

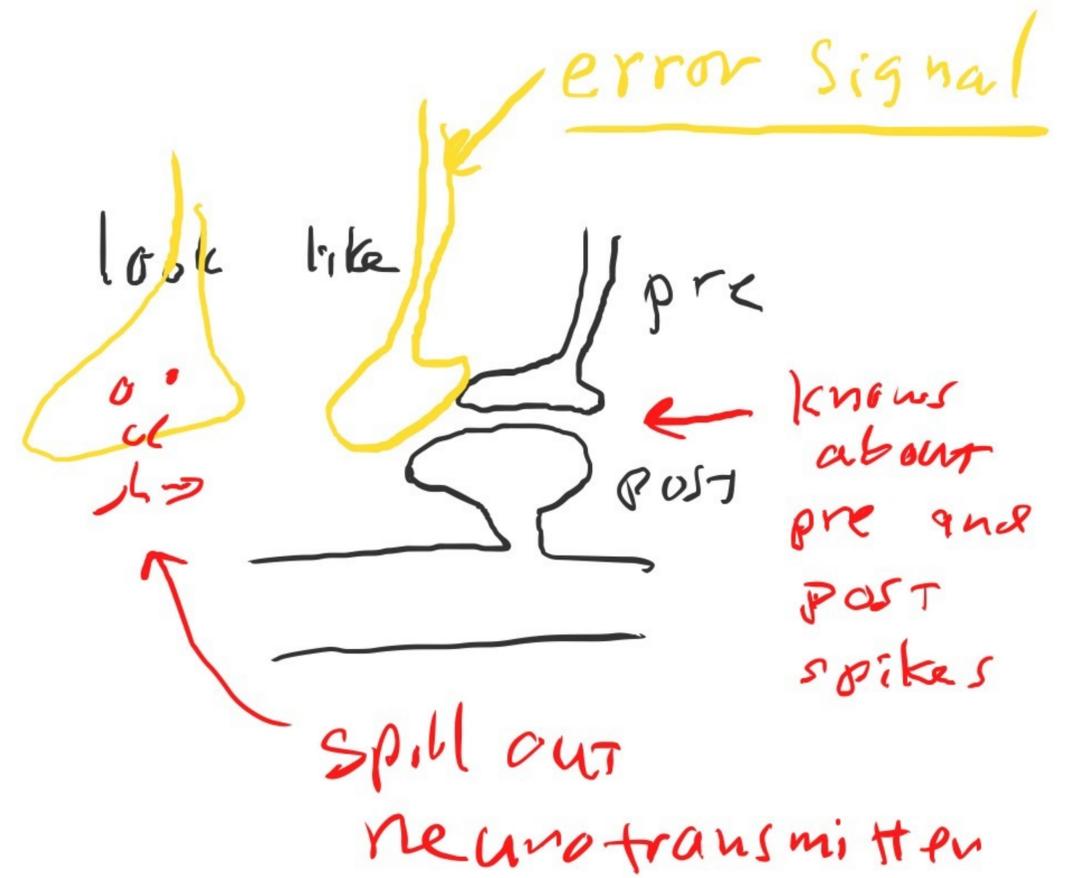
$$V_i = \phi \left(\sum_j W_{ij} V_j \right) - V_i$$

biologically plausible learning rules

$$\Delta W_{ij} = f(V_i, V_j)$$

← pre
← post

is this a powerful enough learning rule to do anything useful?



more correct:

$$y = \phi(\underline{w} \cdot \underline{x})$$

$$y = \underline{w} \cdot \underline{x}$$

→ Oja's rule

$$\Delta \underline{w} = \eta y (\underline{x} - y \underline{w})$$

$$= \eta \underline{w} \cdot \underline{x} (\underline{x} - \underline{w} \cdot \underline{x} \underline{w}) = \eta \underline{w} \cdot \underline{x} \underline{x} - \eta \underline{w} \underline{w} \cdot \underline{x} \underline{x} \cdot \underline{w}$$

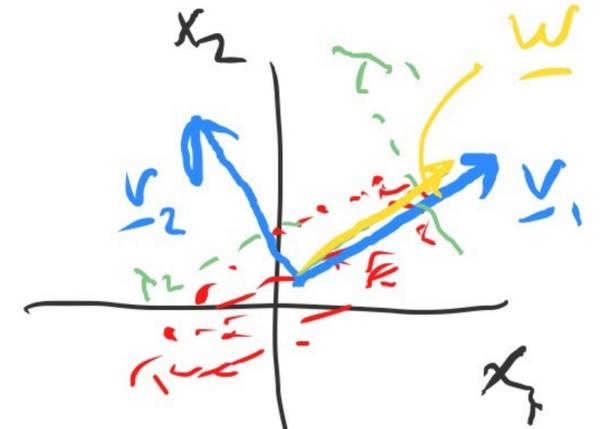
$$\eta \underline{w} = \eta \left[\underline{w} \cdot \langle \underline{x} \underline{x} \rangle - \underline{w} \underline{w} \cdot \langle \underline{x} \underline{x} \rangle \cdot \underline{w} \right]$$

$$\nabla \tilde{w} = \eta [\underline{c} \cdot \underline{w} - \underline{w} \cdot \underline{c} \cdot \underline{w} \cdot \underline{w}]$$

$$\underline{w} = \sum_k a_k \underline{v}_k$$

$$\underline{c} \cdot \underline{v}_k = \lambda_k \underline{v}_k$$

$$\underline{v}_k \cdot \underline{v}_l = \delta_{kl}$$



$$\nabla \sum_k a_k \underline{v}_k = \eta \left[\sum_k a_k \lambda_k \underline{v}_k - \left[\sum_e a_e^2 \lambda_e \right] \sum_k a_k \underline{v}_k \right]$$

$$\nabla a_k = \eta \left[\lambda_k - \sum_e a_e^2 \lambda_e \right] a_k$$

Assume all eigenvalues are different

$$\lambda_k = \sum_e a_e^2 \lambda_e \quad \forall k$$

$$a_k \geq 1 \quad a_{e \neq k} = 0$$

typically,
largest
eigenvalue
wms

$$\begin{aligned} & \sum_{e,m} a_e \underline{v}_e \cdot \underline{c} \cdot a_m \underline{v}_m \\ &= \sum_{e,m} a_e a_m \overbrace{\underline{v}_e \cdot \underline{v}_m}^{\delta_{em}} \lambda_m \\ &= \sum_e a_e^2 \lambda_e \end{aligned}$$

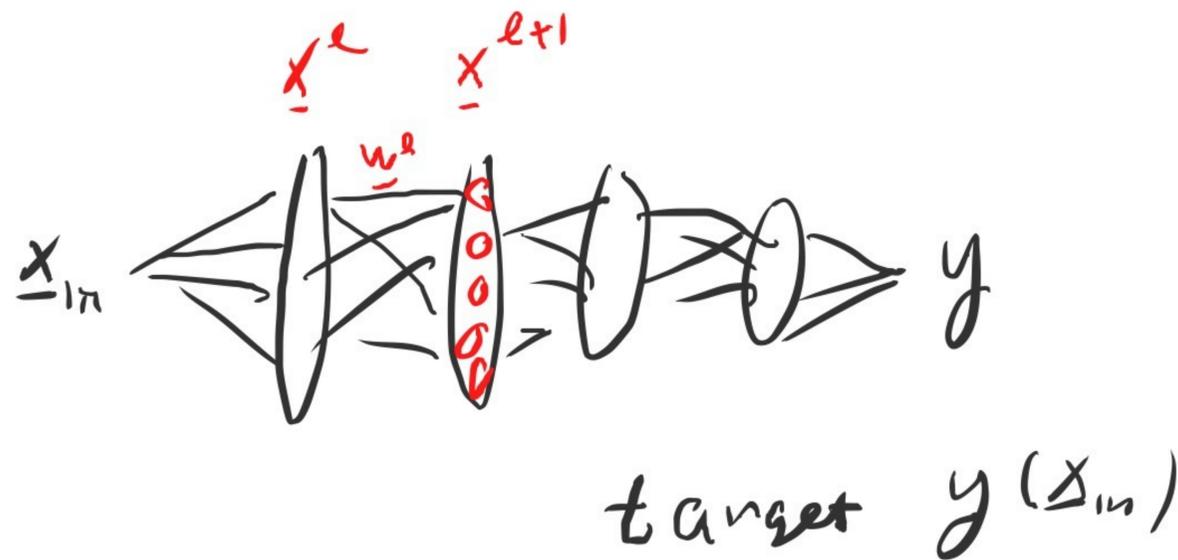
$$\underline{w} = \underline{v}_k$$

learning in real networks

1. vanilla deep networks

$$\underline{x}^{l+1} = \phi(\underline{h}^l)$$

$$\underline{h}^l = \underline{w}^l \cdot \underline{x}^l$$



$$\mathcal{L}(\underline{x}_{in}, y(\underline{x}_{in}))$$

$$\Delta \underline{w}^l = -\gamma \frac{\partial \mathcal{L}}{\partial \underline{w}^l}$$

$$\frac{\partial \mathcal{L}}{\partial \underline{w}^l} = \frac{\partial \mathcal{L}}{\partial \underline{h}^l} \cdot \frac{\partial \underline{h}^l}{\partial \underline{w}^l} = \frac{\partial \mathcal{L}}{\partial \underline{h}^l} \underline{x}^l$$

$$= \delta^l \underline{x}^l$$

$$\frac{\partial \mathcal{L}}{\partial \underline{w}^l} = \frac{\partial \mathcal{L}}{\partial \underline{h}^l} \underline{x}^l$$

$$\frac{\partial \mathcal{L}}{\partial \underline{h}^l} = \frac{\partial \mathcal{L}}{\partial \underline{h}^{l+1}} \cdot \frac{\partial \underline{h}^{l+1}}{\partial \underline{h}^l}$$

$$\underline{h}^{l+1} = \underline{w}^{l+1} \cdot \underline{x}^{l+1}$$

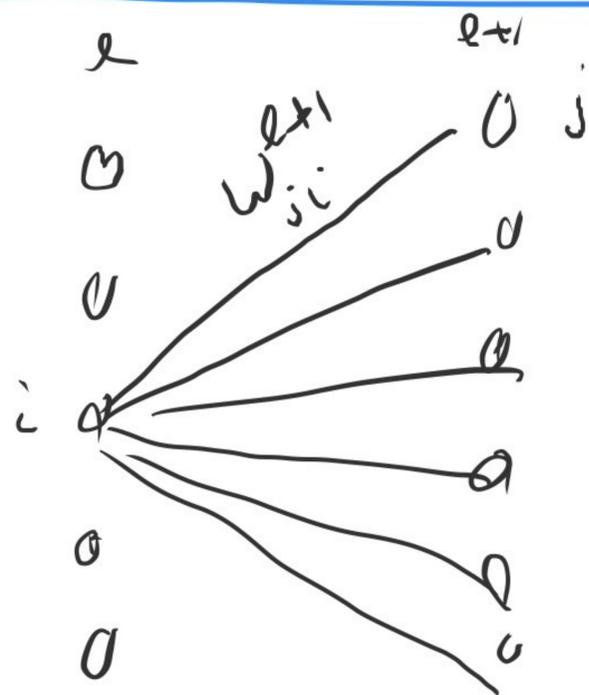
$$= \underline{w}^{l+1} \cdot \phi(\underline{h}^l)$$

$$\frac{\partial \underline{h}^{l+1}}{\partial \underline{h}^l} = \underline{w}^{l+1} \odot \phi'(\underline{h}^l)$$

$$\delta^l = \frac{\partial \mathcal{L}}{\partial \underline{h}^l}$$

$$\delta^l = \delta^{l+1} \cdot \underline{w}^{l+1} \odot \phi'(\underline{h}^l)$$

$$\delta_i^l = \sum_j \delta_j^{l+1} w_{ji}^{l+1} \phi'(h_i^l)$$



$$\delta^l = \delta^{l+1} \cdot \underline{w}^{l+1} \odot \phi'(h^l)$$

↑
outgoing weights

① Tim Lillicrap: replace \underline{w}^{l+1} with a random matrix \underline{A}
(feedback alignment)

- works well on simple problems
- doesn't work on hard problems

② upgrade: train

\underline{A}^l		\underline{w}^l	\underline{w}^{l+1}
x	→	b	l+1
0	→	0	
0	←	0	
0	←	0	
0	←	0	
0	←	0	

training:

$$\Delta \underline{A}^l = \eta \langle \underline{x}^l \underline{x}^{l+1} \rangle$$

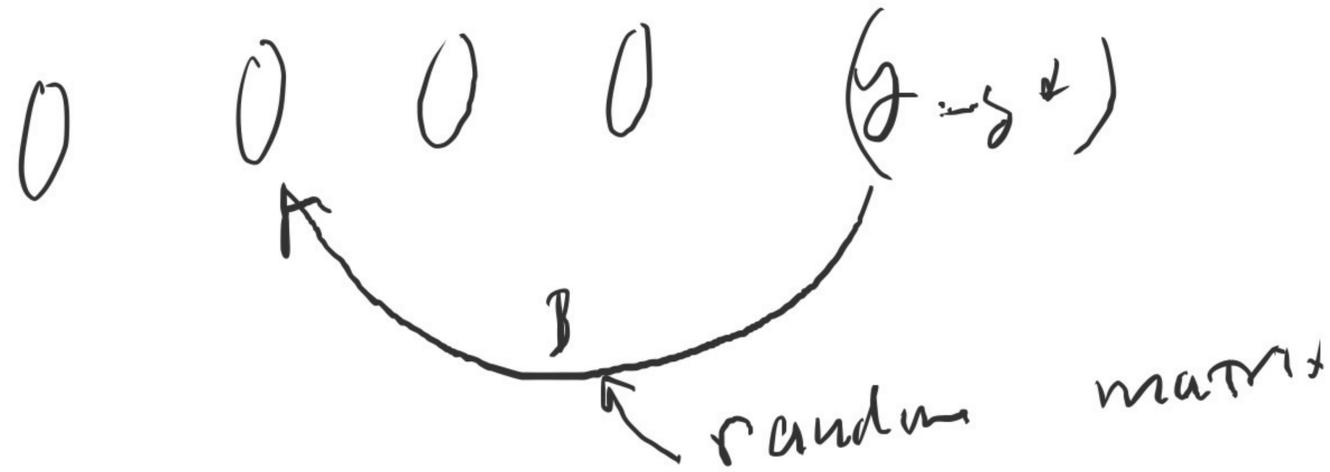
during training: add white noise to \underline{x}^l

$$\Delta \underline{A}^l = \eta \langle \underline{x}^l \underline{x}^l \cdot \underline{w}^{l+1 T} \rangle - \Sigma \underline{A}^l$$

$$= \eta \langle \underline{x}^l \underline{x}^l \rangle \cdot \underline{w}^{l+1 T} - \Sigma \underline{A}^l$$

$$= \eta \underline{w}^{l+1 T} - \Sigma \underline{A}^l$$

Direct feedback alignment



$$\Delta \underline{w}^l = \underline{B}^r (y - y^*)$$

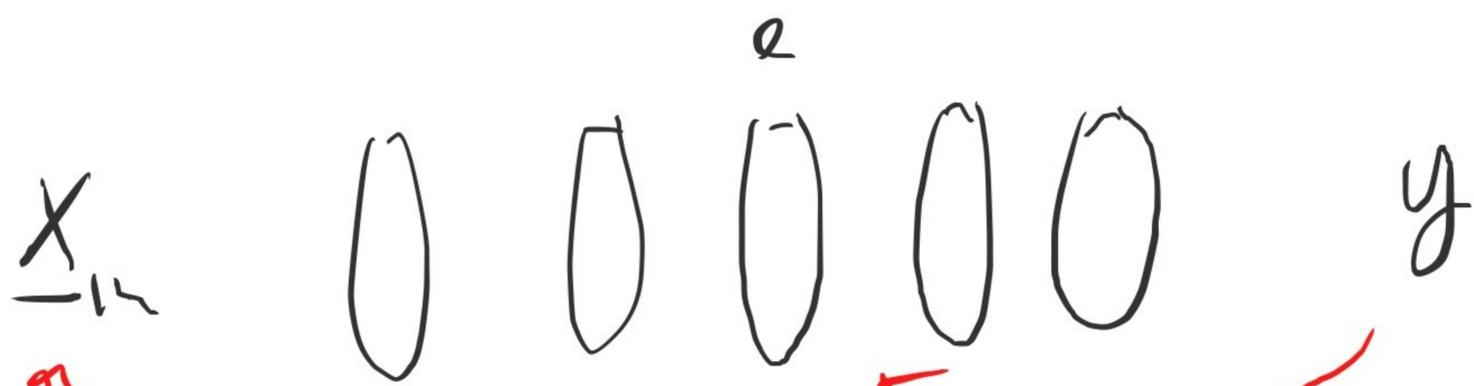
$$x^l \quad \Delta w^l = -\gamma \sum (y - y^*)$$

A diagram showing a vertical column of five zeros on the left. To the right of this column is a horizontal line labeled w^l . Further to the right is a target value y . Red arrows point from each zero in the column towards the target value y .

