## Machine Learning Techniques



Lecture 12：Neural Network
Hsuan－Tien Lin（林軒田）
htlin＠csie．ntu．edu．tw

Department of Computer Science
\＆Information Engineering
National Taiwan University
（國立台灣大學資訊工程系）

## Roadmap

(1) Embedding Numerous Features: Kernel Models
(2) Combining Predictive Features: Aggregation Models

## Lecture 11: Gradient Boosted Decision Tree

 aggregating trees from functional gradient and steepest descent subject to any error measure(3) Distilling Implicit Features: Extraction Models

## Lecture 12: Neural Network

- Motivation
- Neural Network Hypothesis
- Neural Network Learning
- Optimization and Regularization


## Linear Aggregation of Perceptrons: Pictorial View



$$
G(\mathbf{x})=\operatorname{sign}(\sum_{t=1}^{T} \alpha_{t} \underbrace{\operatorname{sign}\left(\mathbf{w}_{t}^{\top} \mathbf{x}\right)}_{g_{t}(\mathbf{x})})
$$

- two layers of weights: $\mathbf{w}_{t}$ and $\alpha$
- two layers of sign functions: in $g_{t}$ and in $G$
what boundary can $G$ implement?

Logic Operations with Aggregation

$g_{2}$


$$
G(\mathbf{x})=\operatorname{sign}\left(-1+g_{1}(\mathbf{x})+g_{2}(\mathbf{x})\right)
$$



- $g_{1}(\mathbf{x})=g_{2}(\mathbf{x})=+1$ (TRUE): $G(\mathbf{x})=+1$ (TRUE)
- otherwise:
$G(\mathbf{x})=-1$ (FALSE)
- $G \equiv \operatorname{AND}\left(g_{1}, g_{2}\right)$

OR, NOT can be similarly implemented

Powerfulness and Limitation


8 perceptrons


16 perceptrons

target boundary

- 'convex set' hypotheses implemented: $d_{\mathrm{vc}} \rightarrow \infty$, remember? :-)
- powerfulness: enough perceptrons $\approx$ smooth boundary

$g_{1} \quad g_{2}$

$\operatorname{XOR}\left(g_{1}, g_{2}\right)$
- limitation: XOR not 'linear separable' under $\phi(\mathbf{x})=\left(g_{1}(\mathbf{x}), g_{2}(\mathbf{x})\right)$
how to implement $\operatorname{XOR}\left(g_{1}, g_{2}\right)$ ?

Multi-Layer Perceptrons: Basic Neural Network

- non-separable data: can use more transform
- how about one more layer of AND transform?

$$
\operatorname{XOR}\left(g_{1}, g_{2}\right)=\operatorname{OR}\left(\operatorname{AND}\left(-g_{1}, g_{2}\right), \operatorname{AND}\left(g_{1},-g_{2}\right)\right)
$$


perceptron (simple)
$\Longrightarrow$ aggregation of perceptrons (powerful) $\Longrightarrow$ multi-layer perceptrons (more powerful)

## Connection to Biological Neurons


by UC Regents Davis campus-brainmaps.org.
Licensed under CC BY 3.0 via Wikimedia Commons

by Lauris Rubenis. Licensed under CC BY 2.0 via
https://flic.kr/ p/fkVuZX

by Pedro
Ribeiro Simōes. Licensed under CC BY 2.0 via https://flic.kr/ p/adiv7b

## neural network: bio-inspired model

## Fun Time

Let $g_{0}(\mathbf{x})=+1$. Which of the following $\left(\alpha_{0}, \alpha_{1}, \alpha_{2}\right)$ allows
$G(\mathbf{x})=\operatorname{sign}\left(\sum_{t=0}^{2} \alpha_{t} g_{t}(\mathbf{x})\right)$ to implement $\operatorname{OR}\left(g_{1}, g_{2}\right)$ ?
(1) $(-3,+1,+1)$
(2) $(-1,+1,+1)$
(3) $(+1,+1,+1)$
(4) $(+3,+1,+1)$

## Fun Time

Let $g_{0}(\mathbf{x})=+1$. Which of the following $\left(\alpha_{0}, \alpha_{1}, \alpha_{2}\right)$ allows
$G(\mathbf{x})=\operatorname{sign}\left(\sum_{t=0}^{2} \alpha_{t} g_{t}(\mathbf{x})\right)$ to implement $\operatorname{OR}\left(g_{1}, g_{2}\right)$ ?
(1) $(-3,+1,+1)$
(2) $(-1,+1,+1)$
(3) $(+1,+1,+1)$
(4) $(+3,+1,+1)$

## Reference Answer: 3

You can easily verify with all four possibilities of $\left(g_{1}(\mathbf{x}), g_{2}(\mathbf{x})\right)$.

