Deep Learning Wrap-Up

Andrea Danyluk

April 24, 2017

Parts of this, as well as parts of the earlier neural net slides, were adapted from Tom Mitchell. Other parts draw from *Deep Learning* by Goodfellow, Bengio, and Courville

Announcements

- AI/Ethics discussion on Wednesday

Today's Lecture

- Deep Learning Wrap-up
- Discussion: Do deep convolutional nets need to be deep and convolutional?

What is Deep Learning?

- Represents the world as a nested hierarchy of concepts
  - Each concept defined in relation to simpler concepts
  - More abstract representations computed in terms of less abstract ones
- An artificial neural network with many layers
- Success generally not due to simply to the fact that they have many layers
  - Autoencoding
  - Convolution
  - Recurrence

Autoencoders

- Trained to copy their input to their output
  - Typically not interested in the decoding piece
- Traditionally used for dimensionality reduction or feature learning

Autoencoders

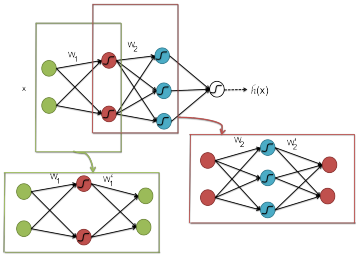

But more typically, we reduce dimensionality at each level.

Deep Belief Networks

- Autoencoder networks learn low dimensional encodings
- With more layers, can learn better encodings
- After each individual encoding layer has been learned, put them together and backpropagate to tune the entire encoder-decoder network

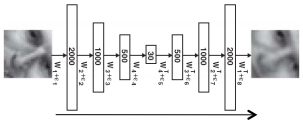

Very Large Scale DBNs

[Quoc Le et al., ICML 2012]

- Data: 10 million 200x200 unlabeled images, sampled from YouTube
- Training: 1000 machines (16000 cores) for one week
- Learned network: 3 multi-stage layers, 1.15 billion parameters
- Achieves 15.8% accuracy classifying 1 of 22K ImageNet items (sota at the time was 9.5%)

One Layer

Local Contrast Normalization

L2 pooling

Local Receptive Fields – not convolutional

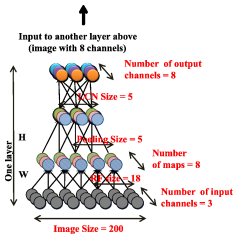

Examining Specific Learned Features

Real images that most excite the feature:

Image synthesized to most excite the feature: