

Machine Learning and Pattern Recognition MC886/MO444

University of Campinas (UNICAMP), Institute of Computing (IC)
Assignment #2, 2018s2, Prof. Sandra Avila

Objective

Explore classification techniques and come up with the best possible model to the problem, avoiding overfitting. In particular, build an object recognition system to accurately classify images using the Fashion-MNIST dataset (<https://github.com/zalandoresearch/fashion-mnist>).

Activities

1. Perform Logistic Regression as the baseline (first solution) to learn the 10 classes in the dataset. Use one-vs-all strategy to build a classification model.
Keep in mind that you should obtain 10 classification models.
2. Perform Multinomial Logistic Regression (i.e., Softmax regression). It is a generalization of Logistic Regression to the case where we want to handle multiple classes. What are the conclusions?
3. Move on to Neural Networks, using one hidden layer. You should implement your solution.
4. Extend your Neural Network to two hidden layers. Try different activation functions. Does the performance improve?
5. Pick your best model and plot the confusion matrix in the test set. What are the conclusions?
6. Prepare a 4-page (max.) report with all your findings. It is UP TO YOU to convince the reader that you are proficient on Logistic Regression and Neural Network, and the choices it entails.

Dataset

Fashion-MNIST is a dataset of Zalando's article images, consisting of a training set of 60,000 examples and a test set of 10,000 examples. Each example is a 28×28 grayscale image, associated with a label from 10 classes.

Dataset Information:

- You should respect the following training/test split: 60,000 training examples, and 10,000 test examples. Avoid overfitting.
- The data is available at:
<https://www.dropbox.com/s/qawunrav8ri0sp4/fashion-mnist-dataset.zip>:
'train' folder (fashion-mnist_train.csv.zip) + 'test' folder (fashion-mnist_test.csv.zip).
- Each training and test example is assigned to one of the following labels: 0 t-shirt/top, 1 trouser, 2 pullover, 3 dress, 4 coat, 5 sandal, 6 shirt, 7 sneaker, 8 bag, 9 ankle boot.
- Each row is a separate image. Column 1 is the class label. Remaining columns are pixel numbers (784 total). Each value is the darkness of the pixel (1 to 255). Dataset was converted to CSV with this script: <https://pjreddie.com/projects/mnist-in-csv>.




Figure 1: Classes in the dataset, each class takes three-rows. Figure reproduced from <https://github.com/zalandoresearch/fashion-mnist>.

Deadline

Tuesday, **October 4th** in the beginning of the class, 7 pm.

Penalty policy for late submission: You are not encouraged to submit your assignment after due date. However, in case you did, your grade will be penalized as follows:

- October 5th 7pm : grade * 0.75
- October 6th 7pm : grade * 0.5
- October 7th 7pm : grade * 0.25

Submission

On the deadline day, bring your 4-page printed report. The template for report is available at <https://www.dropbox.com/s/nc6d89otr8ekvjd/report-model.zip>. Please, print on both sides of the page. The report should be written in Portuguese or English.

Submit a zip file, with the code and the report (PDF file), via Moodle.

This activity is NOT individual, it must be done in pairs (two-person group).