## Neural Network Hypothesis: Output



- OUTPUT: simply a linear model with $s=\mathbf{w}^{T} \phi^{(2)}\left(\phi^{(1)}(\mathbf{x})\right)$
- any linear model can be used-remember? :-)
linear regression

$$
h(\mathbf{x})=s
$$


err $=$ squared

## logistic regression

$$
h(\mathbf{x})=\theta(s)
$$


err = cross-entropy
will discuss 'regression' with squared error

Neural Network Hypothesis: Transformation

- $\lceil$ : transformation function of score (signal) $s$
- any transformation?
- : whole network linear \& thus less useful
- 」 : discrete \& thus hard to optimize for w
- popular choice of transformation: $\int=\tanh (s)$
- 'analog' approximation of - : easier to optimize
- somewhat closer to biological neuron
- not that new! :-)

will discuss with tanh as transformation function


$$
\begin{aligned}
& d^{(0)}-d^{(1)}-d^{(2)} \ldots-d^{(L)} \\
& w_{i j}^{(\ell)}:\left\{\begin{array}{ll}
1 \leq \ell \leq L & \text { layers } \\
0 \leq i \leq d^{(\ell-1)} & \text { inputs } \\
1 \leq j \leq d^{(\ell)} & \text { outputs }
\end{array}, \text { score } s_{j}^{(\ell)}=\sum_{i=0}^{d^{(\ell-1)}} w_{i j}^{(\ell)} x_{i}^{(\ell-1)},\right. \\
& \text { transformed } x_{j}^{(\ell)}= \begin{cases}\tanh \left(s_{j}^{(\ell)}\right) & \text { if } \ell<L \\
s_{j}^{(\ell)} & \text { if } \ell=L\end{cases}
\end{aligned}
$$

apply $\mathbf{x}$ as input layer $\mathbf{x}^{(0)}$, go through hidden layers to get $\mathbf{x}^{(\ell)}$, predict at output layer $x_{1}^{(L)}$

## Physical Interpretation



- each layer: transformation to be learned from data
- $\phi^{(\ell)}(\mathbf{x})=\tanh \left(\left[\begin{array}{c}\sum_{i=0}^{(\ell-1)} w_{i 1}^{(e)} x_{i}^{(\ell-1)} \\ \vdots\end{array}\right]\right)$
-whether $\mathbf{x}$ 'matches' weight vectors in pattern
NNet: pattern extraction with
layers of connection weights


## Fun Time

How many weights $\left\{w_{i j}^{(\ell)}\right\}$ are there in a 3-5-1 NNet?
(1) 9
(2) 15
(3) 20
(4) 26

## Fun Time

How many weights $\left\{w_{i j}^{(\ell)}\right\}$ are there in a 3-5-1 NNet?
(1) 9
(2) 15
(3) 20
(4) 26

## Reference Answer: (4)

There are $(3+1) \times 5$ weights in $w_{i j}^{(1)}$, and $(5+1) \times 1$ weights in $w_{j k}^{(2)}$.


- goal: learning all $\left\{w_{i j}^{(\ell)}\right\}$ to minimize $E_{i n}\left(\left\{w_{i j}^{(\ell)}\right\}\right)$
- one hidden layer: simply aggregation of perceptrons
-gradient boosting to determine hidden neuron one by one
- multiple hidden layers? not easy
- let $e_{n}=\left(y_{n}-\operatorname{NNet}\left(\mathbf{x}_{n}\right)\right)^{2}$ : can apply (stochastic) GD after computing $\frac{\partial e_{n}}{\partial w_{i j}^{(L)}}$ !
next: efficient computation of $\frac{\partial e_{n}}{\partial w_{i j}^{(\ell)}}$

Computing $\frac{\partial e_{n}}{\partial w_{i t}^{(L)}}$ (Output Layer)

$$
e_{n}=\left(y_{n}-\operatorname{NNet}\left(\mathbf{x}_{n}\right)\right)^{2}=\left(y_{n}-s_{1}^{(L)}\right)^{2}=\left(y_{n}-\sum_{i=0}^{d(L-1)} w_{i 1}^{(L)} x_{i}^{(L-1)}\right)^{2}
$$

