

Certified Entry-Level AI Specialist with Python (PCEI)

Turn Python basics into AI skills that pay. This PCEI certification training prepares you for the Entry-level AI Specialist with Python certification, perfect for AI beginners who want to level up. This PCEI certification training will show you how to analyze datasets, implement basic AI logic, and evaluate the capabilities and limitations of AI models. You'll write real AI-style algorithms in Python, work with small datasets, and explore machine learning, neural networks, and generative AI at a practical level. Most importantly, you'll build portfolio-ready projects you can showcase for employers, all while preparing for the Python AI certification exam.

[CBT Nuggets course material](#) →

WEEK 1

Build Your First Machine Learning Algorithm

Loading CSV Data And Calculating Distances

Completing The Algorithm

Challenge & Solution

Measuring and Improving Performance

Measuring Performance With A Train/Test Split

Calculating An Accuracy Score

Testing Model Variations

Challenge & Solution

Basics of Data Normalization

What is Normalization?

How Does Normalization Work?

Normalizing Data For K-Nearest-Neighbor

Other Normalization Strategies

Challenge & Solution

Creating Reusable Machine Learning Models

Designing Reusable Data Structures And APIs

Creating A Reusable KNN Class

Testing Our Reusable KNN Class

Making Train-Test Splitting Reusable

Challenge & Solution

Basics of Linear Regression

What Is Linear Regression?

Implementing A Simple Linear Regression Model

The Least Squares Method

Challenge & Solution

Introduction to Gradient Descent

What Is Gradient Descent?

The Gradient Descent Process

Implementing Gradient Descent

Challenge & Solution

Basics of Decision Trees

How Do Decision Trees Work?

Implementing Impurity Measures

The Structure Of A Decision Tree

Implementing a Decision Tree

Challenge & Solution

Basics of Clustering Algorithms

Supervised vs. Unsupervised Learning

The K-Means Clustering Algorithm

Implementing K-Means Clustering

Recalculating Centroids

Challenge & Solution

Clustering Algorithms In-Depth

How Good Are Our Clusters Really?

Comparing Clusters With Inertia

Implementing The Silhouette Score Metric

Challenge & Solution

Introduction To Neural Networks

Understanding Neurons

Implementing Neurons in Python

Playing Around With Neurons

Challenge & Solution

Connecting Neurons Into Simple Networks

Visualizing Neuron Decision Boundaries

Connecting Neurons (and XOR Revisited)

The Anatomy Of A Neural Network

Challenge & Solution

Training Individual Neurons With Gradient Descent

The Fundamental Ideas Of Training Neurons

Implementing Gradient Descent For Neurons

Training Neurons Through Multiple Epochs

Challenge & Solution

Applying Neurons To Real-World Data

Predicting Credit Risk (and Sigmoid Saturation)