Works pretty well
on problems that
aren't too hard

Bottleneck approach

3-factor Hebb rule



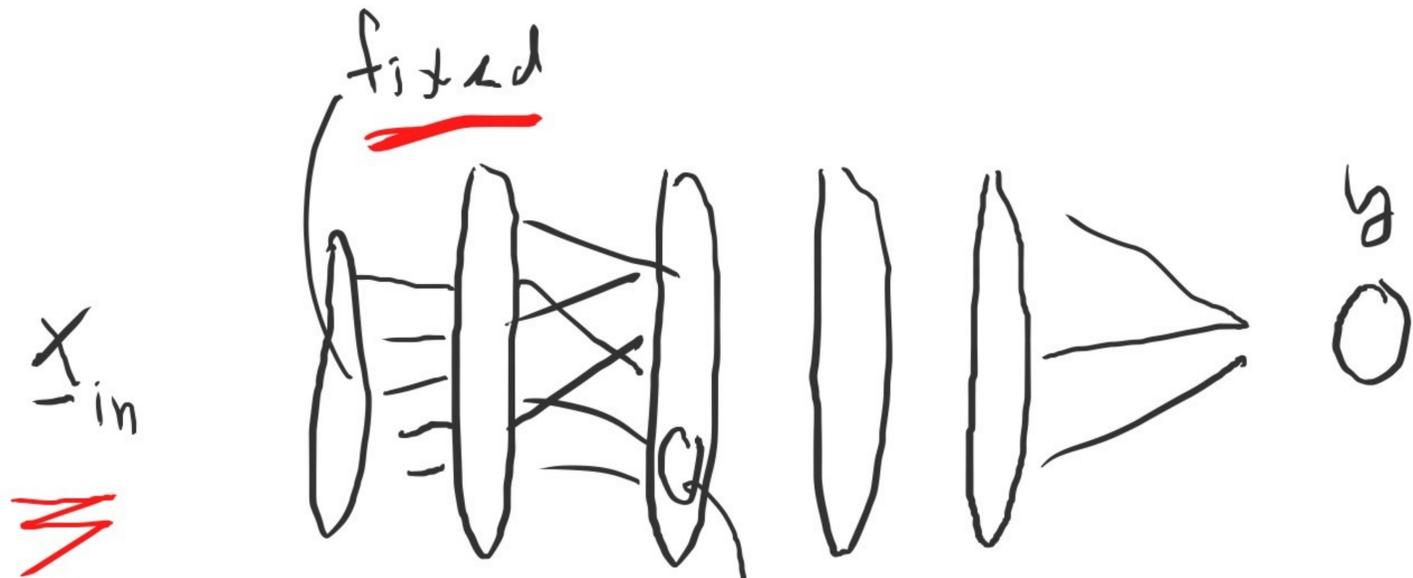
Tali Tishby: irrelevant information is getting squeezed out.

maximize

$$I(x^r, y) - \beta I(x_{in}, x^r)$$

replace

Info w/ some other measure



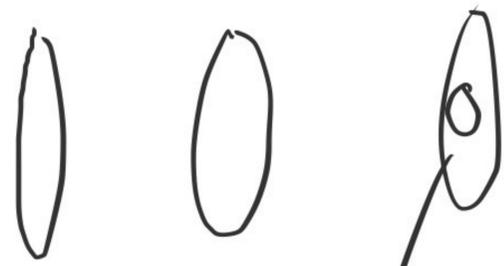
X_i^l : goal is to predict y

$$\mathcal{L} = \frac{1}{2} \sum_{l,j} (x_i^l - y^*(x_{in}))^2$$

local: Hebb rule

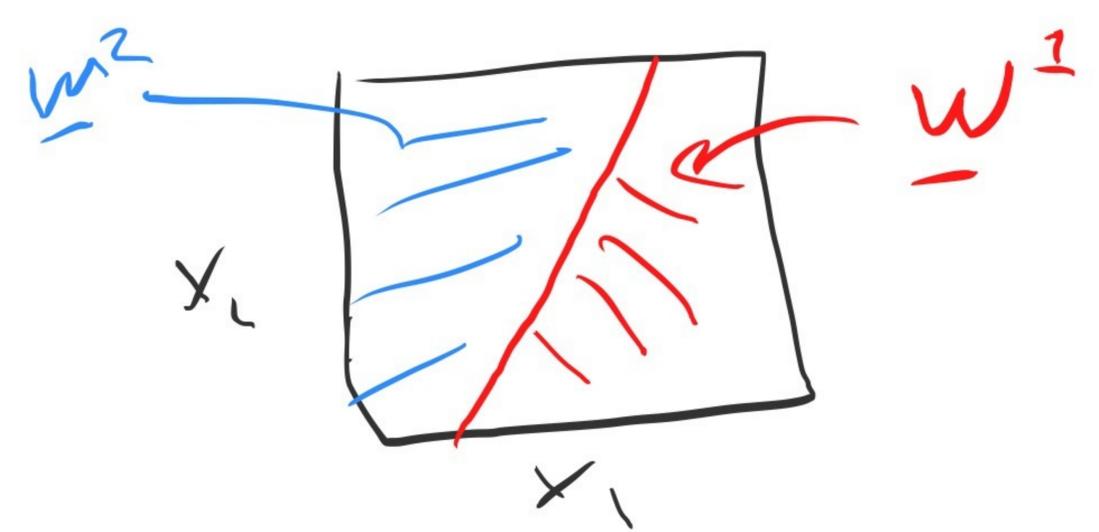
$$= \frac{1}{2} \sum_{l,j} \left(\sum_j w_{ij}^{l-1} x_j^{l-1} - y^*(x_{in}) \right)^2$$

