


K6312 Information Mining & Analysis

Deep Learning/Deep Neural Networks

Neural Network Representation



Given input x :

$$z^{[1]} = W^{[1]}x + b^{[1]}$$

$(4,1)$ $(4,3)$ $(3,1)$ $(4,1)$

$$a^{[1]} = \sigma(z^{[1]})$$

$(4,1)$ $(4,1)$

$$z^{[2]} = W^{[2]}a^{[1]} + b^{[2]}$$


$(1,1)$ $(1,4)$ $(4,1)$ $(1,1)$

$$a^{[2]} = \sigma(z^{[2]})$$


$(1,1)$ $(1,1)$

Shallow vs Deep

Simple Neural Network



Deep Learning Neural Network




● Input Layer

● Hidden Layer


● Output Layer

End-to-End Learning




From Aporras

Representation Learning in DL




From Deep Learning (Goodfellow)

Representation Matters




Task: Draw a line to separate the **green triangles** and **blue circles**.




We want to project the data into the **new** feature/vector space that data is **linearly separated**


“Trick” in Deep Learning



Low-dim, Original Space




High-dim, **Linearly Separated** Space



Softmax Classifier
(Linear Model)

Why Deep Learning



From Andrew Ng

Deep Learning

- Deep learning is suitable for big data
- Deep learning is able to address unstructured data, which can learn representations from these unstructured data.

Structured

- Structured: Table (Matrix) or Tensor

Player	Height (inches)	Weight (pounds)	Position
Player 1	76	225	C
Player 2	75	195	PG
Player 3	72	180	SF
Player 4	82	231	PF

Features (points to Height and Weight columns)

Labels (points to Position column)


Feature Values (points to 225, 195, 180, 231)

Data Sample (points to the entire row for Player 4)

Unstructured

- The original data can not be stored in an “table”
- More abstract, more fuzzy, and more high-dimensionality

Images



Audio



Video



Text

Content

This module provides students a deep overview of various advanced machine learning techniques applied to business analytics tasks. The focus of this course will be the key and intuitive idea behind machine learning models and hands-on examples instead of theoretical analysis. The tentative topics include machine learning pipeline, unsupervised learning, structure learning, Bayesian learning, deep learning and generative models. The programming languages used will be Python.

Environment around agent





Deep Learning

- Deep learning is a subfield of machine learning
- Most machine learning methods work well because of high-quality feature engineering/representation learning.
- Deep learning is an **end-to-end** structure, which supports automatic representation learning
- Different network structures: CNN, RNN, LSTM, GRU, Attention model, etc


Applications of DL

Deep Learning for Speech

The first real-world tasks addressed by deep learning is speech recognition



Recurrent Neural Network for Sequential Data




Deep Learning for Computer Vision


- Computer vision may be the most well-known breakthrough of DL.
- ImageNet Classification with Deep Convolutional Neural Networks.



ImageNet Scoreboard



Convolutional Neural Network For Image



Extracting useful features of data


Perform a ML task (like classification based on the vectorized data)



https://www.youtube.com/watch?v=FwFduRA_L6Q

Deep Learning For Data Generation

Given training data, generate new data samples from same distribution



Examples of Photorealistic GAN-Generated Faces.

Generative Models

- Given training data, generate new samples from same distribution




Training data $\sim p_{\text{data}}(\mathbf{x})$



Generated samples $\sim p_{\text{model}}(\mathbf{x})$

- Want to learn $p_{\text{model}}(\mathbf{x})$ similar to $p_{\text{data}}(\mathbf{x})$
- Address density estimation, a core problem in ? learning





2014



2015



2016



2017




2018

GAN progress on face generation from Ian Goodfellow

DL/NN is not New

Deep Learning Timeline



Why is Deep Learning Powerful Now?

- Feature engineering require high-level expert knowledge, which are easily over-specified and incomplete.
- Large amounts of training data
- Modern multi-core CPUs/GPUs/TPUs
- Better deep learning 'tricks' such as regularization, optimization, transfer learning etc.

When DL may not Work

You need to get off your non-motor vehicle when u pass the pedestrian crossing.



detected offender

The Challenge of Deep Learning

- Ask the right question and know what the answer means:
Image classification is not scene understanding.
- Select, collect, and organize the right data to train on:



Efficient Teaching/Efficient Learning

- Humans can learn from few examples
- DL/machine require thousands/millions of examples
 - Data augmentation



Limitations

- DL always requires a large amount of annotated data



14 million


Pre-training, Transfer Learning, Data Augmentation

- Generalization capability is low, e.g. the model that perform well on benchmarked datasets fail badly on real world images




- Easily got attacked by random, tiny noise
- How to explain such huge black box

Attack Machine Learning




Attack Machine Learning

Adversarial Examples




Open AI

These stickers made an artificial-intelligence system read this stop sign as 'speed limit 45'.



Scientists have evolved images that look like abstract patterns — but which DNNs see as familiar objects.



Why deep-learning AIs are so easy to fool

Three points behind Successful ML Application

- Deep algorithms, i.e., deep learning
- Strong supervision information (data with high quality labels)
- Stable learning environment



Zhihua ZHOU

Limitations of DL

Key Takeaways

- Neural Network is: 1 linear transformation 2 non-linear activation
- Gradient Descent plus Back-Propagation is used to find the model parameters of neural networks
- Deep learning: neural network with a deep structure (many layers)
- Deep learning is the method which tries to learn features by the model itself without human efforts