

Machine learning basics

Tom Rochette <tom.rochette@coreteks.org>

November 2, 2024 — [36c8eb68](#)

0.1 Context

0.2 Learned in this study

0.3 Things to explore

- Is it currently possible to batch process in parallel? If so, how is the network updated?

1 Overview

math basics

matrices

vectors

operations (addition, subtraction, multiplication by scalar, matrix/vector and matrix/matrix multiplication, inverse, transpose)

cost function

gradient descent

normal equation

classification (binary vs multiclass)

- one vs all/rest
 - linear/non-linear boundaries
 - under/over-fitting
- reduce the number of features
- regularization
 - neural network
- layer
- bias
- weight
- input/output
- activation
- forward/backward propagation
 - training/test/validation sets
 - diagnosis of learning algorithm
- bias (underfit) vs variance (overfit)
 - learning curve (error based on training set size)

fixes

- get more training examples -> fix high variance
- try smaller sets of features -> fix high variance
- try getting additional features -> fix high bias
- try adding polynomial features -> fix high bias
- try decreasing lambda -> fix high bias

- try increasing lambda -> fix high variance

neural networks architecture

- small neural network -> prone to underfitting
- large neural network -> prone to overfitting

how to prioritize your work

- start with a simple algorithm that can be implemented quickly
 - implement it and test it on cross-validation data
- plot learning curves to decide if more data, features, etc. are likely to help
- do error analysis: manually examine the examples (in the cross validation set) that your algorithm made errors on. See if you spot any systematic trend in what type of examples it is making errors on
- use pareto rule (80/20)
- use numerical evaluation to determine if a change is an improvement or not

skewed classes

- precision/recall
- precision = true positive/# predicted positive = (true+/(true+ + false+))
- recall = true positive/# actual positive = (true +/(true+ + false-))
- f1 score = $2pr/(p+r)$

support vector machine

kernels

- linear
- gaussian

unsupervised learning

- clustering
 - k-means (cluster assignment, move centroid to the cluster's mean location)
 - elbow method (choose the number of clusters automatically)
- dimensionality reduction (data compression)
 - principal component analysis

To avoid overfitting, the number of parameters estimated from the data must be considerably less than the number of data points

- Linear regression: estimate the coordinates of a value
- Logistic regression: answer a yes/no question
- Softmax classification: answer a multiple choice question

2 Procedure

- Define the task
- Obtain data (write a data fetcher)
- Prepare data
- Implement network architecture
 - Define cost, error and optimize computation methods
- Feed data into network

3 Drawing a rectangle over an image

4 See also

5 References

- <http://www.erogol.com/machine-learning-pathway/>