


Lecture 7: Artificial Neural Networks

- Sigmoid units
- Networks of sigmoid units
- Backpropagation
- Bells and whistles

1

Recall: Perceptron



Sometimes we will add a fixed component $x_0 = 1$ to all the instances and use simpler vector notation:

$$o(\vec{x}) = \begin{cases} 1 & \text{if } \vec{w} \cdot \vec{x} > 0 \\ -1 & \text{otherwise.} \end{cases}$$

2

Building networks of individual units

- Perceptrons have very simple decision surfaces

If we connect them into networks, the error surface for the network is not differentiable (because of the hard threshold)


So we cannot apply gradient descent to find a good set of weights...

- *We would like a “soft” threshold!*

Nicer math, and closer to biological neurons...

3

Sigmoid unit



$\sigma(x)$ is the sigmoid function: $\frac{1}{1 + e^{-x}}$

Nice property: $\frac{d\sigma(x)}{dx} = \sigma(x)(1 - \sigma(x))$

We can derive gradient decent rules to train

- One sigmoid unit
- *Multi-layer networks* of sigmoid units → Backpropagation

4

Error gradient for a sigmoid unit

$$\begin{aligned}\frac{\partial E}{\partial w_i} &= \frac{\partial}{\partial w_i} \frac{1}{2} \sum_{d \in D} (t_d - o_d)^2 = \frac{1}{2} \sum_d \frac{\partial}{\partial w_i} (t_d - o_d)^2 \\ &= \frac{1}{2} \sum_d 2(t_d - o_d) \frac{\partial}{\partial w_i} (t_d - o_d) \\ &= \sum_d (t_d - o_d) \left(-\frac{\partial o_d}{\partial w_i} \right) \\ &= - \sum_d (t_d - o_d) \frac{\partial o_d}{\partial net_d} \frac{\partial net_d}{\partial w_i}, \text{ where} \\ net_d &= \sum_{i=0}^n w_i x_i\end{aligned}$$

5

Error gradient for a sigmoid unit (2)

But we know:

$$\frac{\partial o_d}{\partial net_d} = \frac{\partial \sigma(net_d)}{\partial net_d} = o_d(1 - o_d)$$

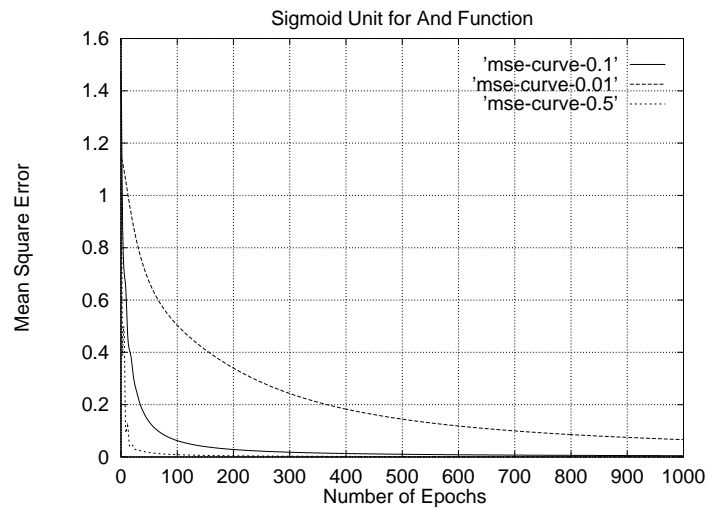
$$\frac{\partial net_d}{\partial w_i} = \frac{\partial (\vec{w} \cdot \vec{x}_d)}{\partial w_i} = x_{i,d}$$

So:

$$\frac{\partial E}{\partial w_i} = - \sum_{d \in D} (t_d - o_d) o_d (1 - o_d) x_{i,d}$$


