

Python Machine Learning

Python Machine Learning Training Course

Prerequisites and Further Training

You should definitely be very well familiar with our [Python AdvancedTopics](#).

Intended Audience

One of the fastest growing areas in the industry at the moment is Data Science and Machine Learning. Join us now!

Further Training

Also have a look at our [Python Data Science Bootcamp](#) .

Course Material

Supplied in electronic format unless specified otherwise

Course Content

- Giving Computers the Ability to Learn from Data
- Building intelligent machines to transform data into knowledge
- The three different types of machine learning
- An introduction to the basic terminology and notations
- A roadmap for building machine learning systems
- Using Python for machine learning

Training Machine Learning Algorithms for Classification

- Artificial neurons – a brief glimpse into the early

history of machine learning

- Implementing a perceptron learning algorithm in Python
- Adaptive linear neurons and the convergence of learning

A Tour of Machine Learning Classifiers Using Scikit-learn

- Choosing a classification algorithm
- First steps with scikit-learn
- Modeling class probabilities via logistic regression
- Maximum margin classification with support vector machines
- Solving nonlinear problems using a kernel SVM
- Decision tree learning
- K-nearest neighbors – a lazy learning algorithm

Building Good Training Sets – Data Preprocessing

- Dealing with missing data
- Handling categorical data
- Partitioning a dataset in training and test sets
- Bringing features onto the same scale
- Selecting meaningful features
- Assessing feature importance with random forests

Compressing Data via Dimensionality Reduction

- Unsupervised dimensionality reduction via principal component analysis
- Supervised data compression via linear discriminant analysis
- Using kernel principal component analysis for nonlinear

mappings

Learning Best Practices for Model Evaluation and Hyperparameter Tuning

- Streamlining workflows with pipelines
- Using k-fold cross-validation to assess model performance
- Debugging algorithms with learning and validation curves
- Fine-tuning machine learning models via grid search
- Looking at different performance evaluation metrics

Combining Different Models for Ensemble Learning

- Learning with ensembles
- Implementing a simple majority vote classifier
- Evaluating and tuning the ensemble classifier
- Bagging – building an ensemble of classifiers from bootstrap samples
- Leveraging weak learners via adaptive boosting

Applying Machine Learning to Sentiment Analysis

- Obtaining the IMDb movie review dataset
- Introducing the bag-of-words model
- Training a logistic regression model for document classification
- Working with bigger data – online algorithms and out-of-core learning

9: Embedding a Machine Learning Model into a Web Application

- Serializing fitted scikit-learn estimators
- Setting up a SQLite database for data storage
- Developing a web application with Flask
- Turning the movie classifier into a web application
- Deploying the web application to a public server

Predicting Continuous Target Variables with Regression Analysis

- Introducing a simple linear regression model
- Exploring the Housing Dataset
- Implementing an ordinary least squares linear regression model
- Fitting a robust regression model using RANSAC
- Evaluating the performance of linear regression models
- Using regularized methods for regression
- Turning a linear regression model into a curve – polynomial regression

Working with Unlabeled Data – Clustering Analysis

- Grouping objects by similarity using k-means
- Organizing clusters as a hierarchical tree
- Locating regions of high density via DBSCAN

Training Artificial Neural Networks for Image Recognition

- Modeling complex functions with artificial neural networks

- Classifying handwritten digits
- Training an artificial neural network
- Developing your intuition for backpropagation
- Debugging neural networks with gradient checking
- Convergence in neural networks
- Other neural network architectures
- A few last words about neural network implementation

Parallelizing Neural Network Training with Theano

- Building, compiling, and running expressions with Theano
- Choosing activation functions for feedforward neural networks
- Training neural networks efficiently using Keras

Duration and pricing

- *Full-time* over 10 days (R19 995)
- *Part-time* over 4 weeks (2 nights per week, 3 hour sessions) (R11995)
- *Part-time* over 8 Saturdays, 3 hour sessions (R11995)
- Please note : For *part-time* courses we do not have a fixed schedule and you will be placed on a waiting list until we get a group of 4+ together. Please book with no dates on the bookings form. This will automatically put you on the waiting list. We will confirm with you as soon as we have a part-time group together.
- [Distance-learning](#) over up to 3 months (R9995)
- International exams are not included in the course price.
- Prices exclude Vat

Certificate

1. Upon completion of this course we will issue you with attendance certificate to certify your attendance and /

or completion of the prescribed minimum examples.

2. You may sit for our competency assessment test and on passing you will obtain our competency certificate.
3. Our competency assessment can be booked and taken by someone who has not attended the course at a cost of R950.

Bookings

You can download the course registration form on our home page or by clicking [here](#)

Brochure

You may download a pdf copy of this page by clicking on the pdf icon at the top of the page.

Questions

Please [email us](#)

Schedule

On the calendar below. If your browser doesn't display the calendar below, please click on [this link](#) or try using [Google Chrome](#), alternatively please enquire via our [Contact Us](#) page.

We are a member of the Python Software Foundation

