Overview of RBMs (L101, lecture 6.5)

Ann Copestake

Computer Laboratory
University of Cambridge

October 2016

Outline

Overview of Restricted Boltzmann Machines

Deep Learning

Artificial versus biological NNs

Outline.

Overview of Restricted Boltzmann Machines

Deep Learning

Artificial versus biological NNs

Introduction to RBMs

- Boltzmann machine: arbitrary interconnections between units. Not effectively trainable in general.
- Restricted Boltzmann Machine (RBM): one input and one hidden layer, no intra-layer links.

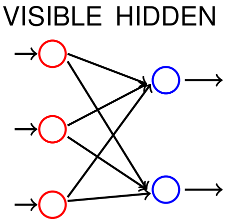

$w_1, \ldots w_6$ $\quad$ b (bias)

Introduction to RBMs

- Usually fully connected between two layers but sparse RBMs are possible.
- The layers allow for efficient implementations — weights can be described by a matrix, fast computation.
- Generative probabilistic model: bipartite graph units in hidden layer conditionally independent given input layer and vice versa.
- RBMs allow efficient Gibbs sampling for training (as a step in the overall procedure).
- Goodfellow et al 2016 (http://www.deeplearningbook.org)
  Murphy 'Machine Learning: a Probabilistic Perspective'

  https://deeplearning4j.org/restrictedboltzmannmachine

Training RBMs: reconstruction of input

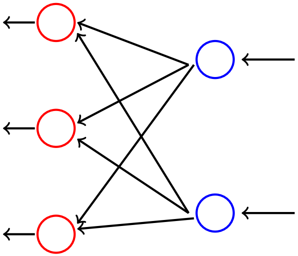

b' (bias) $\qquad w_1, \ldots w_6$

- Forward pass: $P(\text{output}|\text{input}; w)$
- Backprop: $P(\text{input}|\text{output}; w)$
- Overall, joint probability: $P(\text{input}, \text{output})$

Graphical models

- graphical models: show dependence of variables
- unlinked nodes A and B are conditionally independent (so constrain the model)
- directed graphical model (belief network, Bayesian network): A $\rightarrow$ B means B is directly dependent on A
  if A $\rightarrow$ B $\rightarrow$ C, then C is only indirectly dependent on A
- undirected graphical model (Markov Random Field (MRF), Markov network): appropriate when variables interact but causality unclear or operates bidirectionally
- an immorality: a directed graph with links from A to C and B to C, but no link between A and B (cannot be converted to an undirected graph)

Some (hopefully) intuitive explanations of terminology

- regularization: methods of choosing the priors to avoid overfitting (less necessary if lots of data). e.g., fitting a smooth curve rather than a wiggly one.
- energy function: approximation to probabilities of states (always > 0) in undirected models. Close connection with physics (hence terminology).
- back-propagation aka backprop: information about the cost flowing backward through the network (e.g., computing the gradient).
- stochastic gradient descent: performing learning using the gradient.

Outline.

Overview of Restricted Boltzmann Machines

Deep Learning

Artificial versus biological NNs

Combining RBMs

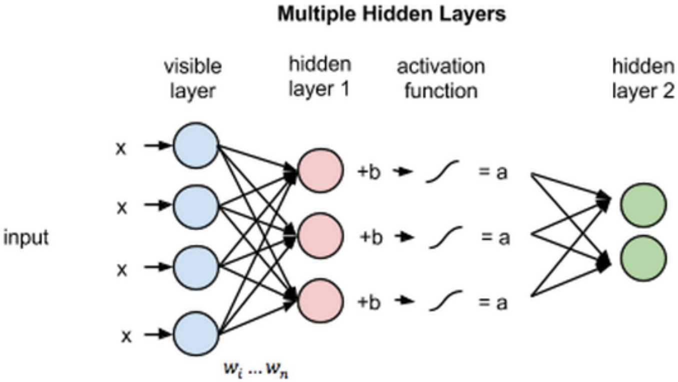

`https://deeplearning4j.org/restrictedboltzmannmachine`

Deep Learning

- One of the most successful deep learning architectures involves combining RBMs, so the output from one RBM is the input to the next.
- RBMs can be trained separately and then fine-tuned in combination.
- The layers allow for efficient implementations and successive approximations to concepts.
- Unlike LDA (and other similar models), there is no predefined interpretation for the latent variables.
- Different architecture needed for sequences and most language problems (RNN/LSTM).

Deep Learning (continued)

- Deep learning: models are made as powerful as possible to the point they are“barely possible to train or use” (`http://www.deeplearningbook.org` 16.7).
- Training for non-standard tasks and tuning hyperparameters is a matter of experimentation . . .
- Myths in the literature, but little publication of negative results.
- The lack of predefined interpretation of the latent variables is what makes the models more flexible/powerful.
- But the models are usually not interpretable by humans after training: potential for serious practical and ethical issues.

Outline.

Overview of Restricted Boltzmann Machines

Deep Learning

Artificial versus biological NNs

Artifical vs biological NNs

- ANNs and BNNs both take input from many neurons and carry out simple processing (e.g., summation), then output to many neurons.
- ANNs are still tiny: biggest c160 billion parameters. Human brain has tens of billions of neurons, each with up to 100,000 synapses.
- Brain connections are much slower than ANNs: chemical transmission across synapse. Incresed size and parallelism (more than) makes up for this.
- Neurotransmitters are complex and not well understood: biological neurons are only crudely approximated by on/off firing.

Artifical vs biological NNs (continued)

- Brains grow new synapses and lose old ones: individual brains evolve (Hebbian Learning: “Neurons which fire together wire together”).
- Brains are embodied: processing sensory information, controlling muscles. There is no hard division between these parts of the brain and concepts/reasoning (e.g., experiments with *kick* vs *hit*).
- Brains have evolved over (about) 600 million years (more if we include nerve nets, as in jellyfish).
- Brains are expensive (about 20% of a person’s energy), but much more efficient than ANNs.
- and . . .

Next lectures

- Nov 10: Overview of RNNs and LSTMs
- Final lecture (Nov 24): using LSTMs, reviewing the state of the art.