

# Deep Belief Networks

Presented by Joseph Nunn

Psych 149/239 Computational Models of Cognition

University of California, Irvine  
Winter 2015



# Talk Structure

- Connectionist Background Material
- To Recognize Shapes, First Learn to Generate Images [Hinton 2006]
- Learning Hierarchical Category Structure in Deep Neural Networks [Saxe et al 2013]
- Letting Structure Emerge: Connectionist and Dynamical Approaches to Cognition [McClelland et al 2010]
- Intriguing Neural Network Properties [Szegedy et al 2013]
- Future of Connectionism




# Connectionist Background

- Neural Plausibility
- Pandemonium - [Selfridge 1958]
- Perceptrons - [Minsky & Papert 1969]
- Backpropagation - Hinton and many others
- AI Winter(s) - 1974-80 and 1987-93
- MINST and other types of Test Data



# Neural Plausibility


- Connectionist models are only vaguely related to actual neurons and brains.
- Many simplifications or patently unreal properties exist in Connectionist models and algorithms.
- Although at Marr's **Algorithmic** level of analysis, Connectionist model details are 'inspired' by neuroscience not rooted in it.





# Pandemonium Model


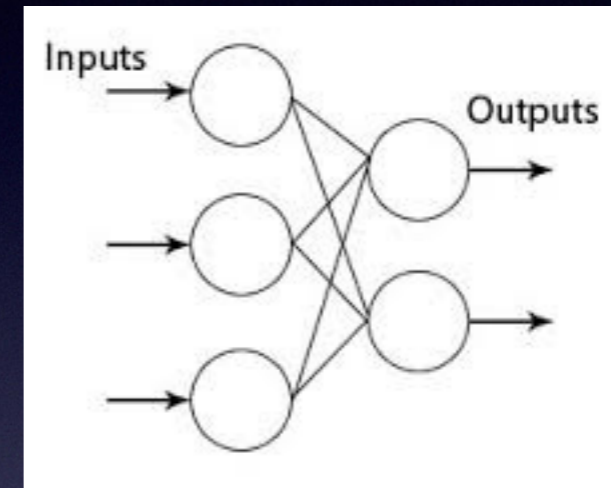
- Each layer comprises many independent agents, or demons, running concurrently.
- Demons become more or less vocal depending on input they see in previous layer.
- Most active top level demons get represented in active conscious mind.
- An early model of Parallel Distributed Processing (PDP) [Selfridge 1958]





# Perceptrons


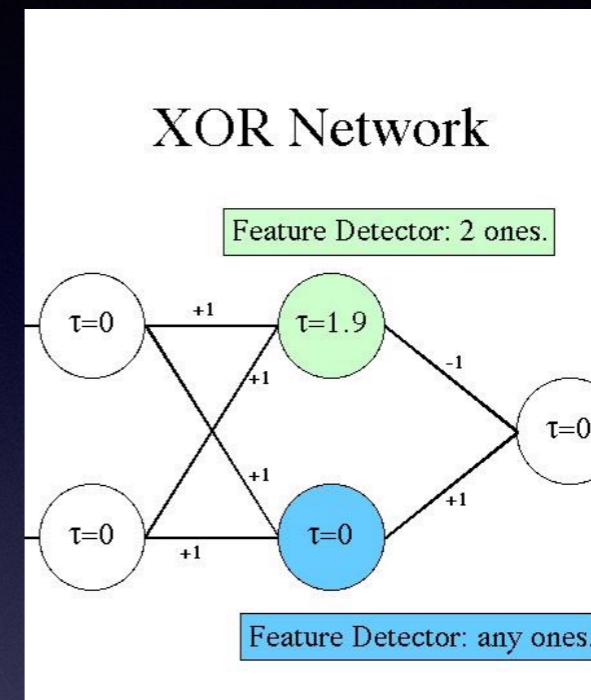
- Early type of Neural Network consisting of an input layer and output layer.
- Easily trainable.
- Shown to be incapable of learning functions not linearly separable in *Perceptrons* book [Minsky & Papert 1969].
- *Perceptrons* book contributed to the 'death' of connectionist research vs symbolist and the first AI Winter 1974-1980.





