

Neural network





## Backpropagation: intuition

Gradient descent method for learning weights by optimizing a loss function

1. calculate output of all nodes
2. calculate the weights for the output layer based on the error
3. "backpropagate" errors through hidden layers


Backpropagation: the details

Gradient descent method for learning weights by optimizing a loss function

1. calculate output of all nodes
2. calculate the updates directly for the output layer
3. "backpropagate" errors through hidden layers

$$
\text { loss }=\sum_{x} \frac{1}{2}(y-\hat{y})^{2} \quad \text { squared error }
$$




## Backpropagation: the details

Gradient descent method for learning weights by optimizing a loss function

$$
\operatorname{argmin}_{w, v} \sum_{x} \frac{1}{2}(y-\hat{y})^{2}
$$

1. calculate output of all nodes
2. calculate the updates directly for the output layer
3. "backpropagate" errors through hidden layers

Finding the minimum


You're blindfolded, but you can see out of the bottom of the blindfold to the ground right by your feet. I drop you off somewhere and tell you that you're in a convex shaped valley and escape is at the bottom/minimum. How do you get out?

Finding the minimum


## One approach: gradient descent

Partial derivatives give us the slope (i.e. direction to move) in that dimension


## One approach: gradient descent

Partial derivatives give us the slope (i.e. direction to move) in that dimension

## Approach:

$\square$ pick a starting point (w)

- repeat:

- move a small amount towards decreasing loss (using the derivative)


## One approach: gradient descent

Partial derivatives give us the slope (i.e. direction to move) in that dimension

## Approach:

$\square$ pick a starting point (w)
$\square$ repeat:

- pick a dimension
- move a small amount in that dimension towards decreasing loss (using the derivative)


Output layer weights



Backpropagation: the details

Gradient descent method for learning weights by optimizing a loss function

$$
\operatorname{argmin}_{w, \nu} \sum_{x} \frac{1}{2}(y-\hat{y})^{2}
$$

1. calculate output of all nodes
2. calculate the updates directly for the output layer
"backpropagate" errors through hidden layers

Learning rate
Output layer update:
$v_{k}=v_{k}+r_{i} h_{k}(y-f(v \cdot h)) f^{\prime}(v \cdot h)$
Hidden layer update:
$\quad w_{k j}=w_{k j}+\eta x_{j} f^{\prime}\left(w_{k} \cdot x\right) v_{k} f^{\prime}(v \cdot h)(y-f(v \cdot h))$

- Adjust how large the updates we'll make (a parameter to
the learning approach - like lambda for n-gram models)
- Often will start larger and then get smaller

Output layer update:
$v_{k}=v_{k}+\eta_{k}(y-f(v \cdot h)) f^{\prime}(v \cdot h)$
dden layer update:
$w_{k j}=w_{k j}+\eta x_{j} f^{\prime}\left(w_{k} \cdot x\right) v_{k} f^{\prime}(v \cdot h)(y-f(v \cdot h))$

- Adjust how large the updates we'll make (a parameter to the learning approach - like lambda for n -gram models)
- Often will start larger and then get smaller

$$
\begin{array}{rlrl}
\frac{d l o s s}{d v_{k}} & =\frac{d}{d v_{k}}\left(\frac{1}{2}(y-\hat{y})^{2}\right) & \frac{d l o s s}{d w_{k j}} & =\frac{d}{d w_{k j}}\left(\frac{1}{2}(y-\hat{y})^{2}\right) \\
& =\frac{d}{d v_{k}}\left(\frac{1}{2}\left(y-f(v \cdot h)^{2}\right)\right. & & =\frac{d}{d w_{k j}}\left(\frac{1}{2}\left(y-f(v \cdot h)^{2}\right)\right) \\
& =(y-f(v \cdot h)) \frac{d}{d v_{k}}(y-f(v \cdot h)) & & =(y-f(v \cdot h)) \frac{d}{d w_{k j}}(y-f(v \cdot h)) \\
& =-(y-f(v \cdot h)) \frac{d}{d v_{k}} f(v \cdot h) & & =-(y-f(v \cdot h)) \frac{d}{d w_{k j}} f(v \cdot h) \\
& =-(y-f(v \cdot h)) f^{\prime}(v \cdot h) \frac{d}{d v_{k}} v \cdot h & & =-(y-f(v \cdot h)) f^{\prime}(v \cdot h) \frac{d}{d w_{k j}} v \cdot h \\
\text { There's a bit of math to } & & =-(y-f(v \cdot h)) f^{\prime}(v \cdot h) \frac{d}{d w_{k j}} v_{k} h_{k} \\
\text { show this, but it's mostly } & & =-(y-f(v \cdot h)) f^{\prime}(v \cdot h) v_{k} \frac{d}{d w_{k j}} h_{k} \\
\text { just calculus... } & & =-(y-f(v \cdot h)) f^{\prime}(v \cdot h) v_{k} \frac{d}{d w_{k j}} f\left(w_{k} \cdot x\right) \\
& =-(y-f(v \cdot h)) f^{\prime}(v \cdot h) v_{k} f^{\prime}\left(w_{k} \cdot x\right) \frac{d}{d w_{k j}} w_{k} \cdot x
\end{array}
$$

| Backpropagation implementation |
| :--- |
| for some number of iterations: |
| randomly shuffle training data |
| for each example: |
| - Compute all outputs going forward |
| Calculate new weights and modified errors at output |
| layer |
| Recursively calculate new weights and modified errors on |
| hidden layers based on recursive relationship |
| Update model with new weights |

## Challenges of neural networks?

Picking network configuration

Can be slow to train for large networks and large amounts of data

Loss functions (including squared error) are generally not convex with respect to the parameter space

## Many variations

Momentum: include a factor in the weight update to keep moving in the direction of the previous update

Mini-batch:

- Compromise between online and batch
$\square$ Avoids noisiness of updates from online while making more educated weight updates

Simulated annealing:
$\square$ With some probability make a random weight update
$\square$ Reduce this probability over time
$\qquad$

