DS265 Jan 3;1

Deep Learning for Computer Vision

Instructor

R. Venkatesh Babu
Email: venky@iisc.ac.in

Teaching Assistant

Ram Prabhakar
Email: kramprabhakar@gmail.com

Department: CDS

Course Time: Tue., Thu., 11:30 - 1:00 PM

Lecture venue: CDS 102

Detailed Course Page:

Announcements

Brief description of the course

In the recent years, Deep Learning has pushed to boundaries of research in many fields. This course focuses on the application of Deep Learning in the field of Computer Vision. The first half of the course formulates the basics of Deep Learning, which are built on top of various concepts from Image Processing and Machine Learning. The second half highlights the various flavors of Deep Learning in Computer Vision, such as Generative Models, Recurrent Models, and Deep Reinforcement Learning Models.

Introduction

While Computers were invented for 'compute' , the desire to have machines which can 'think' has been an aspiration for researchers and philosophers for many decades. Although our current machines crunch numbers at a super human level, they have traditionally been very poor at problems that humans solve intuitively, like walking, understanding spoken words, recognizing objects in an image etc. With the advent of Deep Learning,

it seems machines might finally meet human-level performance (or even surpass them!), in such intuitive tasks.

What is Deep Learning?

Traditional Artificial Intelligence has been successful in solving abstract and formal tasks, i.e. tasks that can be completely described by a brief list of formal rules, such as 'chess', even as early as 1997, when IBM's Deep Blue defeated chess world champion Gary Kasparov. Based on the rules of the task, the problem was solved using methods such as 'Tree-Search' or 'Dynamic Programming'. However, these traditional machines/algorithms failed in solving subjective and intuitive task, like image/language comprehension.

An approach taken to avert 'hard-coding' of rules in these machines, was Machine Learning, where machines acquired their own 'rules' by extracting patterns from raw data. The approach consisted algorithms like Naive Bayes, and Logistic Regression, which helped solve subjective task such as detection of spam-emails. However, these methods were strongly dependent on the representation of the data provided. In most of the Intuitive tasks, such as recognizing objects in an image, it is difficult to comprehend what data representation should be used.

For solving such tasks, a different paradigm of 'learning' was required, which could build the right representation of data using multiple compositions of data representations, where each representation was learnt by extracting patterns in raw data. This is where Deep Learning plays a major role. Deep Learning is a type of machine learning which enables machines to exactly do this.

Formally: Deep Learning consist of Machine Learning paradigms, which allow machines to learn from

experience, and understand the world in terms of a hierarchy of concepts, where each concept is defined in terms of its relation to other simpler concepts. This 'Concept-graph' can consist of complex concepts built on top of many simple concepts, which again might be built on top of even simpler concepts, leading to a graph with many layers, i.e. a graph which is "Deep".(Hence the name!)

Origins of Deep Learning

While "Deep Learning" has been portrayed as a new technology, its origins dates back to as early as the 1940s. Artificial Neural Networks(ANNs), the predecessors of 'Deep Learning', were first created in the early 1950s and were inspired from Computational models for Biological Learning. ANNs were simple linear models motivated from a nonscientific perspective, and introduced the general principle of learning multiple levels of compositions.

While the idea of composing many layers of ANNs, to form a 'Deep' model is not new, multiple issues such as lack of training methods, lack of data, and insufficient computational power led to failure of such attempts for many years. In 2012, in an image recognition challenge(ILSVRC), a 'Deep Learning' based model was able to surpass all competing AI systems by a huge margin. This lead to the current surge in popularity of 'Deep Learning'. 'Deep Learning' has impacted almost every field of science, and led to many technological breakthroughs such as language translation, speech recognition, scene understanding etc.

Summary

This course aims to provide strong foundation in Deep Learning for student-researchers who are deeply interested in Machine Learning/Artificial Intelligence. As this field has massive industrial impact, students

would be trained not only for theoretical understanding, but also for practically implementing Deep learning based AI systems.

Prerequisites

Primary crucial prerequisites : Machine Learning and Computer Vision/Image Processing

Secondary Prerequisites(familiarity preferred): Statistics and Linear Algebra.

Syllabus

The Course is divided into 2 parts,

Part I : Introduction to CNNs

1) Introduction to Deep Learning and Computer Vision

2) Feed Forward Neural Networks

3) Introduction to CNNs

4) Optimization for training Deep neural networks

5) Deep Neural Networks

6) Tricks for Improving the Learning

Part II : Advanced Topics in Deep Learning

1) Introduction to DL packages/ Important architectures

2) Visualizing CNNs

3) Recurrent Neural Networks

4) Generative Modelling using Deep networks

5) Deep Reinforcement Learning

6) Invited Talks from Researchers in Industry.

Course outcomes

1) Thoroughly Understanding the fundamentals of Deep Learning.

2) Gaining knowledge of the different modalities of Deep learning currently used.

3) Gaining Knowlegde about State-of the art models and Other Important Works in recent years.

4) Learning the skills to develop Deep Learning based AI Systems(Use of Multiple packages etc.)

Grading policy

1) Class Project: 60%

2) Tutorials/Assignments: 30%

3) Data Science Challenge (Using CNNs): 10%

Assignments

5 Common tutorials

1 project tutorial (milestones in project)

Resources

There is no prescribed textbook for this course. However, We hope that the following books would be very useful guides along the course:

1) Deep Learning by Ian Goodfellow, Yoshua Bengio, and Aaron Courville

2) Neural Networks and Deep Learning by Michael Nielson

Additional Resources would be released along with the Lectures.