# Backpropagation

- Neural Networks with 1 or more hidden layers are capable of learning linearly separable functions, but no algorithm was known that could train them.
- Back propagation is an algorithm that can train multilayer networks 'rediscovered' and popularized in the mid 80's by several people including Hinton.
- Algorithm works by computing the error between the expected output and the actual output and distributing that error over the previous connections, correcting the connection weights by a small amount.
- Works by gradient descent over a number of training epochs on labeled data.





# AI Winter(s)

- Twice in the history of Artificial Intelligence has research progress and funding dried up, these are referred to as the 'AI Winters', 1974-1980 and 1987-1993.
- Precipitated by overpromises of early researchers and infighting between Connectionist and Symbolicist approaches to AI, each of which at times has been ascendant.
- Much promising research was delayed or had funding cut.
- Each time algorithmic discoveries from either approach has brought AI back in vogue.
- Lesson: Both Connectionist and Structured Probabilistic modeling approaches should be encouraged in Cognitive Science in order to avoid a similar fate. Both approaches have much to contribute.
- Now entering another boom in AI research instigated by the successes with Deep Learning.



# Test Data

- Several standard data sets are used in AI in order to compare the performance of various algorithms.
- Contests are also held, both academic and commercial (Kaggle).
- **MNIST** - **M**ixed **N**ational Institute of **S**tandards and **T**echnology - handwriting database used in the papers reviewed.
- Best performance today with Deep Learning is within a few percent of what humans can do.





# Talk Structure

- Connectionist Background Material
- **To Recognize Shapes, First Learn to Generate Images [Hinton 2006]**
- Learning Hierarchical Category Structure in Deep Neural Networks [Saxe et al 2013]
- Letting Structure Emerge: Connectionist and Dynamical Approaches to Cognition [McClelland et al 2010]
- Intriguing Neural Network Properties [Szegedy et al 2013]
- Future of Connectionism




# 5 Strategies for Learning Multilayer Networks

- Support Vector Machines - Perceptrons
- Evolutionary exploration of weight space
- Multilayer Feature Detectors
- Backpropagation
- Generative Feedback - 'Wake-Sleep'



# Evolutionary exploration of weight space

- Starting from initial configuration, perturb a random weight and evaluate.
- In a fully connected network, any single weight changed could affect the output for any input in the test data.
- Computationally impractical, I know of no model of any size that uses should an algorithm.





