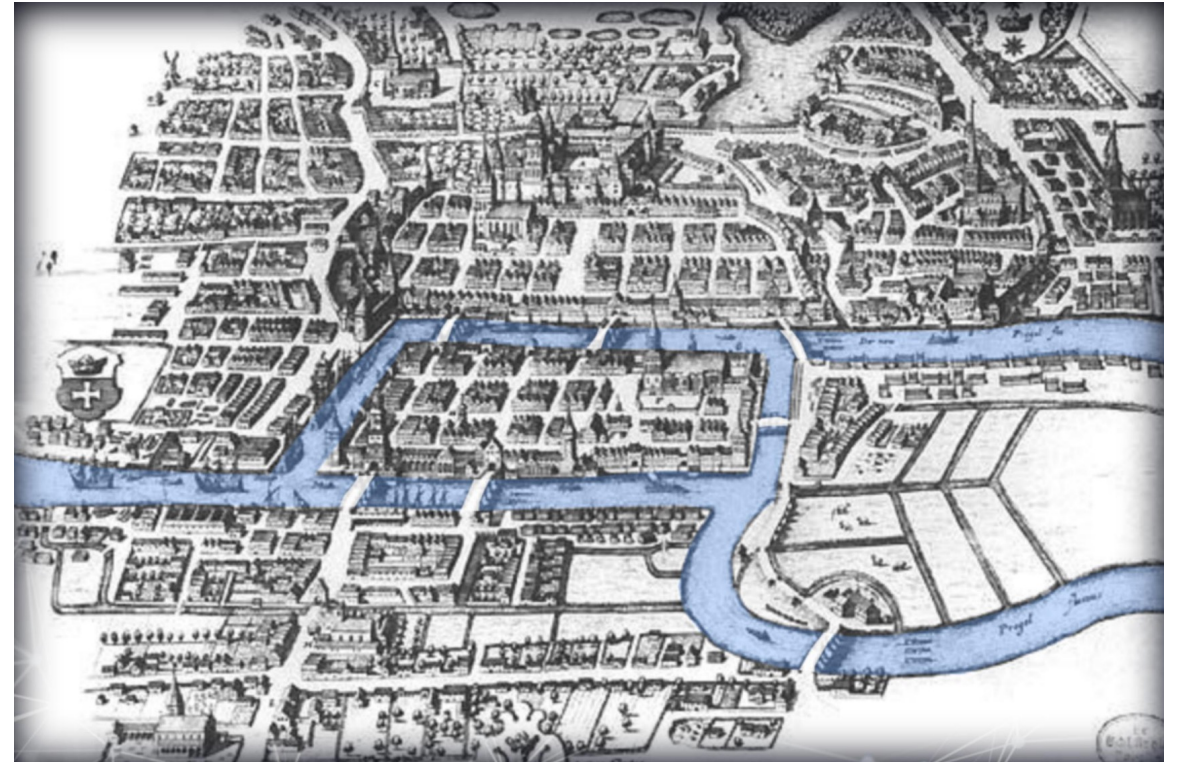



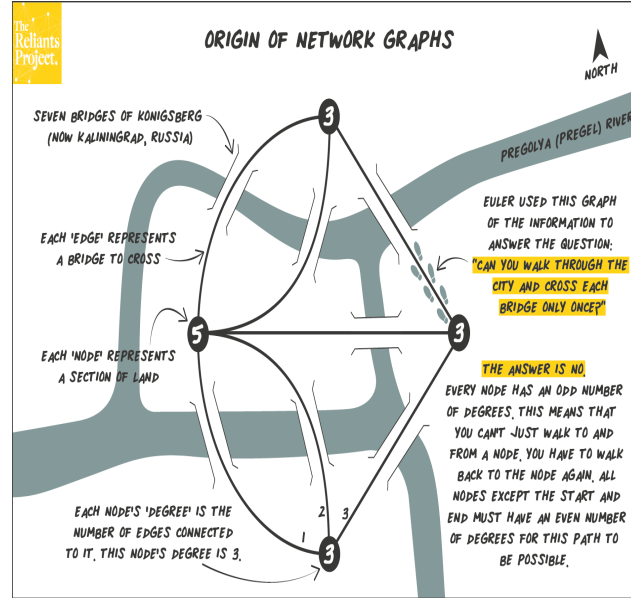

# From Graphs to Graph Neural Networks: Foundations and Applications in Healthcare

Soheila Molaei

# Seven Bridges of Königsberg


Can you walk through the city crossing all bridges only **once**?







