Introduction to Deep Learning

Sargur N. Srihari srihari@cedar.buffalo.edu

This is part of lecture slides on <u>Deep Learning</u>: http://www.cedar.buffalo.edu/~srihari/CSE676

Topics

- Artificial Intelligence Paradigms
- Origins of Deep Learning
- Representation Learning
- Applications in Vision and NLP
- Depth in Deep Learning
- Generative models in Deep Learning
 - VAE
 - GAN

Summary of AI Models



Shaded boxes indicate components that can learn from data

AI Paradigm Shift

- Physics paradigm shift
 - Newtonian Physics





- Cannot explain black-body radiation
- Quantum Mechanics
- Al paradigm shift
 - Knowledge-based systems
 - Cannot perform simple recognition tasks
 - Simple machine learning methods
 - Cannot perform complex recognition tasks
 - Deep Learning methods

Beginnings of Deep Learning



Perceptron Invented at Calspan Buffalo, NY Rosenblatt, Frank, The Perceptron--a perceiving and recognizing automaton. Report 85-460-1, 1957 Cornell Aeronautical Laboratory





Machine Learning

Minsky and Papert dedicated book to him Minsky M. L. and Papert S. A. 1969 *Perceptrons*, MIT Press Srihari

Designing right set of features

- Useful feature for speaker ID: size of vocal tract – Provides strong clue for: man, woman or child
- Difficult to know what set of features are good for detecting a car in photographs
 - Presence of wheel as a feature
 - Has a simple geometric shape but difficult to describe in terms of pixel values
 - Shadows, glare from metal parts, fender or object in front obscures





Representation Learning

- Solution: use ML to not only learn mapping from representation to output but representation itself
 - Better results than hand-coded representations
- Allows AI systems to rapidly adapt to new tasks
 - Designing features can take great human effort
 - Can take decades for a community of researchers

Feature Learning for Classification

- Function to map pixels to object identity is complicated
- Series of hidden layers extract increasingly abstract features
- Final decision made by a simple classifier



Unsupervised Representation Learning

- Autoencoder:
 - Quintessential example of representation learning
 - Encoder:
 - Converts input into a representation with nice properties
 - Decoder:
 - Converts the representation back to input



Deep Style

New designs from representation



Characteristics of Deep Learning

- 1. A type of Machine Learning that improves with experience and data
- 2. Only viable approach to building AI systems in real-world environments
- 3. Obtains its power by a nested hierarchy of concepts
 - each concept defined by relationship to simpler concepts
 - More abstract representations computed in terms of less abstract ones

¹¹Natural Language Processing

- Training Data
- Word-to-vec
 - One-hot vector mapped to vector of 300
- Word embedding
 - Similar words are close together





Depth in Deep Learning

- Perspective 1
 - Learning the right representation
- Perspective 2
 - Depth allows computer to learn a multi-step computer program
 - Each layer is state of computer memory after executing another layer of instructions
 - More depth means more instructions in sequence
 - Sequential instructions offer greater power since later instructions can refer back to results of earlier ones

Measuring Depth of a Model

- View 2 (Computational Graph View)
 - No. of instructions to be executed
 - Longest path through the flowchart
 - Two equivalent computer programs will have different lengths depending on programming language



Measuring Depth of a Model

- View 1 (Deep Probabilistic Model View)
 - Depth of graph describing how concepts are related to each other
 - Each concept may have a more complex underlying graph
 - Simpler concepts may be refined based on more complex concepts
 - Eye is more complex than face (which consists of two eyes)



Which depth is correct?

- No single correct value for depth of architecture just as no single correct length of a program
- Deep learning differs from traditional ML in that models involve
- (a) composition of learned functions
- (b) learned concepts

Depth versus Width

- Going deeper makes network more expressive
 It can capture variations of the data better.
 - Yields expressiveness more efficiently than width
- Tradeoff for more expressiveness is increased tendency to overfit
 - You will need more data or additional regularization
 - network should be as deep as training data allows.
 - But you can only determine a suitable depth by experiment.
 - Also computation increases with no. of layers.