specially (output layer)
$\left(0 \leq i \leq d^{(L-1)}\right)$
$\frac{\partial e_{n}}{\partial w_{i 1}^{(L)}}$
$=\frac{\partial e_{n}}{\partial s_{1}^{(L)}} \cdot \frac{\partial s_{1}^{(L)}}{\partial w_{i 1}^{(L)}}$
$=-2\left(y_{n}-s_{1}^{(L)}\right) \cdot\left(x_{i}^{(L-1)}\right)$
generally $(1 \leq \ell<L)$
$\left(0 \leq i \leq d^{(\ell-1)} ; 1 \leq j \leq d^{(\ell)}\right)$
$\frac{\partial \boldsymbol{e}_{n}}{\partial w_{i j}^{(l)}}$
$=\frac{\partial e_{n}}{\partial s_{j}^{(\ell)}} \cdot \frac{\partial s_{j}^{(\ell)}}{\partial w_{i j}^{(\ell)}}$
$=\delta_{j}^{(\ell)} \cdot\left(x_{i}^{(\ell-1)}\right)$

$$
\delta_{1}^{(L)}=-2\left(y_{n}-s_{1}^{(L)}\right) \text {, how about others? }
$$

Computing $\delta_{j}^{(\ell)}=\frac{\partial e_{n}}{\partial s_{j}^{(\ell)}}$

$$
s_{j}^{(\ell)} \stackrel{\tanh }{\Longrightarrow} x_{j}^{(\ell)} \stackrel{w_{j k}^{(\ell+1)}}{\Longrightarrow}\left[\begin{array}{c}
s_{1}^{(\ell+1)} \\
\vdots \\
s_{k}^{(\ell+1)} \\
\vdots
\end{array}\right] \Longrightarrow \cdots \Longrightarrow e_{n}
$$

$$
\begin{aligned}
\delta_{j}^{(\ell)}=\frac{\partial e_{n}}{\partial s_{j}^{(\ell)}} & =\sum_{k=1}^{d^{(\ell+1)}} \frac{\partial e_{n}}{\partial s_{k}^{(\ell+1)}} \frac{\partial s_{k}^{(\ell+1)}}{\partial x_{j}^{(\ell)}} \frac{\partial x_{j}^{(\ell)}}{\partial s_{j}^{(\ell)}} \\
& =\sum_{k}\left(\delta_{k}^{(\ell+1)}\right)\left(w_{j k}^{(\ell+1)}\right)\left(\tanh ^{\prime}\left(s_{j}^{(\ell)}\right)\right)
\end{aligned}
$$

$\delta_{j}^{(\ell)}$ can be computed backwards from $\delta_{k}^{(\ell+1)}$

## Backpropagation (Backprop) Algorithm

## Backprop on NNet

initialize all weights $w_{i j}^{(\ell)}$
for $t=0,1, \ldots, T$
(1) stochastic: randomly pick $n \in\{1,2, \cdots, N\}$
(2) forward: compute all $x_{i}^{(\ell)}$ with $\mathbf{x}^{(0)}=\mathbf{x}_{n}$
(3) backward: compute all $\delta_{j}^{(\ell)}$ subject to $\mathbf{x}^{(0)}=\mathbf{x}_{n}$
(4) gradient descent: $w_{i j}^{(\ell)} \leftarrow w_{i j}^{(\ell)}-\eta x_{i}^{(\ell-1)} \delta_{j}^{(\ell)}$ return $g_{\mathrm{NNET}}(\mathbf{x})=\left(\cdots \tanh \left(\sum_{j} w_{j k}^{(2)} \cdot \tanh \left(\sum_{i} w_{i j}^{(1)} x_{i}\right)\right)\right)$
sometimes (1) to (3) is (parallelly) done many times and average $\left(x_{i}^{(\ell-1)} \delta_{j}^{(\ell)}\right)$ taken for update in (4), called mini-batch
basic NNet algorithm: backprop to compute the gradient efficiently

## Fun Time

According to $\frac{\partial e_{n}}{\partial w_{i 1}^{(L)}}=-2\left(y_{n}-s_{1}^{(L)}\right) \cdot\left(x_{i}^{(L-1)}\right)$ when would $\frac{\partial e_{n}}{\partial w_{i 1}^{(L)}}=0$ ?
(1) $y_{n}=s_{1}^{(L)}$
(2) $x_{i}^{(L-1)}=0$
(3) $s_{i}^{(L-1)}=0$
(4) all of the above

## Fun Time

According to $\frac{\partial e_{n}}{\partial w_{i 1}^{(L)}}=-2\left(y_{n}-s_{1}^{(L)}\right) \cdot\left(x_{i}^{(L-1)}\right)$ when would $\frac{\partial e_{n}}{\partial w_{i 1}^{(L)}}=0$ ?
(1) $y_{n}=s_{1}^{(L)}$
(2) $x_{i}^{(L-1)}=0$
(3) $s_{i}^{(L-1)}=0$
(4) all of the above

## Reference Answer: (4)

Note that $x_{i}^{(L-1)}=\tanh \left(s_{i}^{(L-1)}\right)=0$ if and only if $s_{i}^{(L-1)}=0$.