6

Learning curves for AND function



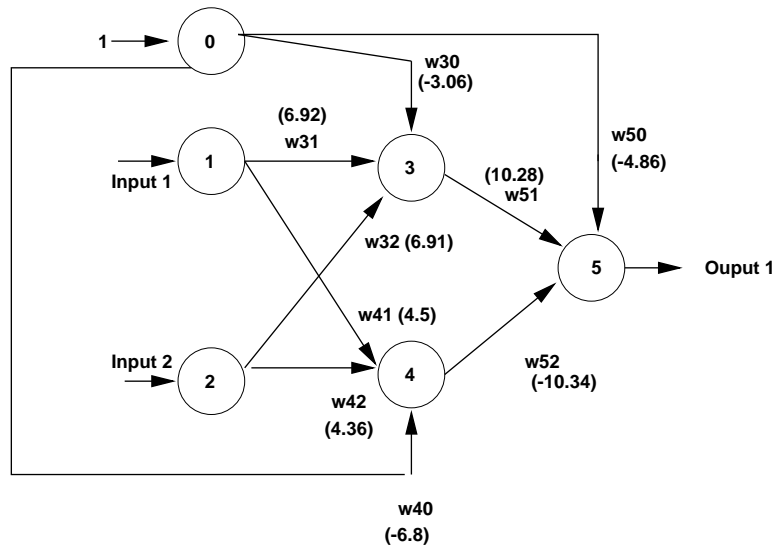
7

A single sigmoid unit cannot learn XOR!



8


A network of sigmoid units can represent XOR...



Input1	Input2	o3	o4	Ouput 1
0	0	0.04	0.001	0.011
0	1	0.98	0.08	0.99
1	0			
1	1			

9

...And it can learn XOR too!



10

Backpropagation algorithm

1. Initialize all weights to small random numbers.
2. Repeat until satisfied:
 - (a) Pick a training example
 - (b) Input the training example to the network and compute the network outputs
 - (c) For each output unit k

$$\delta_k \leftarrow o_k(1 - o_k)(t_k - o_k)$$

- (d) For each hidden unit h

$$\delta_h \leftarrow o_h(1 - o_h) \sum_{k \in \text{outputs}} w_{hk} \delta_k$$

11

0.001in=0.401920.001in0.1in=0.401920.1in

- (e) Update each network weight w_{ij}

$$w_{ij} \leftarrow w_{ij} + \eta \delta_j x_{ij}$$

x_{ij} is the input from unit i into unit j (so for the output neurons, the x 's are the signals received from the hidden layer neurons)

This algorithm is the *incremental* version.

Alternatively, we can do a *batch version*: cycle through the training data, accumulate the weight changes for all instances, then change the weights.

Terminology: *epoch* = one pass through all the training instances

12

Why this algorithm?

For the output units, this is just like the update for a single neuron. The only difference is that now the error function for the whole network is defined over all the outputs:

$$E(\vec{w}) = \frac{1}{2} \sum_{d \in D} \sum_{k \in \text{outputs}} (t_{kd} - o_{kd})^2$$

where t_{kd} and o_{kd} are the target and output values associated with the k th output unit and d th training example.

For the hidden units, we have to compute how much they influence the overall error.

But they only influence the error of the units immediately downstream from them!

The rest is a matter of applying the chain rule.

13

Convergence of backpropagation

Gradient descent to some local minimum

- Perhaps not global minimum...
- Can be much worse than global minimum
- There can be MANY local minima (Auer et al, 1997)

Partial solution: train multiple nets with different initial weights

Restarting is a standard trick in hill-climbing algorithms

More tricks:

- Initialize weights near zero
- Therefore, initial networks near-linear
- Increasingly non-linear functions possible as training progresses
- *Make sure the units start with different weights, to break symmetry!*

14

Expressiveness of feed-forward neural networks

- Every Boolean function can be represented by a network with single hidden layer, but might require exponential (in number of inputs) hidden units
- Every bounded continuous function can be approximated with arbitrarily small error, by a network with one, sufficiently large hidden layer [Cybenko 1989; Hornik et al. 1989]
- Any function can be approximated to arbitrary accuracy by a network with two hidden layers [Cybenko 1988].

