

Course Notes: Deep Learning for Visual Computing

Peter Wonka

August 28, 2021


Contents

1	Introduction to Deep Learning	3
1.1	Why is deep learning so popular now?	4
1.2	Why is deep learning so effective now?	6
1.3	Where is deep learning in the AI hierarchy	7
1.4	What are major topics in AI?	8
1.5	What are typical application areas listed as subfields of AI?	9
1.6	One possible approach to structure Machine Learning	10
1.7	LeCake (analogy from LeCun)	11
1.8	What are traditional Machine Learning tools?	13
1.9	What are major conferences in AI?	15
1.10	What are application examples of deep learning in industry?	16
1.11	Internet Companies	18
1.12	AI Companies	19
1.13	Online Courses about Deep Learning (Outdated)	20

1 Introduction to Deep Learning

1.1 Why is deep learning so popular now?

- It works better than all traditional machine learning algorithms as the data size increases
- Significant improvements on problems where large datasets are available or have been collected




1.2 Why is deep learning so effective now?

- Availability of data
- Hardware
- Community (blogs, github, arXiv, ...)
- Tools (Tensorflow, PyTorch, Keras, ...)
- Investment

1.3 Where is deep learning in the AI hierarchy

- **Deep Learning** (DL) is part of **Machine Learning** (ML)
- **Machine Learning** is part of **Artificial Intelligence** (AI)



1.4 What are major topics in AI?

- **Machine Learning**
- Search
 - A* search
- Planning
- Knowledge Representation
- Game Theory
- Probabilistic Graphical Models
- Logic
- Book: Artificial Intelligence: A Modern Approach, Stuart Russell and Peter Norvig

1.5 What are typical application areas listed as subfields of AI?

- Data mining
- Computer vision
- Robotics
- Natural language processing
- (Computer graphics)
- (Bio-informatics)

1.6 One possible approach to structure Machine Learning

- Supervised Learning vs. Unsupervised Learning vs. Reinforcement Learning
- Regression vs. Classification

1.7 LeCake (analogy from LeCun)

- “Pure” Reinforcement Learning (cherry on top of the cake)
 - The machine predicts a scalar reward given once in a while
 - A few bits for some samples
- Supervised Learning (icing on the cake)
 - The machine predicts a category or a few numbers for each input predicting human-supplied data
 - 10 – 10,000 bits per sample
- Self-Supervised Learning (cake genoise, main part)
 - The machine predicts any part of its input for any observed part
 - E.g., predicting future frames in videos
 - Millions of bits per sample

- Slide from LeCun showing the cake

▶ **“Pure” Reinforcement Learning (cherry)**


- ▶ The machine predicts a scalar reward given once in a while.
- ▶ **A few bits for some samples**

▶ **Supervised Learning (icing)**

- ▶ The machine predicts a category or a few numbers for each input
- ▶ Predicting human-supplied data
- ▶ **10 → 10,000 bits per sample**

▶ **Self-Supervised Learning (cake génoise)**

- ▶ The machine predicts any part of its input for any observed part.
- ▶ Predicts future frames in videos
- ▶ **Millions of bits per sample**



1.8 What are traditional Machine Learning tools?

- **Neural Networks**
- Mixture models and EM (k-means clustering)
- knn classifier
- Decision trees
- Support vector machines
- Logistic regression
- Gaussian process regression
- Graphical models
- Kernel Methods
- PCA, ICA, ...
- Sampling, MCMC

- Markov Models
- Boosting
- Book: Pattern Recognition and Machine Learning, Bishop

1.9 What are major conferences in AI?

- AI: AAAI, IJCAI
- Computer Vision: ECCV, ICCV, CVPR
- Data Mining: SIGKDD
- Computer Graphics: SIGGRAPH, SIGGRAPH ASIA
- Machine Learning: ICML, NeurIPS
- Deep Learning: ICLR
- NLP: ACL, EMNLP, NAACL (not sure)

1.10 What are application examples of deep learning in industry?

- Search engines
- Advertisement selection
- Spam filter
- Voice recognition on mobile phones
- Movie or product recommendation
- Credit card fraud detection
- Language translation
- Self-driving cars, boats, planes, trucks, ... ?
- Facebook
 - users upload 2 billion photos per day
 - within 2 seconds each photo passes through 4 convolutional networks

- image tagging: wedding, birthday party, landscape, indoor scene, dog, breed of dog, sail boat, ... (few thousand tags)
- face recognition for automatic tagging of friends
- captions for visually impaired
- detect objectionable content (violence, pornography)

1.11 Internet Companies

- Some decades ago internet companies were a major trend
- Shopping mall + website \neq internet company
- Internet company is not about if you have a website or not, but if you organized your company to do things that the internet lets you do really well
 - A / B testing
 - Short shipping times
 - You can launch a new product every day
 - Shopping mall is redesigned maybe once a year
 - Push decision making down to engineers or product managers
 - decisions are too complex to be done by upper level management

1.12 AI Companies

- Company + Neural Networks \neq AI company
- What are AI companies focused on
 - Strategic data acquisition
 - Unified data warehouses
 - Spotting pervasive automation opportunities

1.13 Online Courses about Deep Learning (Outdated)

- Stanford's Probability and Statistics
- MIT's Linear Algebra
- Stanford's CS231n: Convolutional Neural Networks for Visual Recognition
- Fastai's Practical Deep Learning For Coders
- Stanford's Natural Language Processing with Deep Learning
- Coursera's Machine Learning
- Coursera's Probabilistic Graphical Models
- DeepMind's Introduction to Reinforcement Learning
- UC-Berkeley's Full Stack Deep Learning
- Kaggle's How to Win a Data Science Competition