Deep Learning in Industrial Systems

Class 1: An Overview of Deep Learning Technologies and Applications

May 8, 2017

Louis W. Giokas

DesignNews

CEC CONTINUING EDUCATION CENTER

Presented by: Digi-Key ELECTRONICS

This Week's Agenda

Monday	Overview
Tuesday	Neural Networks
Wednesday	Inspection and Anomaly Detection
Thursday	Robotics Applications
Friday	Future Trends in Deep Learning

Course Description

- Deep learning generally refers to a set of technologies, primarily Artificial Neural Networks (ANN), that are capable of processing unstructured, as well as structured, data.
- Deep learning is especially important in computer vision applications as well as audio and other types of complex signals.
- In this course we will look at the technologies that make up Deep Learning concentrating on Industrial Applications.

Today's Agenda

- Deep Learning Overview
- Implementation Technologies
- Applications
- Conclusion/Next Class

Deep Learning Overview

- In this class we will look at what Deep Learning is, where it fits in the world of Computer Science. We will also introduce some implementation technologies and applications which we will then go into more detail on during the remainder of the week.

Deep Learning Overview

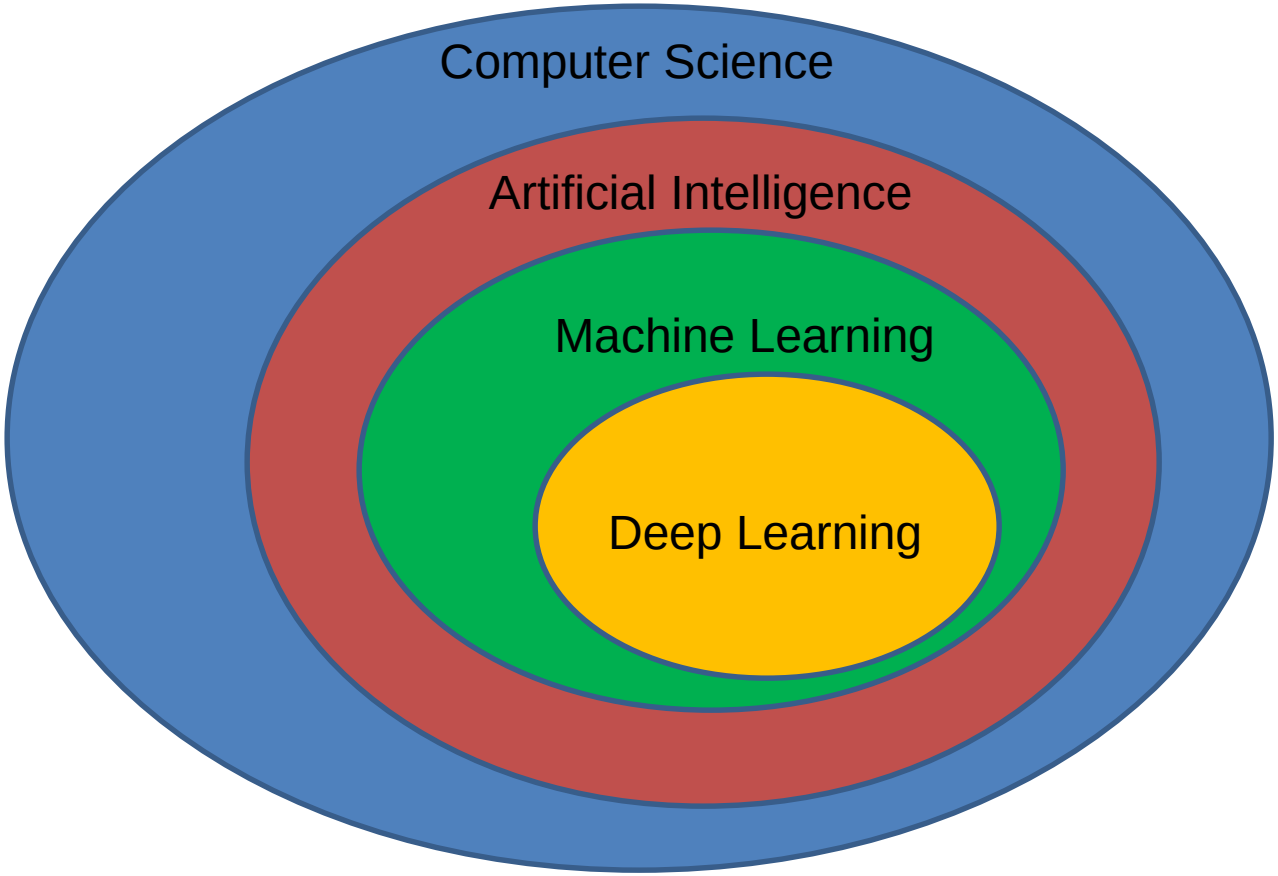

Deep Learning Overview

- Artificial Intelligence (AI):
  - Goal of AI:
    - "To have computers do things, that, if people did them, we would consider intelligent." – Henry Lieberman, MIT
  - Classes of AI
    - Symbolic: Inference, Search, Symbolic Programming, Rule based systems
    - Sub-symbolic: Support Vector Machines (SVM), Linear Discriminant Analysis (LDA), Bayesian Learning, Neural Networks/Deep Learning

Deep Learning Overview

- Machine Learning
  - Learning from the data
  - Statistical Techniques
  - Supervised:
    - Have predictor and response variables. The predictors are used to train the data.
  - Unsupervised
    - No predictors! Directly infer joint probability distribution from the data.