Inductive bias is roughly *smooth interpolation between points*

15

More on backpropagation

- Gradient descent over entire *network* weight vector
- Easily generalized to arbitrary directed graphs (not only two layers)
- In theory it will find a local, not necessarily global error minimum, but in practice, it often works well (can run multiple times)
- Minimizes error over *training* examples
Will it generalize well to subsequent examples?
See the overfitting issue...
- Training can take thousands of iterations → VERY SLOW!
But using network after training is very fast.

16


Example: Learning an encoder function

Input		Output
10000000	→	10000000
01000000	→	01000000
00100000	→	00100000
00010000	→	00010000
00001000	→	00001000
00000100	→	00000100
00000010	→	00000010
00000001	→	00000001

Can this be learned??

17

Learning hidden layer representations

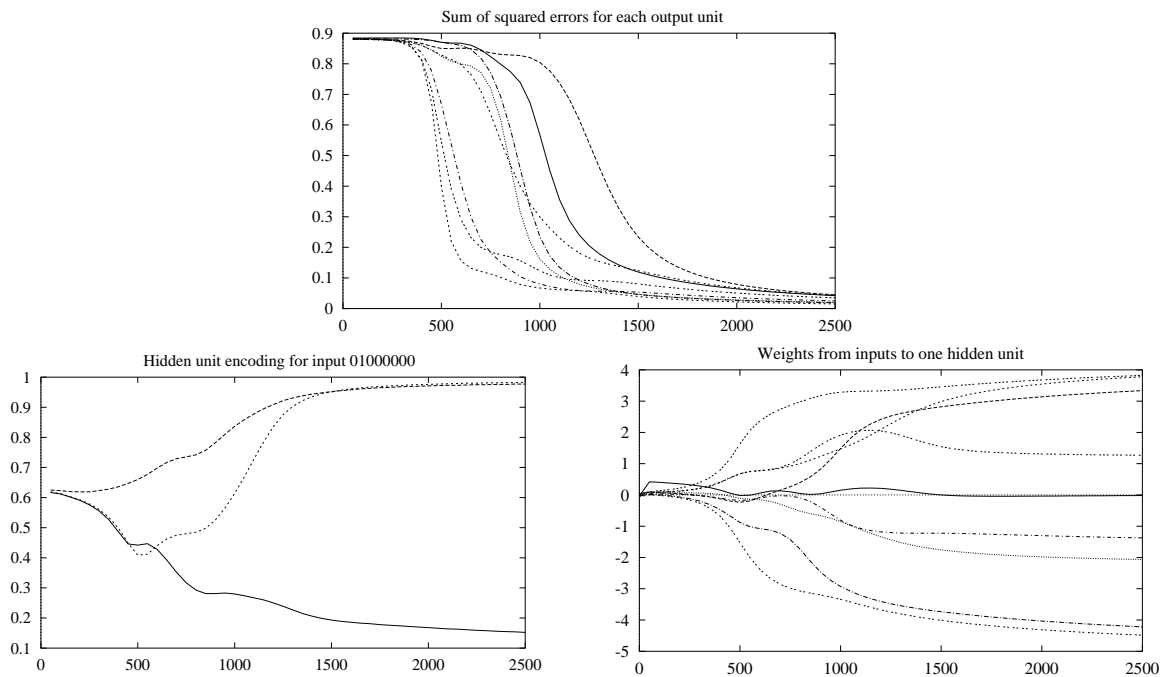


Learned hidden layer representation:

Input		Hidden Layer				Output
10000000	→	.89	.04	.08	→	10000000
01000000	→	.15	.99	.99	→	01000000
00100000	→	.01	.97	.27	→	00100000
00010000	→	.99	.97	.71	→	00010000
00001000	→	.03	.05	.02	→	00001000
00000100	→	.01	.11	.88	→	00000100
00000010	→	.80	.01	.98	→	00000010
00000001	→	.60	.94	.01	→	00000001