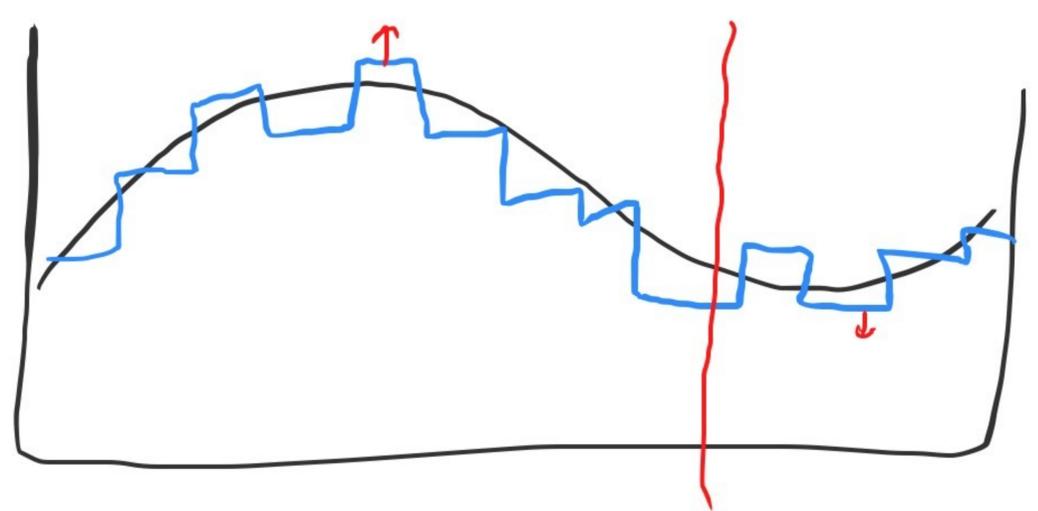
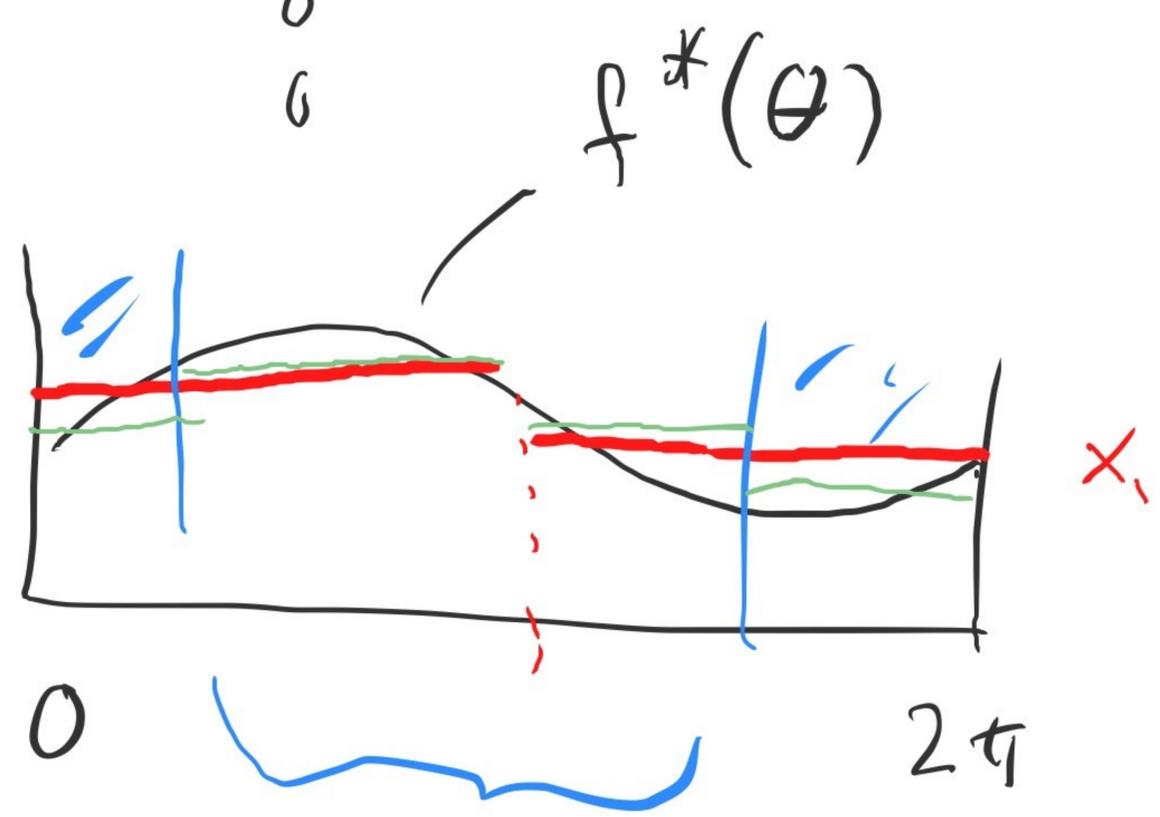
$$\Delta w_{ij}^{l-1} = -\eta \frac{\partial \mathcal{L}^{l-1}}{\partial w_{ij}^{l-1}} = -\eta \underbrace{(x_i^l - y^*)}_{\text{error (post)}} x_j^{l-1} \quad \leftarrow \text{pre-synaptic}$$