Deep Learning Overview

- Machine Learning Techniques
  - Regression
    - Additive Models
    - Tree Based Models
    - Splines
    - Logistic
  - Support Vector Machines (SVM)
  - K-means Clustering
  - Principle Components Analysis
  - Linear Discriminant Analysis

Deep Learning Overview

- Deep Learning
  - Systems that mimic how the “brain” learns
  - Neural Networks (NN), sometimes called Artificial Neural Networks (ANN)

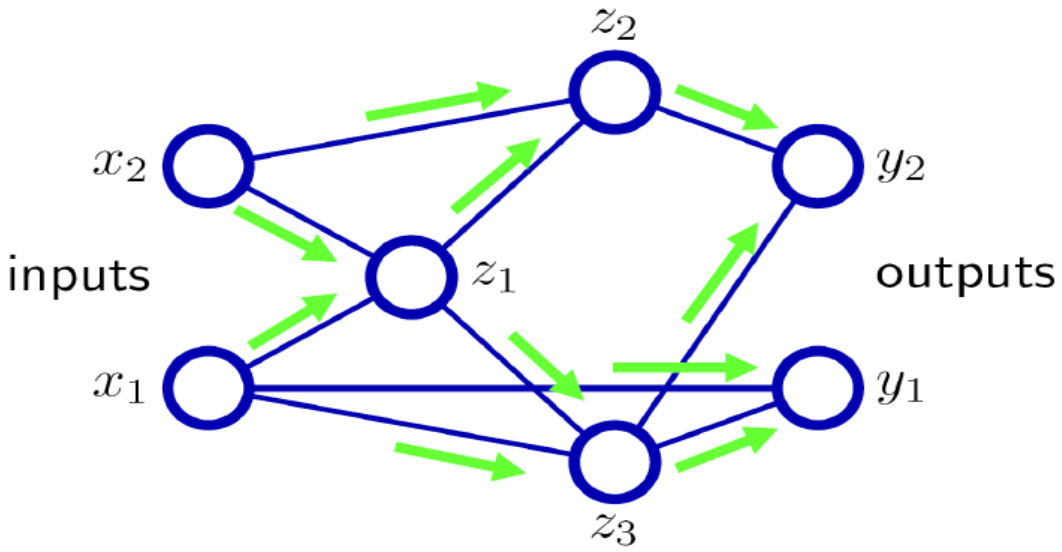

Deep Learning Overview

- Deep Belief Networks

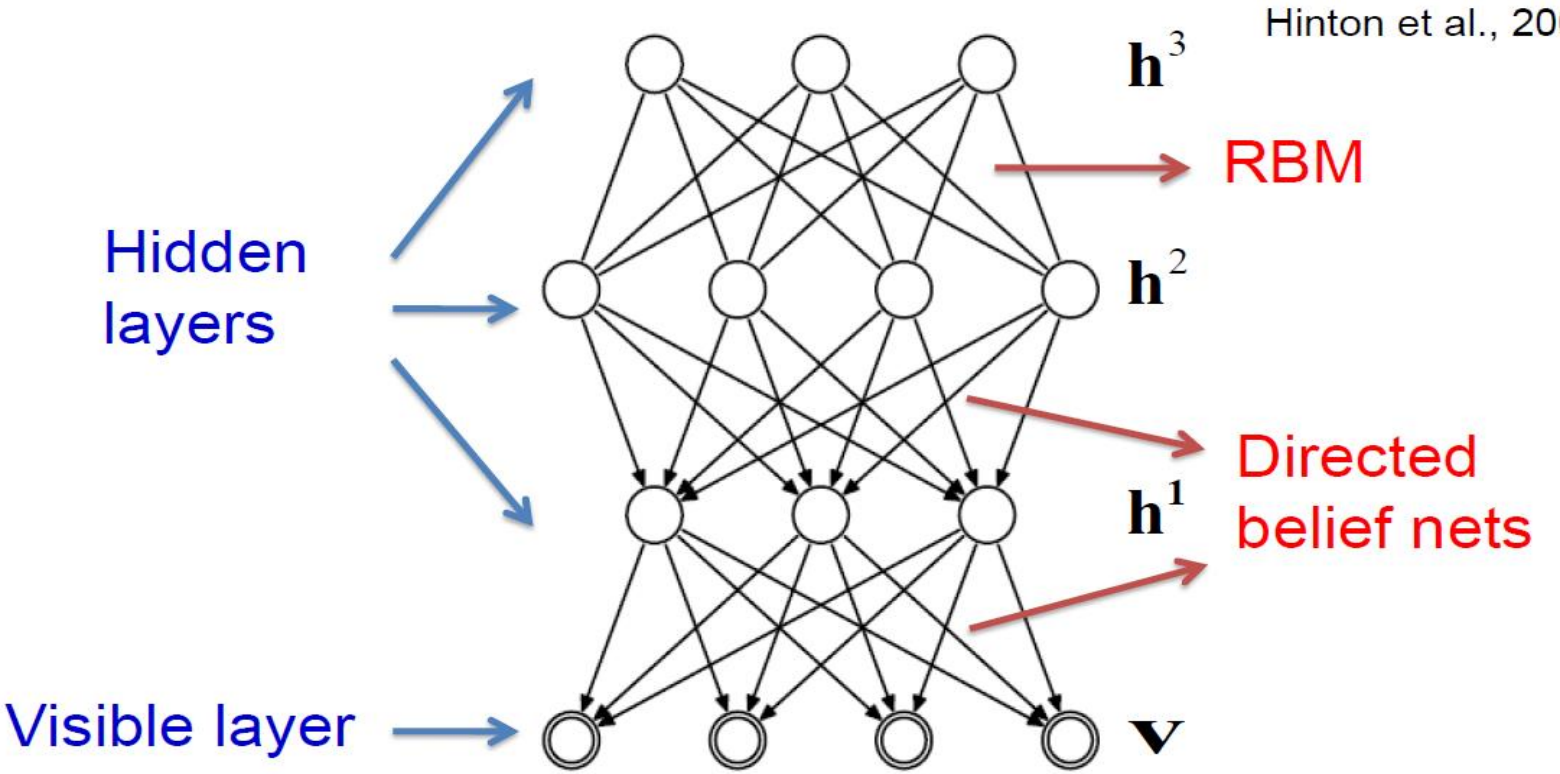

$$P(\mathbf{v},\mathbf{h}^1,\mathbf{h}^2,\ldots,\mathbf{h}^l) = P(\mathbf{v}\,|\,\mathbf{h}^1)P(\mathbf{h}^1\,|\,\mathbf{h}^2)\ldots P(\mathbf{h}^{l-2}\,|\,\mathbf{h}^{l-1})P(\mathbf{h}^{l-1},\mathbf{h}^l)$$

Deep Learning Overview

- Many of the techniques give much better accuracy than alternative machine learning technologies
  - Two layer Boltzman Network
    - Error rate 0.95% compared with a SVM at 1.4%
  - Image recognition (NORB Dataset)
    - 7.2% compared to 11.6% for a SVM and 22.5% for logistic regression

Deep Learning Overview

- The transition to “Deep Learning” generally has to do with the scale of the data and the scale of the neural nets
  - Think Big Data!
- There is also an evolution in the complexity of the neural net
  - Different layers have well defined tasks
  - Many different structures in use
  - Training is both unsupervised and supervised

Implementation Technologies

- Neural nets are typically implemented in software!
  - This is true even though new hardware implementations are evolving
  - Many packages available
    - Some are open source, such as H2O.ai
      - Works in R, Python and other environments
  - Software implementations are often distributed