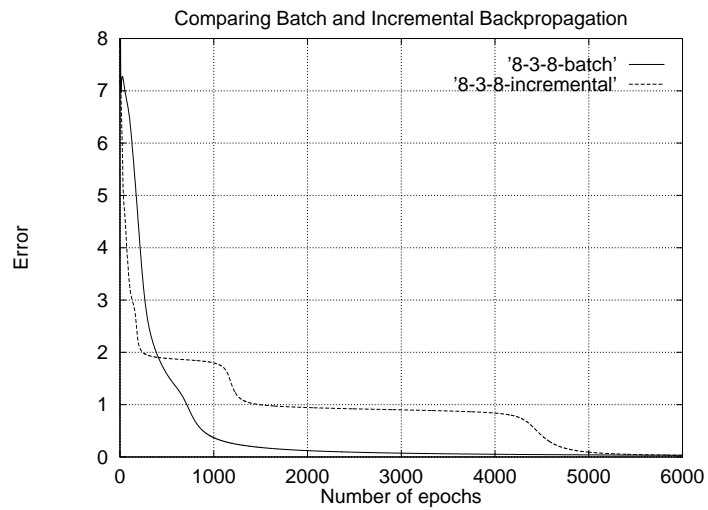
18

Evolution during training



19

Batch vs. incremental learning



20

Adding momentum

On the n -th training sample, instead of the update:

$$\Delta w_{ij} \leftarrow \eta \delta_j x_{ij}$$

we do:

$$\Delta w_{ij}(n) \leftarrow \eta \delta_j x_{ij} + \alpha \Delta w_{ij}(n-1)$$

The second term is called *momentum*

Advantages:

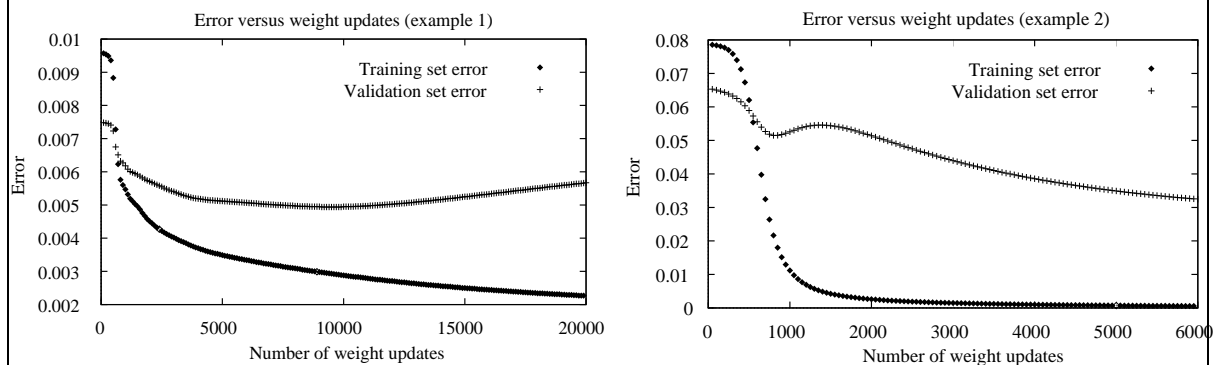
- Easy to pass small local minima
- Keeps the weights in areas where the error is flat
- Increases the speed where the gradient stays unchanged

Disadvantages:

- With too much momentum, it can get out of a global maximum!
- One more parameter to tune, and more chances of divergence

21

Overfitting in feed-forward networks



Use a validation set to decide when to stop training!

22

Practical issues

- The choice of initial weights has great impact on convergence!
If the input size is N , and N is large, a good heuristic is to choose initial weights between $-1/N$ and $1/N$.
- Backpropagation is very sensitive to the size of the learning rate! If it is too large, the weights diverge.
- Sometimes it is appropriate to use different learning rates for different layers.
- The choice of input encoding and network topology can drastically affect learning!
 - It is bad to have inputs of very different magnitude
 - A thermometer encoding can be better than a 1-of-n
 - Too many hidden units hurt (why?!). Good heuristic: $\log(N)$.

23

Alternative error functions

Penalize large weights:

$$E(\vec{w}) \equiv \frac{1}{2} \sum_{d \in D} \sum_{k \in \text{outputs}} (t_{kd} - o_{kd})^2 + \gamma \sum_{i,j} w_{ji}^2$$

Used to avoid overfitting.

