Deep Learning

By Will Woliver-Jones

What is Deep Learning?

- Deep learning is a hierarchical system based on an assembly of artificial neural nets
  - Recently Deep learning has been applied to any implementation of an ANN
- Nets are used to represent more and more specific parts of the problem
- There are several types of structures
  - Deep Belief network
  - Deep Neural Net
  - Convolutional Neural networks

Artificial Neural Nets

- Based on biological systems
- Basically a graph problem
  - Nodes and edges with weights
  - But the weights are updated as the net learns “right” and “wrong”

An ANN

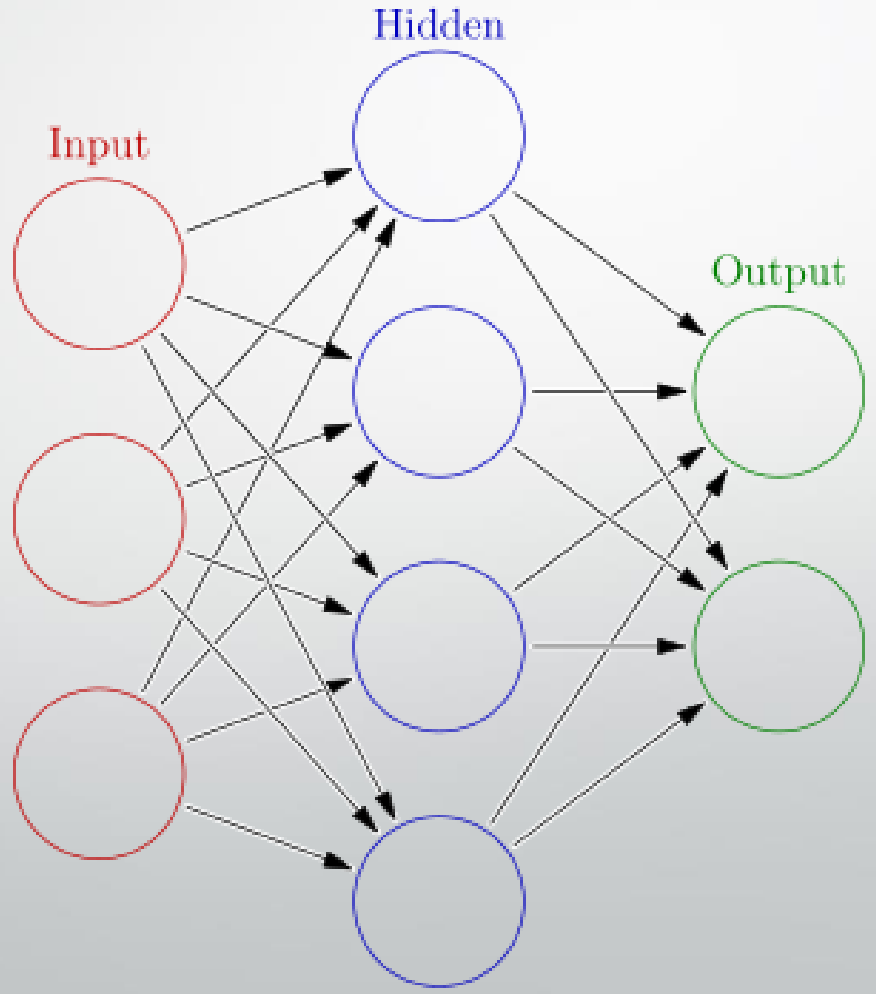

Putting them together

- More layers can be used to find less linear solutions
  - By having more steps more abstraction can be reached
- More nodes can be added to increase the specificity of the result
  - A wider range of values can be measured

Training

- The net must be pre-trained before it can be used
  - This involved several main methods:
    - Supervised
    - Unsupervised
    - Reinforcement

Training

- Unsupervised
  - The net is provided raw, unlabeled data to use
  - Deep Learning Nets learn best through back propagation
    - Wrong answers give the whole chain a decreased weight
    - This also means that the net cannot be randomly initialized
      - Random initialization prevents the effectives use of back propagation to correct for errors which causes the deep net to have a lower success rate than shallow networks
  - Greedy layer-wise initialization works well
    - Each layer learns a more specific version of the layer above it

Unsupervised greedy layer-wise training procedure.