  - Use of new underlying technologies such as TensorFlow (data graph processing)

Implementation Technologies

- Generally run on traditional CPU architectures
- Recent trend is to program onto GPU architecture
  - Massive parallelism is exploited
  - Good match to the massively parallel nature of neural nets and associated technologies
  - GPUs run with traditional CPUs to form a complete system

Implementation Technologies

- Emerging trends
  - ASICs
    - Develop ASICs to directly implement the Neural Net
    - IBM SyNASPE Chip
      - 1M programmable neurons
      - 256M programmable synapses
      - 4,096 neurosynaptic cores
    - Google is working on implementing TensorFlow in a chip
  - FPGAs
    - Direct RTL implementation of algorithms

Applications

- The applications of Deep Learning are many and varied
- Some are already being used in applications we use today on the Internet
  - Search
  - Image search
  - Voice recognition

Applications

- Other applications
  - Network Security Risks
  - Medical Diagnosis
  - Trading Signals
- The common thread here is that these applications consist of analyzing complex signals, often time varying in nature

Applications

- Industrial
  - Industrial applications may include some aspects of the previously mentioned applications, but are more specific
  - Inspection: image recognition with specific application to a part or assembly
  - Anomaly detection: applied to tools, proecsses
  - Robotics: imaging and planning/control

Conclusion/Next Class

- Deep learning is a type of Machine Learning
- Utilizes the basic technology of Neural Networks
- Software implementations now being augmented with hardware
- Applications go beyond industrial applications, but have a common thread
- Net time we will talk about Neural Networks in detail, including implementation