Train on target slopes as well as values:

$$E(\vec{w}) \equiv \frac{1}{2} \sum_{d \in D} \sum_{k \in \text{outputs}} \left[(t_{kd} - o_{kd})^2 + \mu \sum_{j \in \text{inputs}} \left(\frac{\partial t_{kd}}{\partial x_d^j} - \frac{\partial o_{kd}}{\partial x_d^j} \right)^2 \right]$$

Tie together weights: Train each weight individually, but then replace the values with the mean of the weights obtained by backprop.

24

Consider predicting survival probability from patient data $\langle x_i, d_i \rangle$, where d_i is 1 or 0

We want to train neural network to output a *probability* given x_i (not a 0 or 1)

The brute-force approach would be to estimate the probabilities from the data and then train the network using them... but we want to avoid that.

We will do an analysis of the most likely hypothesis, similar to the previous one.

Analysis

$$P(D|h) = \prod_{i=1}^m P(x_i, d_i|h) = \prod_{i=1}^m P(d_i|h, x_i)P(x_i)$$

Since h is our hypothesis about the ²⁴⁻¹probability of each classification:

$$\begin{aligned} P(d_i|h, x_i) &= \begin{cases} h(x_i) & \text{if } d_i = 1 \\ 1 - h(x_i) & \text{if } d_i = 0 \end{cases} \\ &= h(x_i)^{d_i} (1 - h(x_i))^{1-d_i} \end{aligned}$$

The ML hypothesis is:

$$h_{ML} = \arg \max_{h \in H} \prod_{i=1}^m h(x_i)^{d_i} (1 - h(x_i))^{1-d_i} P(x_i)$$

The last factor is a constant independent of h so it can be dropped.

And by taking logs, like before, we have:

$$h_{ML} = \arg \max_{h \in H} \sum_{i=1}^m d_i \ln h(x_i) + (1 - d_i) \ln(1 - h(x_i))$$

We want to maximize the likelihood of a hypothesis $G(h,D)$:

$$\begin{aligned}\frac{\partial G(h, D)}{\partial w_{jk}} &= \sum_{i=1}^m \frac{\partial G(h, D)}{\partial h(x_i)} \frac{\partial h(x_i)}{\partial w_{jk}} \\ &= \sum_{i=1}^m \frac{\partial (d_i \ln h(x_i) + (1 - d_i) \ln(1 - h(x_i)))}{\partial h(x_i)} \frac{\partial h(x_i)}{\partial w_{jk}} \\ &= \frac{d_i - h(x_i)}{h(x_i)(1 - h(x_i))} \frac{\partial h(x_i)}{\partial w_{jk}}\end{aligned}$$

Weight update rule for a sigmoid unit:

$$w_{jk} \leftarrow w_{jk} + \Delta w_{jk}$$

where

$$\Delta w_{jk} = \eta \sum_{i=1}^m (d_i - h(x_i)) x_{ijk}$$

24-3

Constructive methods for neural networks

Meiosis networks (Hanson):

- Start with just one hidden unit, train using backprop
- Compute the variance of each weight during training
- If a unit has one or more weights of high variance, it is split into two units, and the weights are perturbed

Cascade correlation (Fahlman & Lebiere):

- Start with outputs only and train using backprop
- Add a neuron connected to all inputs, and train it to correlate to the residual error
- Connect the neuron to the output node, then freeze its weights and train the output again
- Continue until the residual error is below a threshold

When to consider using neural networks

- Input is high-dimensional discrete or real-valued (e.g. raw sensor input)
- Output is discrete or real valued, or a *vector of values*
- Possibly noisy data
- Training time is unimportant
- Form of target function is unknown
- Human readability of result is unimportant
- The computation of the output based on the input has to be fast

Examples:

- Speech phoneme recognition [Waibel] and synthesis [Nettalk]
- Image classification [Kanade, Baluja, Rowley]
- Financial prediction