Normalizing Neuron Training Data

Initializing Neuron Weights And Biases

Challenge & Solution

Working With Non-Numerical Data

Translating Tasks Into Inputs and Outputs

Converting Transactions Into Multi-Hot Vectors

Training Neurons With Multi-Hot Vectors

Challenge & Solution

Creating Layers and Neural Networks

How Hidden Layers Work

Creating a Layer Class

Creating A Neural Network Class

Challenge & Solution

Training Neural Networks With Backpropagation

Understanding Backpropagation

Implementing The Forward Pass

Updating The Neurons

Challenge & Solution

Applying Neural Networks To Real-World Data

Shirt Sizes Revisited

Choosing Hidden Layer Sizes

Saving Trained Neural Networks

Challenge & Solution

Vanilla Python Data-Wrangling - Basics

Transforming Lists and Dictionaries With Comprehensions

Filtering Lists and Dictionaries With Comprehensions

Generator Expressions

Challenge & Solution

Vanilla Python Data Wrangling - Intermediate

Working With Python's Functional Tools

Lambda Expressions

Reducing Data Structures

Challenge & Solution

Basic Statistics For Machine Learning

Understanding Histograms

The Normal Distribution, Mean, and Standard Deviation

The Median, Quartiles, and Percentiles

Challenge & Solution

The Python Collections Module

Counting And Frequency Analysis Tools

Queue Utilities

Other Useful Data Structures

Challenge & Solution

Working With Randomness

Generating Random Values

Generating Random Integers

Random Sampling From Distributions

Challenge & Solution

The Itertools Module For Machine Learning

Infinite Iterator Functions

Using Infinite Iterators in Machine Learning

Combinatorics Functions

Permutations, Combinations, And More

Challenge & Solution

Slicing And Filtering Stream Data

Slicing Iterators

Filtering Iterators

Compressing Iterators

Challenge & Solution

Transforming and Grouping Stream Data

Understanding Python's "Iterable" Vocabulary

Chaining Iterables

Data Alignment With `zip_longest`

Accumulating, Grouping, and Starmapping Data

Challenge & Solution

Creating Simple Data Pipelines

Batch Processing Datasets

Converting Iterables Into Batches

Creating A Full Pipeline

Challenge & Solution

Advanced Data Pipeline Concepts

Functional Composition In Data Pipelines

Implementing Composition

Partial Application In Data Pipelines

Defining Data Destinations

Challenge & Solution

Incremental Computation With Streams

The Problem With "sum", "len", and Other Functions

Calculating Count, Min, and Max "Online"

Calculating Mean And Standard Deviation "Online"

Challenge & Solution

Working With Vectors

What Are Scalars, Vectors, and Matrices?

Defining Vector Operations

Understanding Vector Operations Visually

Challenge & Solution

Working With Matrices

Basic Matrix Terminology

Matrix-Vector Multiplication

Applying Matrix-Vector Multiplication to Machine Learning

Challenge & Solution

Understanding Derivatives and Gradients

The Basic Idea Of Derivatives

Derivative Notation and Computation

From Derivatives to Gradients

Challenge & Solution

Implementing A Neural Network With Matrix Math

Representing Weights And Biases

Forward Passes As Matrix Multiplication

Batch-Processing Inputs With Matrix Multiplication

Challenge & Solution

Training A Neural Network With Matrix Math

The Hadamard and Outer Products

Adjusting Neuron Weights Mathematically

Computing the "Delta" For The Last Layer

Training A Neural Network With Matrix Math

Challenge & Solution

Building Naive Bayes Classifiers

How Does Naive Bayes Work?

Loading and Preparing Data For Naive Bayes

Training A Naive Bayes Classifier

Applying A Naive Bayes Classifier

Challenge & Solution

Basics of Reinforcement Learning

What Is Reinforcement Learning?

A Tiny Grid-World Demo

Letting Agents Choose Actions

Using Rewards to Update Preferences

Challenge & Solution

Going Deeper With Reinforcement Learning

How Do Q-Table Updates Actually Work?

Exploration Strategies

Implementing Epsilon Decay

Implementing Softmax Exploration

Challenge & Solution

Market Basket Analysis Basics

What Is Market Basket Analysis?

The Apriori Algorithm

Dataset Preparation For Market Basket Analysis

Implementing The Apriori Algorithm

Challenge & Solution

Market Basket Analysis In-Depth

Understanding Confidence

Understanding Lift

Generating Rules From Frequent Itemsets

Challenge & Solution

Basics of Natural Language Processing (NLP)

Tokenization, Cleaning, and Word Frequency

Stopword Removal And Lemmatization

Challenge & Solution

Basics of Generative AI And LLMs

What Is Generative AI?

Using Generative AI APIs

Generative AI API Pricing

Challenge & Solution

Working with NumPy Arrays

NumPy Arrays vs. Python Lists

Creating NumPy Arrays

Accessing Array Values

Challenge & Solution

Working With Pandas

How Are Pandas and NumPy Related?

Array Broadcasting and Data Alignment

WEEK 2

Basic Data Exploration with DataFrames

Challenge & Solution

Pandas Data Manipulation

Selecting DataFrame Columns

Selecting DataFrame Rows

Boolean Broadcasting

Boolean Filtering

Challenge & Solution

Data Cleaning For Machine Learning

The Most Common Symptoms Of Dirty Data

Fixing Data Types

Dealing with Missing Values and Outliers

Challenge & Solution

Pandas Data Cleaning Walkthrough

Identifying Potential Dataset Problems

Renaming Dataset Columns

Modifying Cell Values Efficiently

Challenge & Solution

Data Visualization Basics

Displaying Simple Charts

Working With Different Chart Types

Visualizing Datasets

Challenge & Solution

Measuring Model Performance

Why Accuracy Isn't Enough

Precision Vs. Recall

Performance Metrics for Multi-Class Models

Challenge & Solution

Identify Ethical Risks and Safety Concerns in AI

Bias And Discrimination In Machine Learning

Recognize Privacy Concerns

Safety Issues In Machine Learning

Challenge & Solution