



Neural networks

basic perceptron



non-linear activation

Hooking them up into a networks



Train \rightarrow determine some loss and minimize it

$$\text{loss} = \sum (\sigma_i - t_i)$$

\sum over all training values \rightarrow $i \geq 0$ \uparrow output \uparrow training value

overall training values output value


To actually train:
gradient descent
of same kind

→ updating weights and biases
(back propagation)

Convolutional neural networks

↳ used for images

→ input for CNN → output = no tree



→ output = yes tree

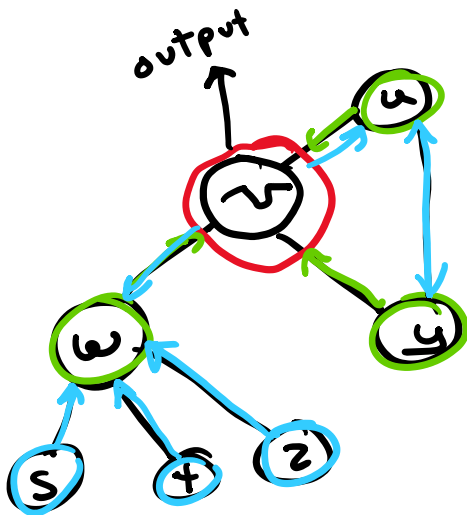
GCN → graph convolutional networks

→ we move "inputs" across
the graph centered on vertices

the graph centered on vertices

Q1: How do we determine inputs/output for training?

Q2: How do we incorporate graph's topology?



From our prior discussion


→ we can consider the "state" of a vertex as a function of its neighborhood

GCN idea: we propagate states/values and outputs along edges in the graph

→ consider the broader k -hop neighborhood around some vertex v



Formulation



Formulation

\bar{O}_v = output

\bar{h}_v = hidden states

\bar{x}_v = features

$\bar{x}_{co[v]}$ = edge features

$k = \# \text{ hops}$

$$\bar{h}_v = f(\bar{x}_v, \bar{x}_{co}, \bar{h}_{ne}, \bar{x}_{ne})$$

↑
transition function
(update states)

↑ ↑
neighbor states
and features

$$\bar{O}_v = g(\bar{h}_v, \bar{x}_v)$$

↑
output function
(updates output)

General approach

we have global H, X

→ Iteratively update H

→ Iteratively update H

$$H^{t+1} = F(H^t, X)$$

$$\text{supervised loss} = \sum_{i=0}^{|\bar{O}|} (t_i - o_i)$$

t_i = i th data in
training set

o_i = i th output
from GCN

→ Update weights / biases W, B
via gradient descent


Note: $H^{t+1} = F(H^t, X)$

$$F(H^t, X) \sim \sigma(H^t, X, W^t, \delta^t)$$

To account for varying degrees

→ aggregation via averaging
or otherwise "reducing"
our neighbor inputs

→ we need to do this at every level



$$\bar{h}_0 = \bar{x}_0$$

$$\bar{h}_k = \sigma \left(\omega_k \frac{\sum_{u \in N(v)} h_u}{|N(v)|} + \beta_k h_{k-1} \right)$$

$$\sigma = \bar{h}_k$$

↑
"avg." reduction

Notes:

- ω and β are consistent across the whole network
- Aggregation can be done in any number of way
- Learning for ω, β is similar to training a regular NN

Overall: similar generalization
and "learning power" of NNs,
but we capture graph topology

Code Mode

We want to predict clusters
on same graph

We have 4 class labels, to
create 4 clusters ↗

Only 1 vertex
labeled per class
(semi-supervised)

Network: Zachary Karate Club

vertices: Karate students

edges: friendships