Very Deep CNNs

CNNs with depth 11 to 19

ConvNet Configuration										
A	A-LRN	B	C	D	E					
11 weight	11 weight	13 weight	16 weight	16 weight	19 weight					
layers	layers	layers	layers	layers	layers					
input (224×224 RGB image)										
conv3-64	conv3-64	conv3-64	conv3-64	conv3-64	conv3-64					
	LRN	conv3-64	conv3-64	conv3-64	conv3-64					
maxpool										
conv3-128	conv3-128	conv3-128	conv3-128	conv3-128	conv3-128					
		conv3-128	conv3-128	conv3-128	conv3-128					
maxpool										
conv3-256	conv3-256	conv3-256	conv3-256	conv3-256	conv3-256					
conv3-256	conv3-256	conv3-256	conv3-256	conv3-256	conv3-256					
			conv1-256	conv3-256	conv3-256					
					conv3-256					
		max	pool							
conv3-512	conv3-512	conv3-512	conv3-512	conv3-512	conv3-512					
conv3-512	conv3-512	conv3-512	conv3-512	conv3-512	conv3-512					
			conv1-512	conv3-512	conv3-512					
					conv3-512					
		max	pool							
conv3-512	conv3-512	conv3-512	conv3-512	conv3-512	conv3-512					
conv3-512	conv3-512	conv3-512	conv3-512	conv3-512	conv3-512					
			conv1-512	conv3-512	conv3-512					
					conv3-512					
	maxpool									
FC-4096										
FC-4096										
FC-1000										
soft-max										

Depth increases from left (A) to right (E) as more layers are added (the added layers are shown in bold)

Convolutional layer parameters denoted "conv (receptive field size) –(no. of channels)"

ReLU activation not shown for brevity

Table 2: Number of parameters (in millions).									
Network	A,A-LRN	В	C	D	E				
Number of parameters	133	133	134	138	144				

Srihari

Generative Model

- Variational autoencoder (VAE) is a *generative model*
 - i.e., able to generate fake samples that look like samples from training data
 - With MNIST data, these fake samples would be synthetic images of handwritten digits
- VAE provides us with a space, the latent space, from which we can sample points
 - Any of these points can be decoded into a reasonable image of a handwritten digit

Standard Autoencoder



A standard autoencoder trained on MNIST digits may not provide a reasonable output when a V image is input

Srihari

Principle of a VAE

A standard normal distribution



This is how we would like points corresponding to MNIST digit images to be distributed in the latent space How well we reconstruct the input is done by a variational measure: KL divergence

Decoder of a VAE



3 s are in first quadrant, 6 s are in third quadrant

Encoder of a VAE



MNIST Variational Autoencoder



Structure of Latent Space



Decoder expects the latent space to be normally distributed Whether the sum of distributions produced by the encoder approximates a standard Normal distribution is measured by the KL divergence

VAE Training

- Due to random variable between input and output it cannot be trained using backprop
- Instead, backprop proceeds through parameters of the latent distribution
 - Called reparameterization trick

$$N(\mu,\Sigma) = \mu + \Sigma N(0, I)$$

where the covariance matrix $\boldsymbol{\Sigma}$ is diagonal



Srihari

Conditional VAE



Do nets have Human-level understanding?

 In many cases, neural networks have begun to reach human level performance when evaluated on an i.i.d. test set

- Have they reached human level understanding?

- To probe the level of understanding we can probe examples that model misclassifies
 - Even neural networks that perform at human level accuracy have a 100% error rate on examples intentionally constructed!

Adversarial examples

- An optimization procedure is used to search for an input x' near data point x such that the model output is very different at x'
 - In many cases, x' can be so similar to x that a human observer cannot tell the difference between the original example and the adversarial example
 - But the network makes a highly different prediction

Adversarial Example Generation

We add to x an imperceptibly small vector

Its elements are equal to the sign of the elements of the gradient of the cost function wrt the input. It changes Googlenet's classification of the image



y ="panda" with 58% confidence

y ="nermatode" With 8.2% confidence y ="gibbon" With 99% confidence

Generative Adversarial Network

 GANs are a way to make a generative model by having two neural networks compete with each other



The discriminator tries to distinguish genuine data from forgeries created by the generator

The generator turns random noise into imitations of the data, in an attempt to fool the discriminator