
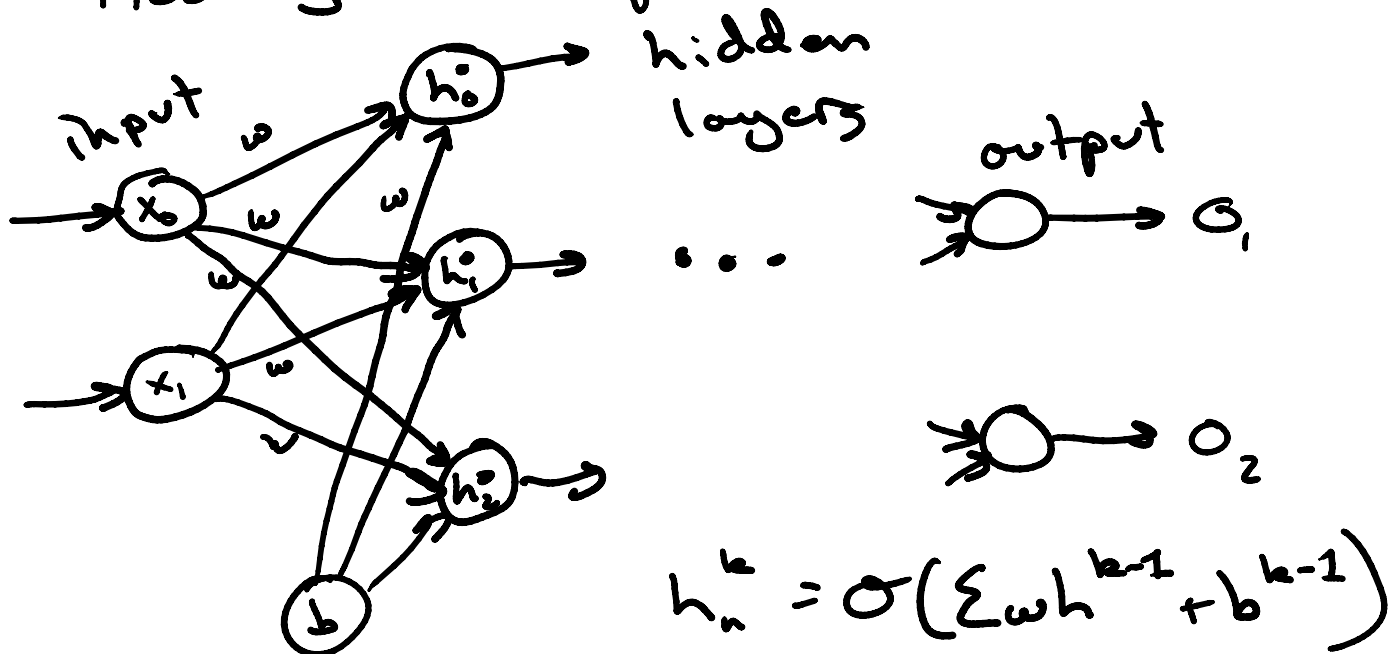


Neural networks

basic perceptron



Hooking them up



Train \rightarrow determine some loss

$$\text{loss} = \sum_{i \in \bar{o}} (o_i - t_i)$$

\leftarrow over all training data
 \uparrow output from NN
 \uparrow expected training value

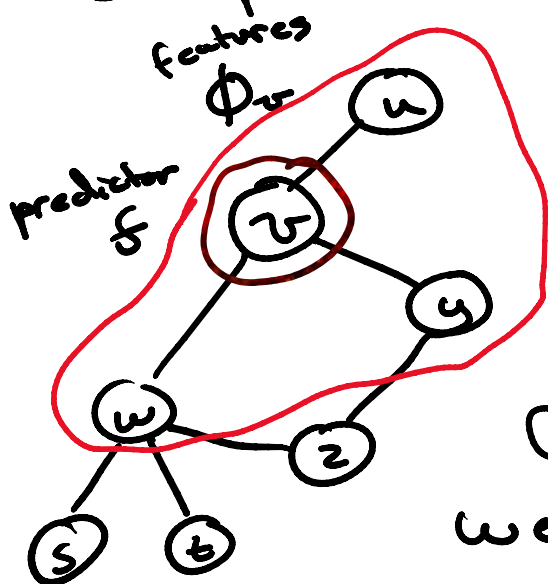
to actually train: gradient descent
=> updating weights and biases
(backpropagation)

Convolutional neural networks

→ used for image processing



Graph Neural networks (GNN)



From a few classes ago:


- we predicted some output on v , which was dependent on $N(v)$

GNN basic idea:

we propagate outputs/states along edges with vertices acting as neurons/NNs.