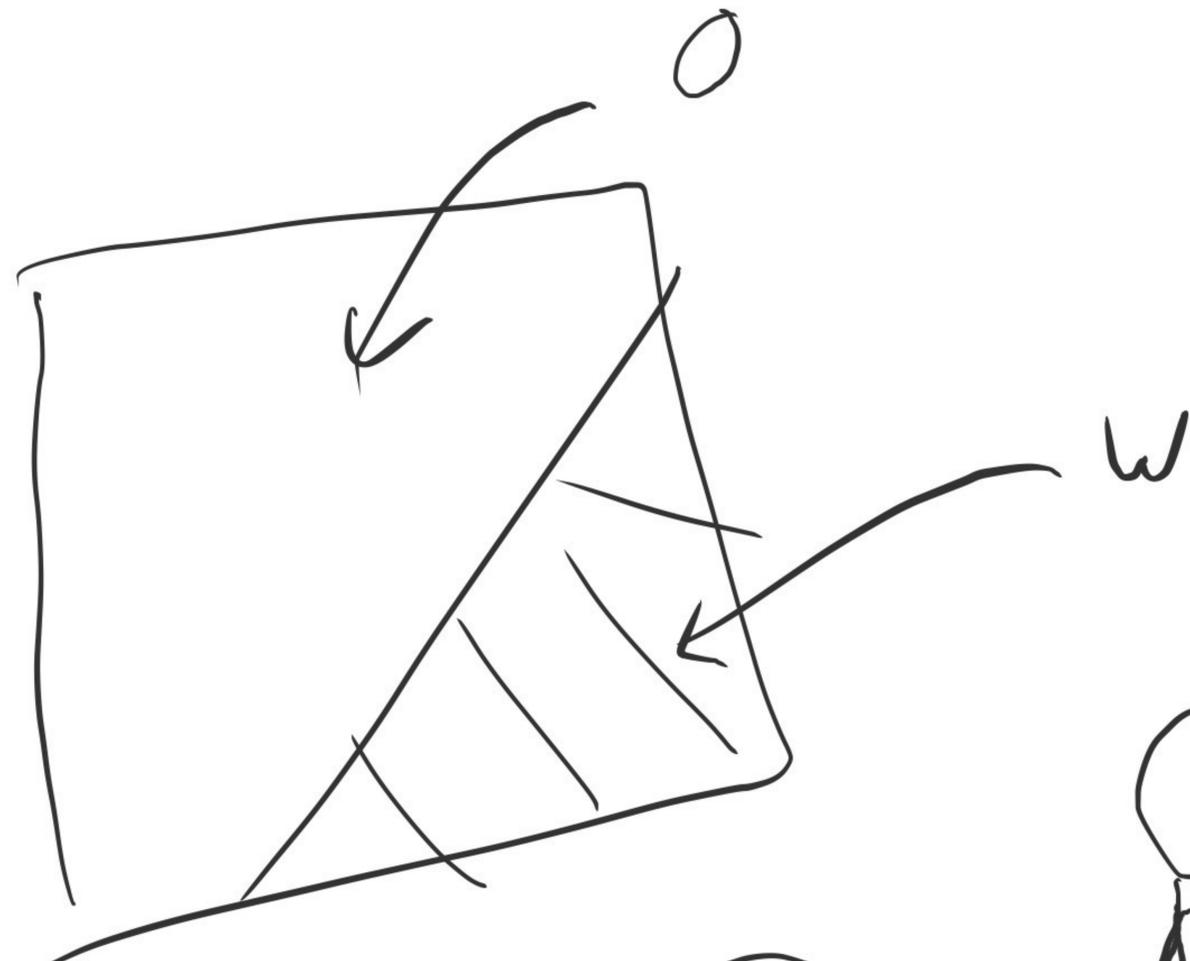
$$x_i^l = \sum_j w_{ij}^{l-1, c(x_m)} x_j^{l-1}$$

- each neuron has a bank of weights
- which one it chooses depends on input
- pre-chosen





1. Convex
2. gets better with depth



GLN
 gated linear networks

$$\underline{r}^l = f(\underline{w}^l \cdot f^{-1}(\underline{r}^{l-1}))$$

$$= f(\underline{w}^l \cdot \underline{w}^{l-1} \cdot f^{-1}(\underline{r}^{l-2}))$$



each branch
 is gated
 on or off
 (inhibits
 neurons)

DGN
 dendritic gated
 network

Summary:

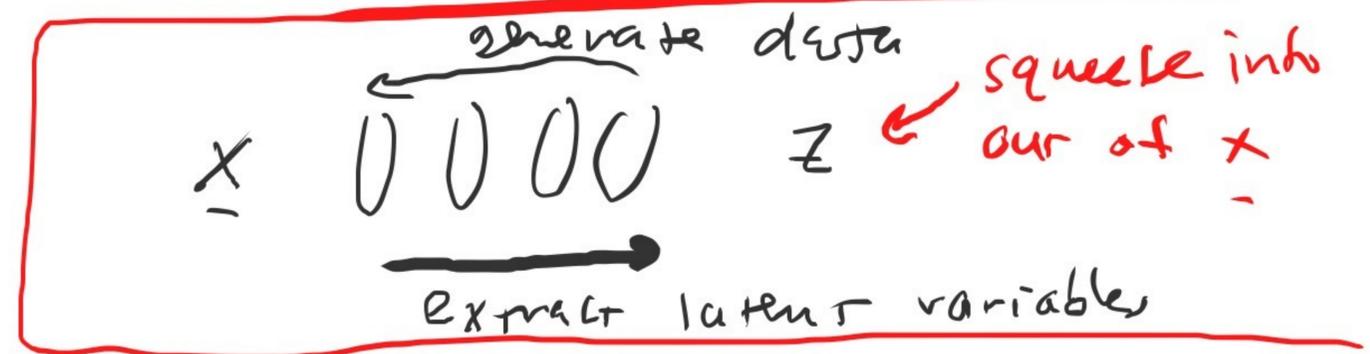
- backprop works really well
- NOT biologically plausible
 1. needs forward and backward pass
 2. neurons need to know their outgoing weights

lots of fixes

- mainly hacks

- GLM (provably converges - convex)

