CS6421: Deep Neural Networks

Gregory Provan

Spring 2020 Lecture 2: Classification and Learning

Based on notes from John Canny, Ismini Lourentzou



Machine Learning Basics

- Key tasks of Deep Learning network
 - Classification
 - Learning
- Classification in DN
- Learning in DN
 - Backpropagation

Machine Learning Basics

Machine learning is a field of computer science that gives computers the ability to **learn without being explicitly programmed**



Methods that can learn from and make predictions on data

Training and testing: Supervised Case



• Training is the process of making the system able to learn.

- No free lunch rule:
 - Training set and testing set come from the same distribution
 - Adequacy of functional approximation



Performance

- There are several factors affecting the performance:
 - Types of training provided
 - The form and extent of any initial background knowledge
 - The type of feedback provided
 - The learning algorithms used
- Two important factors:
 - Modeling
 - Optimization



- The success of machine learning system also depends on the algorithms.
- The algorithms control the search to find and build the knowledge structures.
- The learning algorithms should extract useful information from training examples.

Types of Learning

Supervised: Learning labels **y** with a **labeled training** set **x** Example: email *classification* with already labeled emails

Unsupervised: Discover **patterns** in **unlabeled** data **x** Example: *cluster* similar documents based on text

Reinforcement learning: learn to **act** based on **feedback/reward z** Example: learn to play Go, reward: *win or lose*



Anomaly Detection Sequence labeling y=f(x)

 \mathcal{Z}

Tasks of Different Approaches

- Supervised learning $(\{x_n \in \mathbb{R}^d, y_n \in \mathbb{R}\}_{n=1}^N)$
 - Prediction
 - Classification (discrete labels), Regression (real values)

• Unsupervised learning $(\{x_n \in \mathbb{R}^d\}_{n=1}^N)$

- Clustering
- Probability distribution estimation
- Finding association (in features)
- Dimension reduction
- Semi-supervised learning
- Reinforcement learning
 - Decision making (robot, chess machine)

Visualisation of Approaches





Supervised learning

Unsupervised learning



Semi-supervised learning

Machine Learning Methodology

Supervised learning



Machine Learning Methodology

Unsupervised learning



Objectives (Examples)

• Supervised: Low Eout or maximize probabilistic terms

$$error = \frac{1}{N} \sum_{n=1}^{N} [y_n \neq g(x_n)]$$

Ein: for training set
Eout: for testing set
$$Eout(g) \leq Ein(g) \pm O\left(\sqrt{\frac{d_{VC}}{N} \ln N}\right)$$

Unsupervised: Minimum quantization error, Minimum distance, MAP, MLE(maximum likelihood estimation)



Learning Representations

- Supervised learning representations
 - Linear classifier (numerical functions)
 - Parametric (Probabilistic functions)
 - Naïve Bayes, Gaussian discriminant analysis (GDA), Hidden Markov models (HMM), Probabilistic graphical models
 - Non-parametric (Instance-based functions)
 - *K*-nearest neighbors, Kernel regression, Kernel density estimation, Local regression
 - Non-metric (Symbolic functions)
 - Classification and regression tree (CART), decision tree
 - Aggregation
 - Bagging (bootstrap + aggregation), Adaboost, Random forest

Linear Classifiers

• Linear classifier



$$g(x_n) = sign(w^T x_n)$$

where w is an d-dim vector (learned)

- Techniques:
 - Perceptron
 - Logistic regression
 - Support vector machine (SVM)
 - Ada-line
 - Multi-layer perceptron (MLP)

Linear Classification using NN

Using perceptron learning algorithm(PLA)





Training Error rate: 0.10 Testing Error rate: 0.156

Using logistic regression





Training Error rate: 0.11

Error rate: 0.145

Non-Linear Classification

• Non-linear case



- Support vector machine (SVM):
 - Linear to nonlinear: Feature transform and kernel function