...
acting as neurons/NNs

Note: we can consider the broader network around v to some k -hops



Formulation

\bar{O}_v = output

\bar{h}_v = hidden states

\bar{x}_v = features

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

$$\bar{h}_v = f(\bar{x}_v, \bar{x}_{e[v]}, \bar{h}_{ne}, \bar{x}_{ne}) \quad \bar{h}_{ne}[v] = \text{neighbor state}$$

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

$\bar{x}_{ne}[v]$ = neighbor features

f = transition function

→ updates states

g = output function

→ updating output

General approach:

- consider global H, X matrices

- we iteratively update H

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

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

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

$t_i = i^{\text{th}}$ data in training set for v

$o_i = i^{\text{th}}$ output from GNN on v

To actually update:

we have some weight matrix W and possibly some bias matrix B

W, B are updated via gradient descent or similar


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

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

To account for varying degrees

→ we can aggregate by averaging neighbor inputs (avg. degree)

→ Need to potentially account for this at every level



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

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

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

(actual k)

Note: we train these

Notes:

- ω and β are consistent across the networks
- Aggregation can be done in any number of ways
- Learning for ω, β is similar to B.P. on regular old NNs

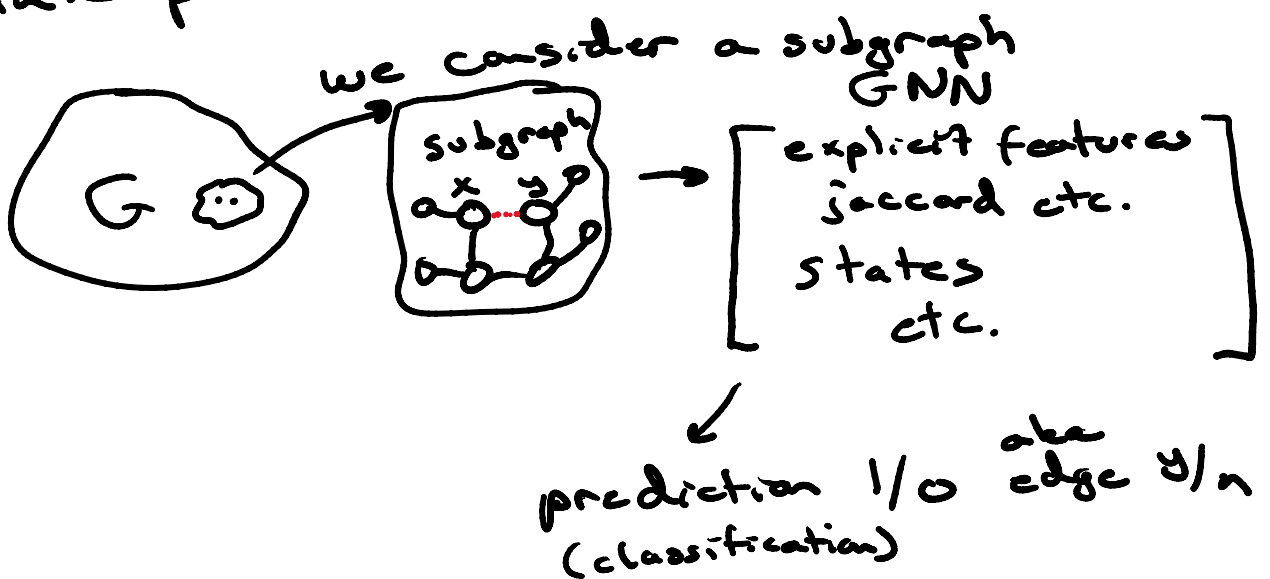
Overall: similar power to NNs but we include the additional per-vertex edge-based

per-vertex edge-based computation graph

Applications

(almost everything we've talked about)

- Link prediction



- Centrality


→ we can use (semi) supervised
along with regression

- Clustering and C.D.

→ we can do classification with
some set of "seed" vertices
(more soon)

- graph measures
(subgraph mining measures for classification)

→ Note: we can include an additional layer: aggregation layer



- graph alignment

Use our aggregation layer to construct some "embedding"

Recall: theoretically fastest isomorphism algorithms convert a graph topology to a string

→ similar idea

Code mode

Data: Korate network with 4 class labels

0 is unlabeled

Problem: classification of unlabeled
set of vertices

Approach: a GNN