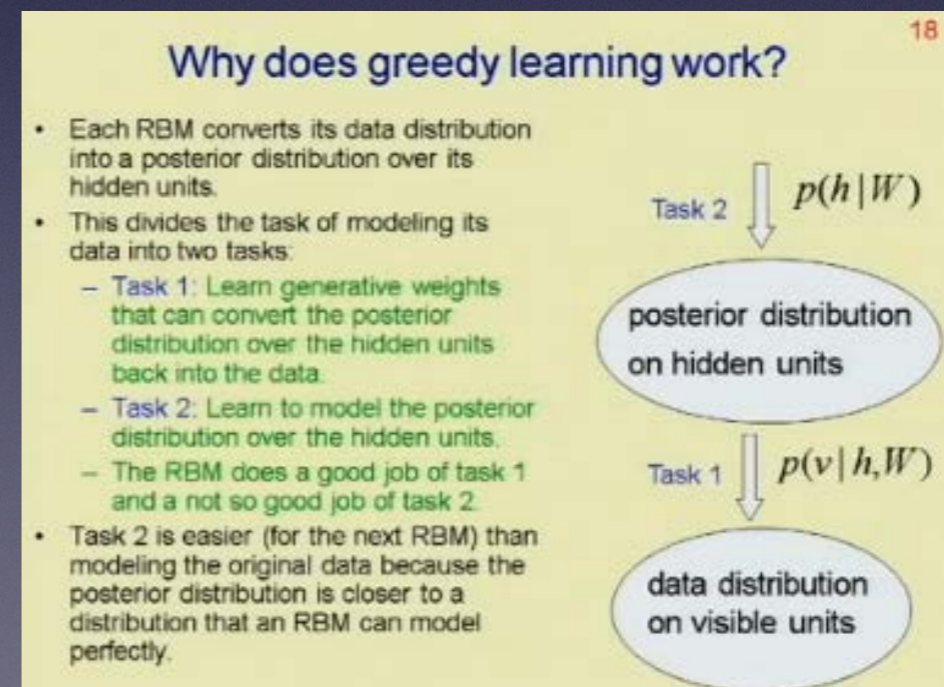
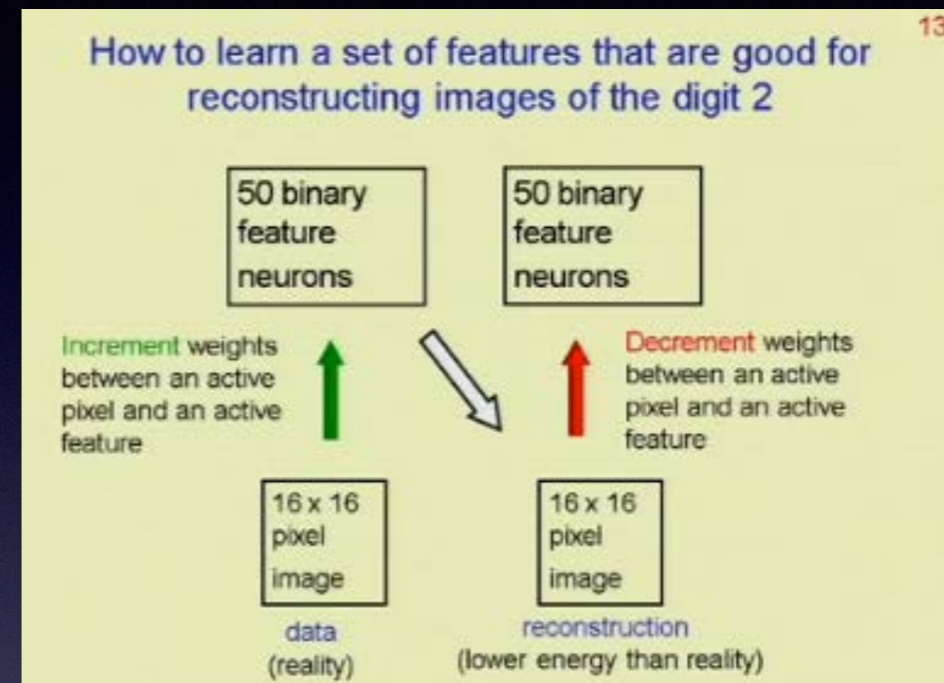
# Multilayer Feature Detectors

- Attempts to learn ‘interesting correlations’ between input elements as features detectors in hidden layers.
- Can be composed hierarchically of many layers, each learning ‘interesting correlations’ between the elements in the previous layer.
- Without guidance by desired output any ‘interesting correlation’ in input could be learned as a feature. At the top level feature detectors learned are hoped to be useful for categorizing the input.
- Vaguely defined, what counts as an ‘interesting correlation’ and why?
- Computationally intractable, equivalent to searching through a vector space for a random basis explaining the input using heuristic methods. May not converge.