## Neural Network Optimization

$$
E_{\text {in }}(\mathbf{w})=\frac{1}{N} \sum_{n=1}^{N} \operatorname{err}\left(\left(\cdots \tanh \left(\sum_{j} w_{j k}^{(2)} \cdot \tanh \left(\sum_{i} w_{i j}^{(1)} x_{n, i}\right)\right)\right), y_{n}\right)
$$

- generally non-convex when multiple hidden layers
- not easy to reach global minimum
- GD/SGD with backprop only gives local minimum
- different initial $w_{i j}^{(\ell)} \Longrightarrow$ different local minimum
- somewhat 'sensitive' to initial weights
- large weights $\Longrightarrow$ saturate (small gradient)
- advice: try some random \& small ones


## NNet: difficult to optimize,

 but practically works
## VC Dimension of Neural Network Model roughly, with tanh-like transfer functions: <br> $$
d_{\mathrm{vc}}=O(V D) \text { where } V=\# \text { of neurons, } D=\# \text { of weights }
$$



- pros: can approximate 'anything' if enough neurons (V large)
- cons: can overfit if too many neurons

NNet: watch out for overfitting!

## Regularization for Neural Network

basic choice:

$$
\text { old friend weight-decay (L2) regularizer } \Omega(\mathbf{w})=\sum\left(w_{i j}^{(\ell)}\right)^{2}
$$

- 'shrink' weights:
large weight $\rightarrow$ large shrink; small weight $\rightarrow$ small shrink
- want $w_{i j}^{(\ell)}=0$ (sparse) to effectively decrease $d_{v c}$
- L1 regularizer: $\sum\left|w_{i j}^{(\ell)}\right|$, but not differentiable
- weight-elimination ('scaled' L2) regularizer: large weight $\rightarrow$ median shrink; small weight $\rightarrow$ median shrink

$$
\text { weight-elimination regularizer: } \sum \frac{\left(w_{i}^{(e)}\right)^{2}}{1+\left(w_{i j}^{(e)}\right)^{2}}
$$

## Yet Another Regularization: Early Stopping

- GD/SGD (backprop) visits more weight combinations as $t$ increases

- smaller $t$ effectively decrease $d_{v c}$
- better 'stop in middle’: early stopping

$\left(d_{\text {vc }}^{*}\right.$ in middle, remember? :-))

when to stop? validation!


## Fun Time

For the weight elimination regularizer $\sum \frac{\left(w_{i}^{(i)}\right)^{2}}{1+\left(w_{i}^{(i)}\right)^{2}}$, what is $\frac{\partial \text { regularizer }}{\partial w_{i j}^{(i)}}$ ?
(1) $2 w_{i j}^{(l)} /\left(1+\left(w_{i j}^{(e)}\right)^{2}\right)^{1}$
(2) $2 w_{i j}^{(l)} /\left(1+\left(w_{i j}^{(e)}\right)^{2}\right)^{2}$
(3) $2 w_{i j}^{(\ell)} /\left(1+\left(w_{i j}^{(\ell)}\right)^{2}\right)^{3}$
(4) $2 w_{i j}^{(\ell)} /\left(1+\left(w_{i j}^{(\ell)}\right)^{2}\right)^{4}$

## Fun Time

For the weight elimination regularizer $\sum \frac{\left(w_{i}^{(i)}\right)^{2}}{1+\left(w_{i}^{(i)}\right)^{2}}$, what is $\frac{\partial \text { regularizer }}{\partial w_{i j}^{(i)}}$ ?
(1) $2 w_{i j}^{(\ell)} /\left(1+\left(w_{i j}^{(\ell)}\right)^{2}\right)^{1}$
(2) $2 w_{i j}^{(\ell)} /\left(1+\left(w_{i j}^{(\ell)}\right)^{2}\right)^{2}$
(3) $2 w_{i j}^{(\ell)} /\left(1+\left(w_{i j}^{(\ell)}\right)^{2}\right)^{3}$
(4) $2 w_{i j}^{(\ell)} /\left(1+\left(w_{i j}^{(\ell)}\right)^{2}\right)^{4}$

## Reference Answer: 2

Too much calculus in this class, huh? :-)

## Summary

(1) Embedding Numerous Features: Kernel Models
(2) Combining Predictive Features: Aggregation Models
(3) Distilling Implicit Features: Extraction Models

## Lecture 12: Neural Network

- Motivation
multi-layer for power with biological inspirations
- Neural Network Hypothesis
layered pattern extraction until linear hypothesis
- Neural Network Learning
backprop to compute gradient efficiently
- Optimization and Regularization
tricks on initialization, regularizer, early stopping
- next: making neural network 'deeper’

