

Machine learning basics

This repository contains implementations of basic machine learning algorithms in plain Python (Python Version 3.6+). All algorithms are implemented from scratch without using additional machine learning libraries. The intention of these notebooks is to provide a basic understanding of the algorithms and their underlying structure, *not* to provide the most efficient implementations.

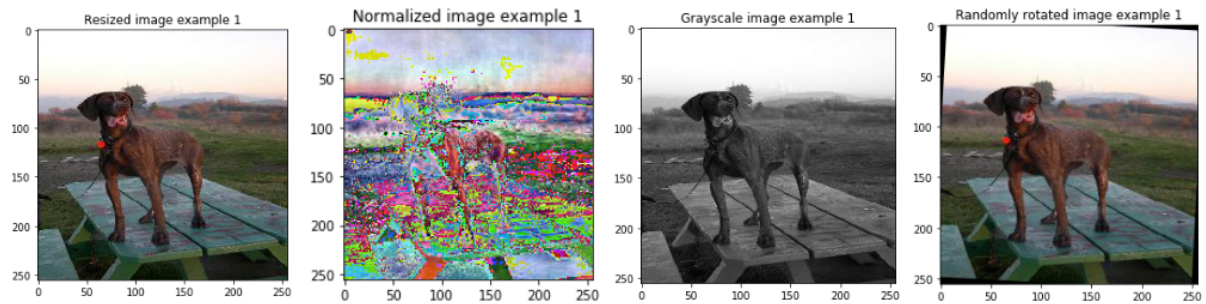
- Bayesian Linear Regression
- Decision tree for classification
- Decision tree for regression
- k-nearest-neighbor
- k-Means clustering
- Linear Regression
- Logistic Regression
- Multinomial Logistic Regression
- Perceptron
- Principal Component Analysis
- Simple neural network with one hidden layer
- Softmax regression
- Support vector machines



Data preprocessing

After several requests I started preparing notebooks on how to preprocess datasets for machine learning. Within the next months I will add one notebook for each kind of dataset (text, images, ...). As before, the intention of these notebooks is to provide a basic understanding of the preprocessing steps, *not* to provide the most efficient implementations.

- Image preprocessing
- Preprocessing a numerical/categorical dataset



Live demo

Run the notebooks online without having to clone the repository or install jupyter: [launch](#) [binder](#).

Note: this does not work for the `data_preprocessing.ipynb` and `image_preprocessing.ipynb` notebooks because they require downloading a dataset first.

Feedback

If you have a favorite algorithm that should be included or spot a mistake in one of the notebooks, please let me know by creating a new issue.

License

See the LICENSE file for license rights and limitations (MIT).