# Wake-Sleep Algorithm

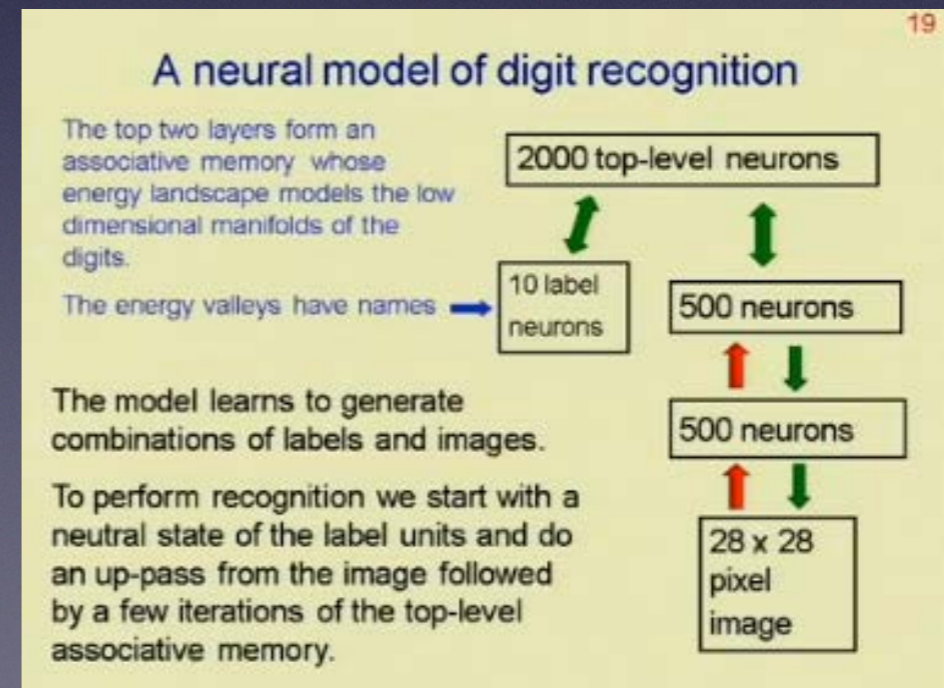
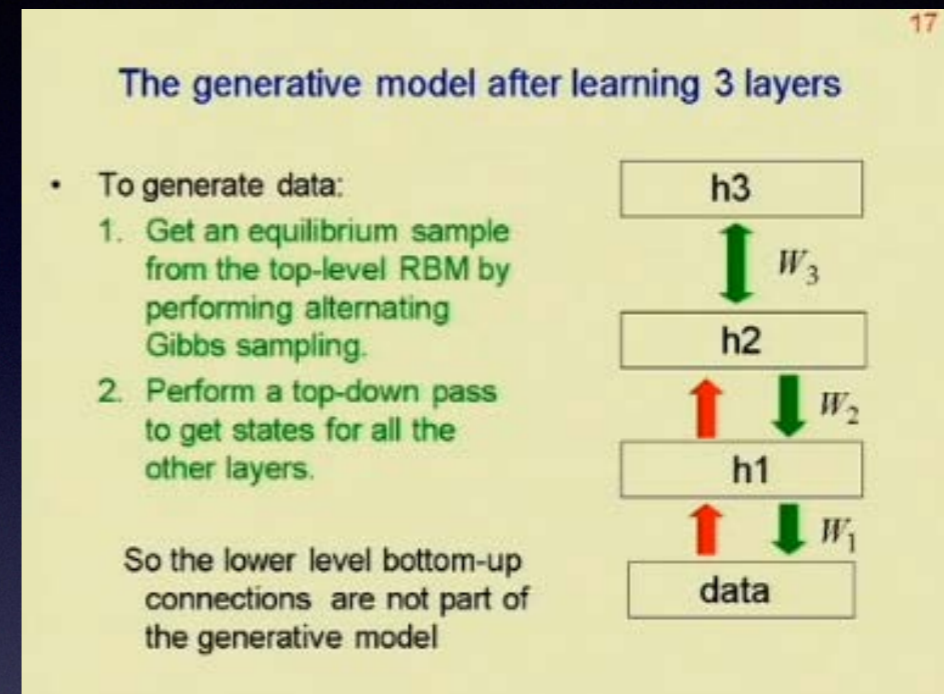
- Hinton's very successful Deep Learning Network.
- Can consist of multiple layers. Latest research shows the more the merrier, some networks 9-10 hidden layers.
- Each layer consists of a Randomized Boltzmann Machine (RBM), top layer with symmetric connections.
- Trains very fast and performs better than Backprop.





# Wake-Sleep Cont.

- Top layer forms an associative memory that settles into stable state.
- Paper discusses augmenting Wake-Sleep with Backprop for fine tuning. AKA 'Bag of Tricks'.
- Hinton's Google Presentation <https://www.youtube.com/watch?v=AyzOUbkUf3M>