- what the brain uses is unknown



the brain mainly uses unsupervised learning

x , model $P_{\theta}(x)$

$$\mathcal{L} = - \sum_{x \sim P} \log P_{\theta}(x)$$
$$\Delta \theta = - \nabla \frac{\partial \mathcal{L}}{\partial \theta}$$

goal of most science:
bridge levels of description

neuroscience

lower level: neural activity
upper level: behavior

goal is to understand behavior in terms of
neural activity

neurons (micro level)

$$C \frac{dV_i}{dt} = \underbrace{\left(\text{single neuron model} \right)}_{\sum_j W_{ij} (V_j - E_j) g_j(t)}$$

conductance changes due to pre-synaptic spikes

$$\Delta W_{ij} = \dots$$

behavior = $f(\underline{S})$

spikes on all the neurons (output neurons)

10^{11} eqns

10-100 dimensional

understand:

solve analytically mapping from spikes to behavior

2 parts:

1. What are the equations?
(including initial conditions)
2. What's the solution?

Weather prediction

Simplified equations

make up kind of reasonable equations and cross our fingers. (usually simpler)

Biophysics

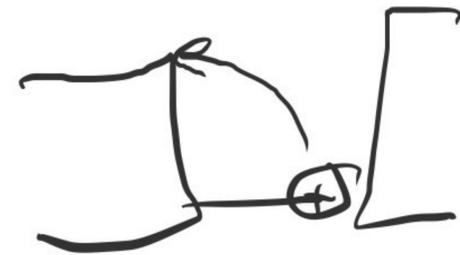
$$C \frac{dv}{dt} = I$$

$$I = - \sum_x g_x (v - E_x)$$

conductance reversal potential

derive all of biophysics once you

know g_x



constant (leak)

voltage dependent (HH)

- add geometry (cable eqn.)

concentration - dependent

axons

synaptic dynamics



networks

understand this quantity

$$\sum_j w_{ij} (V_i - \varepsilon_j) g_j(t)$$

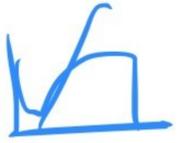
fast fluctuations

static piece



$$\sum_j \tilde{w}_{ij} V_j + \sum_j \tilde{w}_{ij} \delta V_j(t)$$

famous analogy



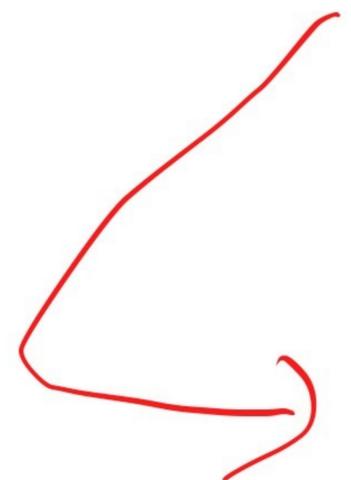
random =>

$$\bar{V}_\varepsilon, \bar{V}_S$$

are all that matters

not random

overlaps



large sums are gaussian

weights fixed

- Switch to learning
- we wrote down learning rules in supervised setting.
- now things get really hard