Deep Learning Representation

- Functional approximation
 - NN is a universal function approximator
 - Can approximate an arbitrary non-linear function

•We focus on *supervised* learning

Neural Network Basics

- Two main operations
 - Forward propagation
 - Backward propagation (weight adjustment)





Forward Propagation

- Forward Propagation :
 - Sum inputs, produce activation, feed-forward











General View: Forward Propagation

 An Artificial Neuron is a non-linear parameterized function with restricted output range



$$y = f \overset{\mathcal{R}}{\underset{e}{\Diamond}} w_0 + \overset{n-1}{\underset{i=1}{\circ}} w_i x_i \overset{"}{\underset{i=1}{\circ}} w_i x$$

 w_0 also called a bias term (b_i)

Multiple Layers: Artificial neuron



 $\mathbf{x} \qquad \mathbf{y}_1 \qquad \mathbf{y}_2 \qquad \mathsf{VECTORS}$

$$y_1 = f_1(W_1 x + b_1)$$
 $y_2 = f_2(W_2 y_1 + b_2)$

Training: Deep Network



Loss (cost) function

 $\theta^* = \arg \min_{\theta} \sum_{(x,y) \subseteq (X,Y)} J[y, fL(x, \theta_1])$

Summary: Deep Network



Deep Network: family of parametric, non-linear, hierarchical representations

$$\mathbf{y}_{N}(\mathbf{x}, \boldsymbol{\theta}_{1}, \boldsymbol{\theta}_{2}, \dots, \boldsymbol{\theta}_{N}) = f_{N}(f_{N-1}(\dots(f_{1}(\mathbf{x}, \boldsymbol{\theta}_{1}), \boldsymbol{\theta}_{N-1}), \boldsymbol{\theta}_{N})$$

Training: optimize network parameters to minimise loss over training set

$$\theta^* = \arg \min_{\theta} \sum_{(x,y) \subseteq (X,Y)} J[y, fL(x, \theta_{1,\dots,L})]$$

Deep Feed Forward Neural Nets (in 1 Slide (③))



Simple learning procedure: *Back Propagation* (of the error signal)

Supervised Training of Neural Networks

- Use set of training data D
- Adjust weights to optimize classification performance over D

A dataset Fields class 1.4 2.7 1.9 0 3.8 3.4 3.2 0 6.4 2.8 1.7 1 4.1 0.1 0.2 0 etc ... 0



Training the neural network Fields class 1.4 2.7 1.9 0 3.8 3.4 3.2 0 6.4 2.8 1.7 1 4.1 0.1 0.2 0 etc ...



Training dataFieldsclass1.42.71.903.83.43.206.42.81.714.10.10.20etc...0

Initialise with random weights





Present a training pattern



Training data			
Fields			<u>class</u>
1.4	2.7	1.9	0
3.8	3.4	3.2	0
6.4	2.8	1.7	1
4.1	0.1	0.2	0
etc	• • •		

Feed it through to get output



Training data Fields class 1.4 2.7 1.9 0 3.8 3.4 3.2 0 6.4 2.8 1.7 1 4.1 0.1 0.2 0 etc ... 1

Compare with target output







Adjust weights based on error





Present a training pattern



Training data Fields class 1.4 2.7 1.9 0 3.8 3.4 3.2 0 6.4 2.8 1.7 1 4.1 0.1 0.2 0 etc ... 1

Feed it through to get output





Compare with target output





Adjust weights based on error





Repeat this thousands, maybe millions of times – each time taking a random training instance, and making slight weight adjustments

Algorithms for weight adjustment are designed to make changes that will reduce the error

Initial random weights













Training: Summary



Optimize (min. or max.) objective/cost function $e(\theta)$ Generate error signal that measures difference between predictions and target values





Use error signal to change the **weights** and get more accurate predictions Subtracting a fraction of the **gradient** moves you towards the **(local) minimum of the cost function**