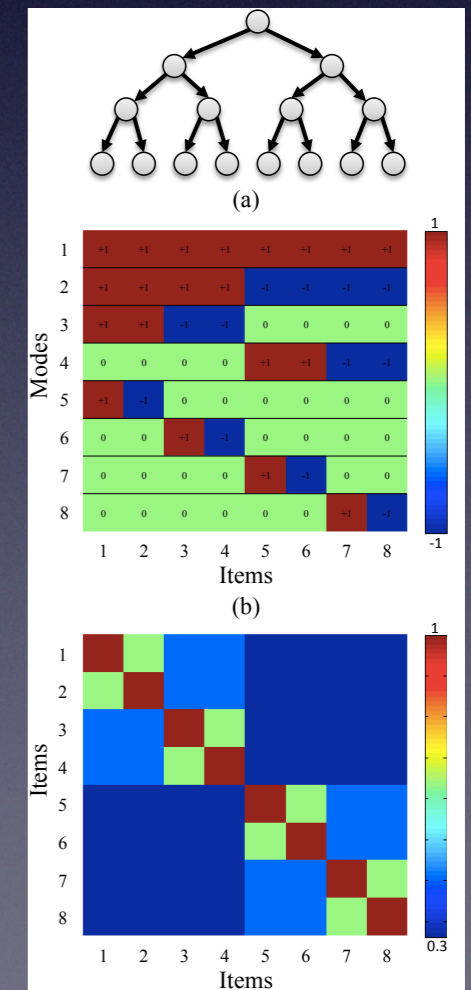
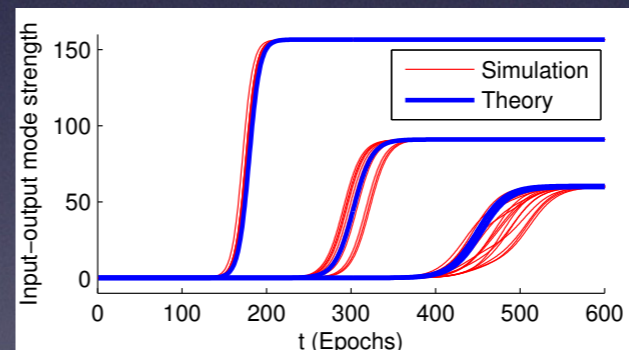
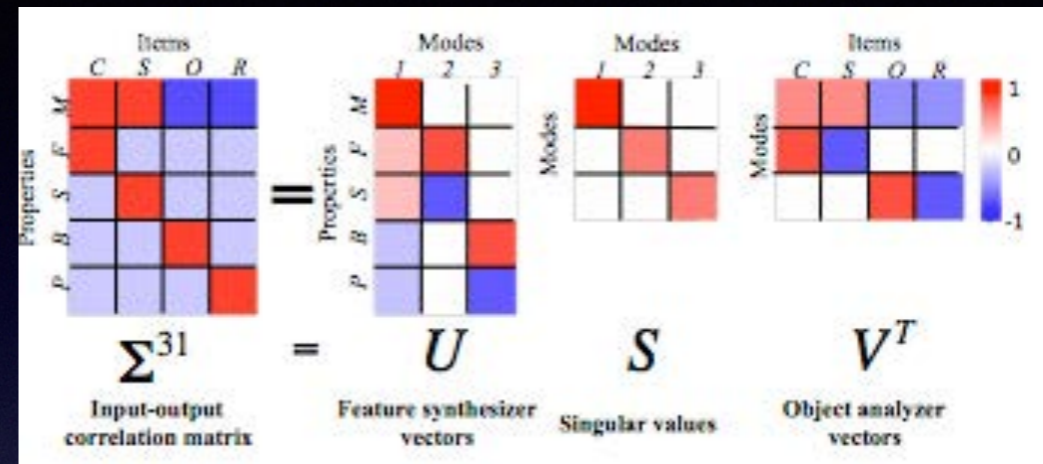
# Talk Structure

- Connectionist Background Material
- To Recognize Shapes, First Learn to Generate Images [Hinton 2006]
- **Learning Hierarchical Category Structure in Deep Neural Networks [Saxe et al 2013]**
- Letting Structure Emerge: Connectionist and Dynamical Approaches to Cognition [McClelland et al 2010]
- Intriguing Neural Network Properties [Szegedy et al 2013]
- Future of Connectionism



# Learning Hierarchical Category Structure

- Uses **S**ingular **V**alue **D**ecomposition (SVD) to investigate efficiency and learning dynamics of backpropagation.
- Singular values show importance relation between matrix dimensions.
- Exhibits non linear learning dynamics including rapid stage like transitions.
- Used a probabilistic generative system to develop arbitrary hierarchical structured data.
- Singular values and their magnitudes reflect hierarchal organized data and the degrees of separation.
- Learning dynamics are strongly correlated with magnitudes of singular values. Stronger input/output correlations described by singular values take less time to learn.





