Optimize TensorFlow Tensors

Squeezing Tensors	11
-------------------	----

WEEK 11

167 min.

One-Hot Encoding	7
Numpy => Friend ☒☒	9
GPU & TPU Tensor Optimization	4
Challenge ☒	11

Explore Neural Network Regression with TensorFlow

What is Regression Analysis?	8
Neural Network Architecture	12
Build a Model	14
Challenge ☒	4
Solution Video	9

Build a Simple Regression Model with TensorFlow

Introduction	2
Build a Small Regression Model from Memory	6
Build Model From Scratch	13
☒ Challenge: Improve Model	24

Evaluate Regression Models with TensorFlow

Introduction	3
Regression Challenge	7
Preprocess Data	13
☒ Challenge: Build Model	3
Challenge Solution	18

WEEK 12

Visualize and Evaluate Performance with TensorFlow

157 min.

Introduction	1
Generate Linear Transformation Data	6
Common Evaluation Metrics: MAE, MSE, & Huber	10
Split Data for Train and Test Datasets	10
Define Basic Model Architecture	6
Make Predictions and Evaluate Model	10
☒ Challenge	1
Solution Video	5

Normalize and Feature Scale Data with TensorFlow

Introduction	4
Handle Imports & Load Dataset	6

One-hot Encode & Separate Features and Target	6
Perform Train/Test Split	4
Define Model Architecture	5
Evaluate Model and Visualize Loss	4
What is Normalization and Standardization?	3
☒ Challenge	7

Explore TensorFlow Neural Network Classification

Introduction	4
What is Classification?	6
What is Binary Classification?	7
What is Multi-Class Classification?	5
What is Multi-Label Classification?	6
Classification Code Example	9
☒ Challenge	16

Build a Neural Network Classifier with TensorFlow

Introduction	4
Pseudocode Image Classification	3
Create Circles Dataset & EDA	9

WEEK 13

160 min.

Build, Compile, and Train Model	6
Visualize and Evaluate Model	7
☒ Challenge	1
Solution Video	7
Bonus Video	6

Build a TensorFlow Classifier with Non-Linearity

Introduction	1
Review Non-Linearly Separable Data	5
Create Circles DataSet	9
Create Second Model	11
Create Third Model	7
Create Fourth Model	9
☒ Challenge	1
Solution	7

Evaluate TensorFlow Classification Models

Introduction	1
Review: Learning Rates	8
Adaptive Learning Rates	31
Big Five Evaluation Metrics	4
☒ Challenge	1
Solution Video	4

Explore Multi-Class Classification with TensorFlow

Introduction	1
Compare Binary and Multi-Class Classification	7
Create a Teachable Machine Multi-Class Classifier	7
Review Model Building Steps	3
Load and Explore MNIST Fashion Dataset	16

☒ Challenge

157 min.

11

Tune Multi-Class Classification TensorFlow Models

Introduction	4
Review MNIST Fashion Multi-Class Classifier	8
Load and Visualize Dataset	8
One-Hot Encode Features and Build Model	11
Softmax and Validation Exploration	6
☒ Challenge	10

Explore Multi-Label Classification with TensorFlow

Introduction	3
Binary, Multi-Class, and Multi-Label Classification	14
Start Building a Multi-Label Classifier	13
Build a Sequential Multi-Label Model	9
Evaluate Model	9
☒ Challenge	7

Explore The Fundamentals of Large Language Models

Introduction	4
What is a Large Language Model (LLM)?	5
How do LLMs work?	9
Two Kinds of LLMs: Base and Instruction Tuned	6
System Messages and Tokens	12
Challenge: Connect Google Colab to ChatGPT via OpenAI's API	8

WEEK 14

WEEK 15

158 min.

Build LLM Apps with ChatGPT and the OpenAI API

Introduction	4
Web Chat Interfaces Vs. Programmatic Notebooks	8
Route Queries Using Classification for Different Cases	9
Evaluate Inputs to Prevent Prompt Injections	4
Implement The OpenAI Moderation API	8
Sanitize and Validate Inputs: Injection Attacks	8
Challenge: Filter Inputs with a Chain of Thought Prompt Filter	10

Design Effective Prompts for Large Language Models

Introduction	4
Iterative Prompt Engineering	17
Build a Summarizer for Interesting Topics	12
Implement Supervised Learning Through Inference	7
Challenge: Build The AutoBot ChatBot To Manage Orders	13

Implement LangChain in Language Model Workflows

Introduction	4
Compare Direct API Calls Vs. API Calls Through LangChain	9
Leverage LangChain Templating for Complex Prompts	12
Leverage Power of Templating for DRY Code	9
Challenge	7

Implement LangChain Memory for Autonomous Tasks

Introduction	4
--------------	---

ConversationBufferMemory	9
--------------------------	---

WEEK 16

162 min.

ConversationBufferWindowMemory	5
ConversationTokenBufferMemory	4
ConversationSummaryBufferMemory	9
The Power of Chaining LangChain Components	7
Challenge: Implement LangChain Memory	11

Combine LangChain Components for Coherent Apps

Introduction	5
Chaining in LangChain	6
LLMChain	7
SimpleSequentialChain	6
SequentialChain	7
RouterChain	12
Challenge	7

Build Task-Driven Autonomous Agents with LangChain

Introduction	5
Leverage LangChain Agents	8
Perform math calculation using an Math LLM	6
Use Wikipedia to Find General Information	6
Program using a Python REPL tool	3
Create new custom agents and tooling (BabyAGI)	2
Debugging with LangChain	9
Challenge	6

Use LangChain to Interact with PDFs and Documents

Introduction	6
Retrieval Augmented Generation (RAG) over 2 Skills	6
Document Loaders	7
Document Separation	12

WEEK 17

153 min.

Embeddings	8
Vector Stores	8
Challenge	1

Use LangChain to Chat with PDFs and Documents

Introduction	3
Similarity Search	6
Maximum Margin Relevance	6
ContextualCompressionRetriever + MMR	7
Chat Q&A	15
Challenge	8

Explore Transformer Encoders and Decoders

Introduction	5
What are Transformers?	9
Attention Is All You Need (Optional)	10
Encoders	4
Decoders	5
Encoder-Decoders	4
What is HuggingFace Again?	4

Challenge	1
Solution Video	2

Examine the Fundamentals of HuggingFace

Introduction	4
What is HuggingFace? ☒	6
Models	15
Datasets	9
Spaces	13

WEEK 18

8 min.

ChatGPT Competitor HuggingChat ☒☒	1
Challenge	7