$$\frac{dw_{ij}}{dt} = f(\underline{w})$$

← averaged over activities

we don't know this one

↑ 10^{14}

$$\frac{dw_{ij}}{dt} = f(w_{ij}, x)$$

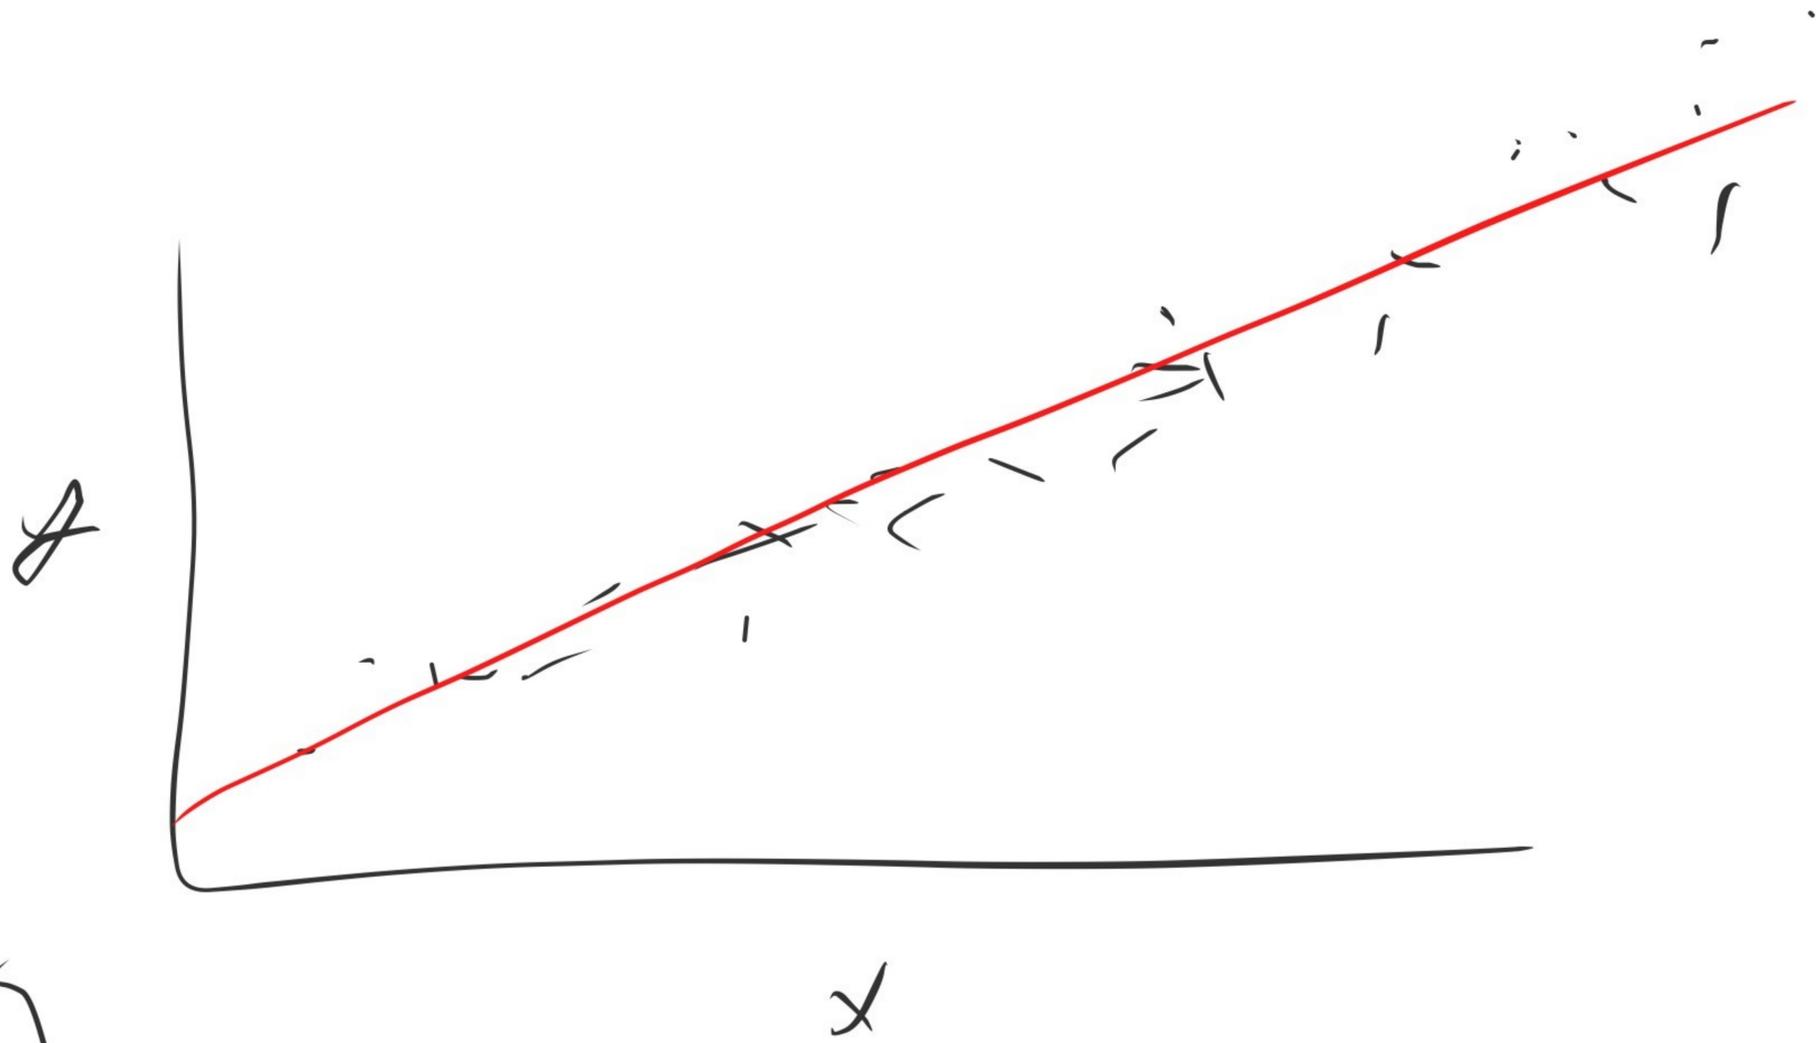
← activities of lots of neurons

$$f(w_{ij}, x_i, x_j)$$

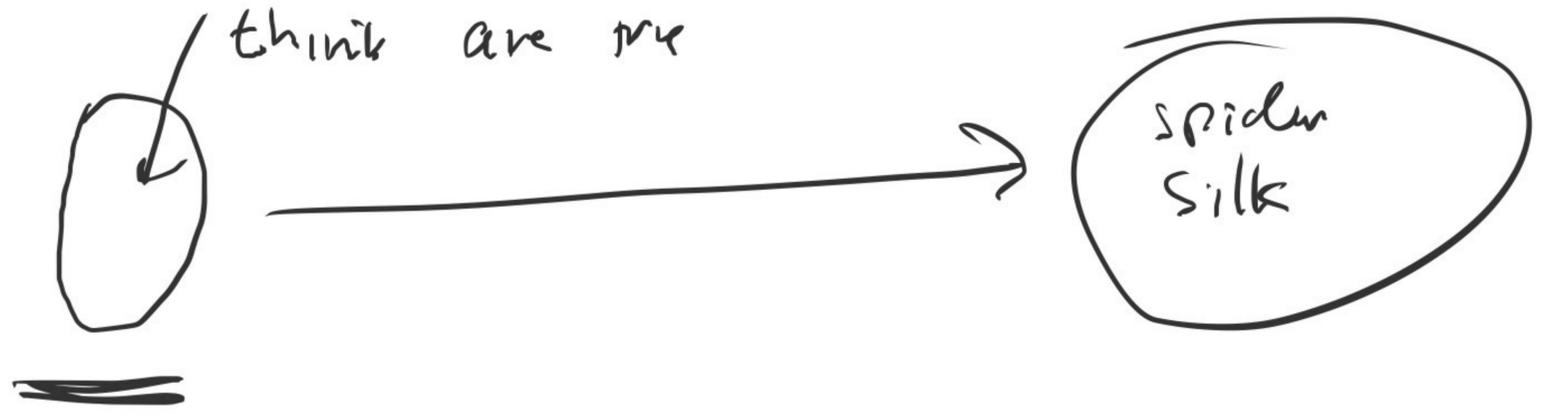
← standard picture: pre-post

$$\frac{dx_i}{dt} = f\left(\sum_j w_{ij} x_j\right)$$

← equations highly nonlinear



if x, y, z true then w



$$\frac{D}{n} = \underline{\underline{0.138}}$$