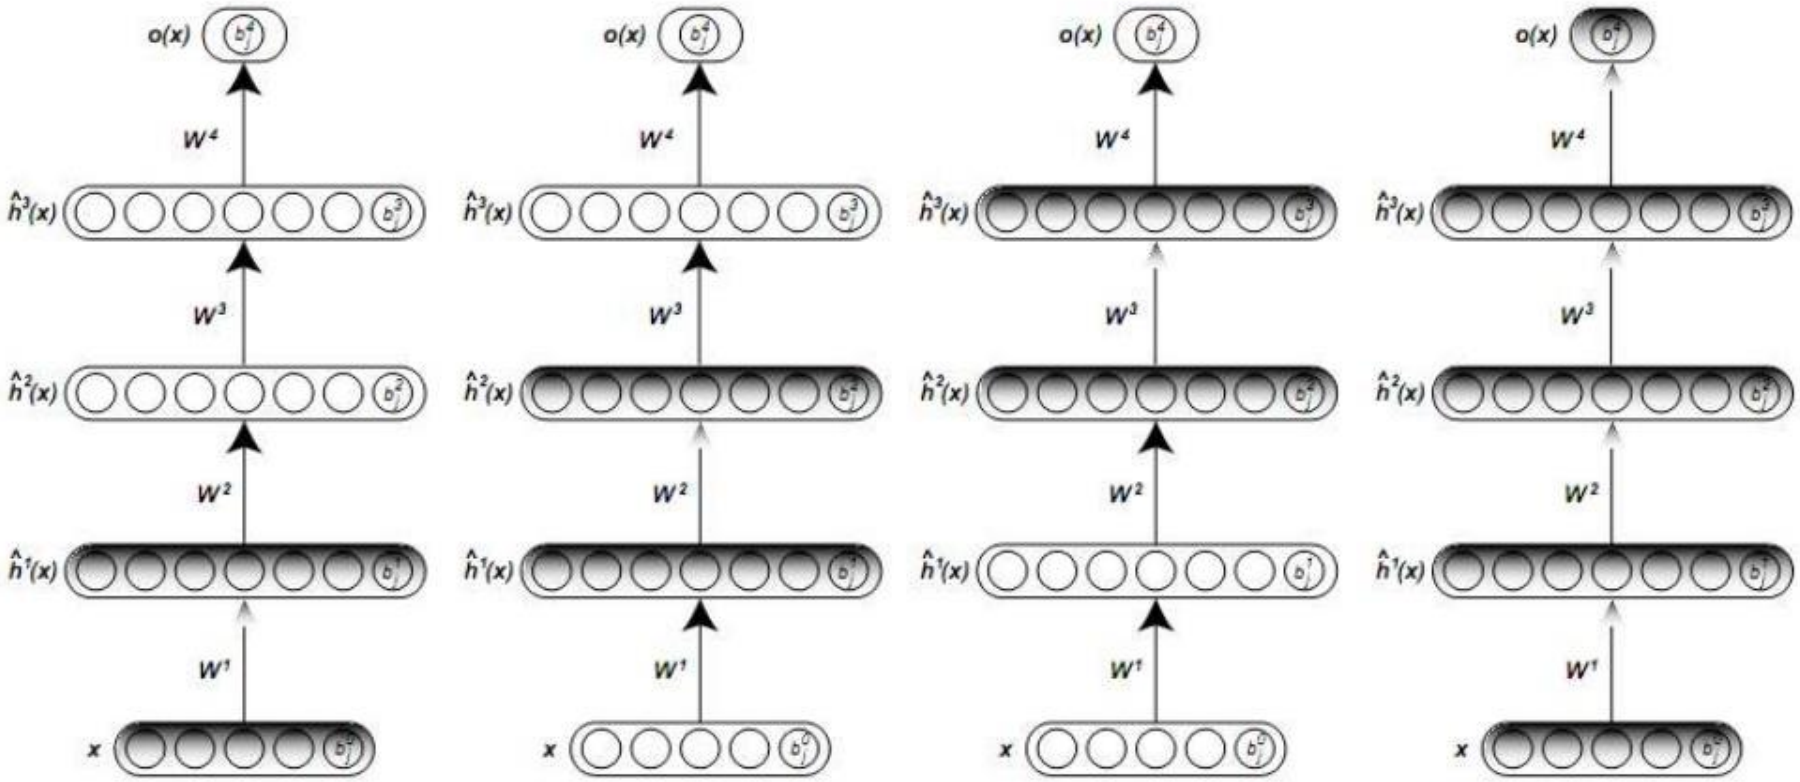

(a) First hidden layer pre-training

(b) Second hidden layer pre-training

(c) Third hidden layer pre-training

(d) Fine-tuning of whole network

Training

- Supervised
  - Program is provided data with labels to read and sort
    - This provides the basis for what the net can identify
  - Used for the output layer to fine tune results

Training

- Reinforcement
  - Used when feed back can be given at intervals, rather than only at completion
  - The computer preforms some action and observes the effect it has
    - This is used to help form a policy that minimizes some cost
  - This is generally the most applicable form of training for games

How Does this apply to games?

- Better AI- Not just for playing against
  - Computer Vision
    - Handwriting interpretation
    - Facial Recognition
  - Voice recognition
    - SIRI
  - Gameplay also provides Reinforcement Training

Utility

- A well designed deep learning program can be used for a wide variety of similar tasks
  - One program can play Atari games
    - http://www.cs.toronto.edu/~vmnih/docs/dqn.pdf
  - One program can determine what is in a photo
    - http://devblogs.nvidia.com/parallelforall/digits-deep-learning-gpu-training-system/

Things to Consider

- The algorithm shouldn't always pick the best strategy
  - Could be hiding a more optimal strategy
  - It can create repetitive game play
  - The player should get to win
    - Especially on easier difficulties

Things to Consider

- Is a deep learning net the correct choice for the problem?
  - Good at identifying things
    - Can overfit- describing noise rather than real relationships
  - Good at turn based decision making
    - But might be able to be replaced with an algorithm
  - High complexity

Additional Reading

- Overview
  - http://neuralnetworksanddeeplearning.com/
- Deep Learning
  - http://research.microsoft.com/pubs/209355/DeepLearning-NowPublishing-Vol7-SIG-039.pdf
  - http://deeplearning.net/
  - http://www.cs.toronto.edu/~vmnih/docs/dqn.pdf
- Restricted Boltzmann Machines
  - http://image.diku.dk/igel/paper/AItRBM-proof.pdf
- Artificial Neural Nets
  - http://ulcar.uml.edu/~iag/CS/Intro-to-ANN.html