# Talk Structure

- Connectionist Background Material
- To Recognize Shapes, First Learn to Generate Images [Hinton 2006]
- Learning Hierarchical Category Structure in Deep Neural Networks [Saxe et al 2013]
- **Letting Structure Emerge: Connectionist and Dynamical Approaches to Cognition [McClelland et al 2010]**
- Intriguing Neural Network Properties [Szegedy et al 2013]
- Future of Connectionism



# Letting Structure Emerge

- McClelland et al argues that Connectionism is a better way forward for cognitive science than structured probabilistic approaches.
- Structured probabilistic approaches require too much specified knowledge such as the form of the hypothesis space, space of concepts and related structures, priors etc, that may not be present in the real world, e.g. taxonomy hierarchies and prey/predator similarities.
- Stresses the relevance of the Algorithmic level in modeling cognition. Places importance on 'integrated accounts' across multiple levels of analysis for cognitive modeling.
- Takes the view that cognitive behavior is 'Emergent' from simpler, lower level processes. AKA patterns of neuronal activations.
- Takes issue with hypothesis testing as primary cognitive task as people appear to vary their algorithm depending on constraints while underlying probabilistic problem remains the same.
- Cannot separate cognition as an emergent phenomena from the underlying mechanism without missing critical aspects.



# Talk Structure

- Connectionist Background Material
- To Recognize Shapes, First Learn to Generate Images [Hinton 2006]
- Learning Hierarchical Category Structure in Deep Neural Networks [Saxe et al 2013]
- Letting Structure Emerge: Connectionist and Dynamical Approaches to Cognition [McClelland et al 2010]
- **Intriguing Neural Network Properties [Szegedy et al 2013]**
- Future of Connectionism



# Semantic Meaning of Individual Units

- Neural Networks classically viewed as hierarchical feature detectors. A classification decision is made based on the features identified in the input.
- Individual neurons or 'units' each learn to recognize one distinguishing feature in the input.
- Features detectors learned form a basis vector set for interpreting data. Researchers like to think these feature units have semantic interpretations.
- Turns out any random transformation of basis still has semantic interpretation.
- Implication - Neural Networks are not semantic feature detectors, but instead encode data regularities across whole activation surface.

## Natural Basis




(a) Unit sensitive to lower round stroke.



(c) Unit sensitive to left, upper round stroke.

## Random Basis



(a) Direction sensitive to upper straight stroke, or lower round stroke.



(c) Direction sensitive to round top stroke.



# Adversarial Examples

- Minor perturbations to input cause wild misclassifications.



- Effect appears independent of any particular Connectionist model or set of hyperparameters.

- Seems to be 'In the data'.

- Implications

- Connectionist models and algorithms are not cognitively plausible.

- If issues are 'In the data' perhaps analysis at Marr's **Computational** level is best.







# Talk Structure

- Connectionist Background Material
- To Recognize Shapes, First Learn to Generate Images [Hinton 2006]
- Learning Hierarchical Category Structure in Deep Neural Networks [Saxe et al 2013]
- Letting Structure Emerge: Connectionist and Dynamical Approaches to Cognition [McClelland et al 2010]
- Intriguing Neural Network Properties [Szegedy et al 2013]
- **Future of Connectionism**



# Future of Connectionism

- Deep Learning networks currently dominate AI competitions and performance on tasks such as MNIST data.
- Success is being driven by hardware advances in programmable GPUs due to parallel nature of Connectionist algorithms. Progress now 'plugged in' to Moore's Law.
- Software tools growing in sophistication.
- Google Voice Recognition on Android phones.
- Convolutional Networks.
- **Prediction:** Performance increases will continue to dominate other forms of AI not amplifiable by parallelization on GPUs.





Thoughts?