**Not possible!**

# Why Graphs are important?!




## Graphs from social networks

- people and their interactions
  - directed (Twitter) and undirected (Facebook)
- typical ML tasks
  - Node prediction
  - Link(edge) prediction






# Biomedical Data: Molecular Scale



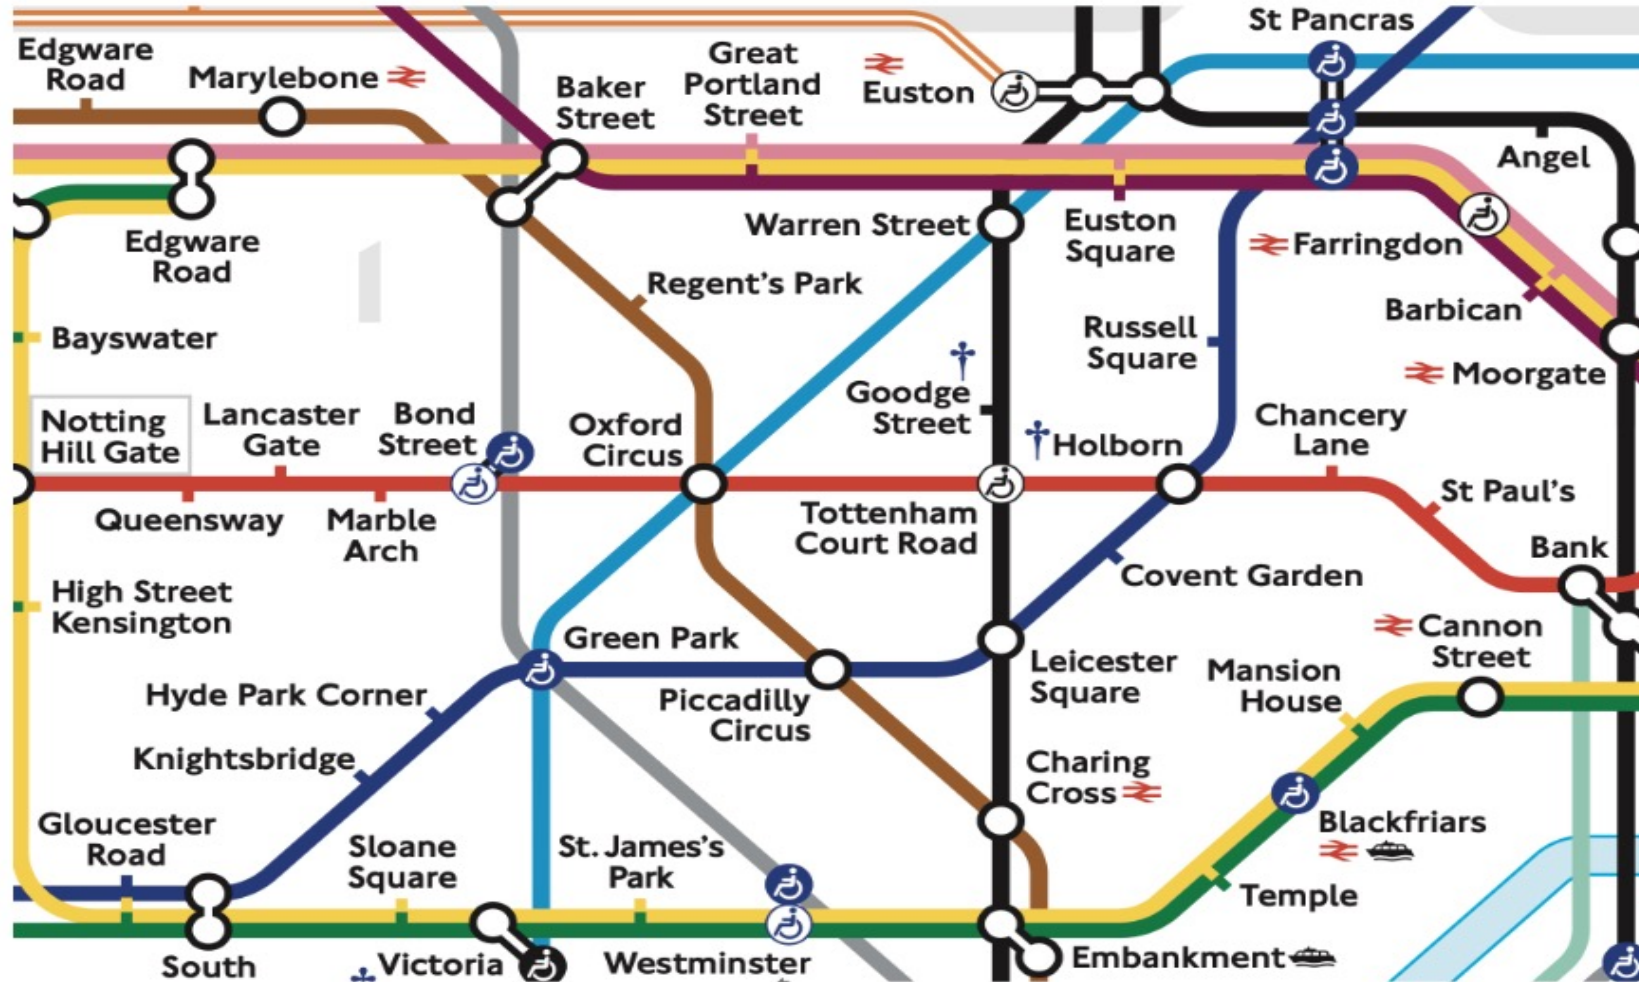
**Molecular scale:** Small molecule drugs can be represented as graphs relating their **constituent atoms** and **chemical bonding** structure. Complex molecules, such as proteins, can be represented as graphs capturing spatial and structural relationships between their amino acid residues.

## Changing Molecule to Graph



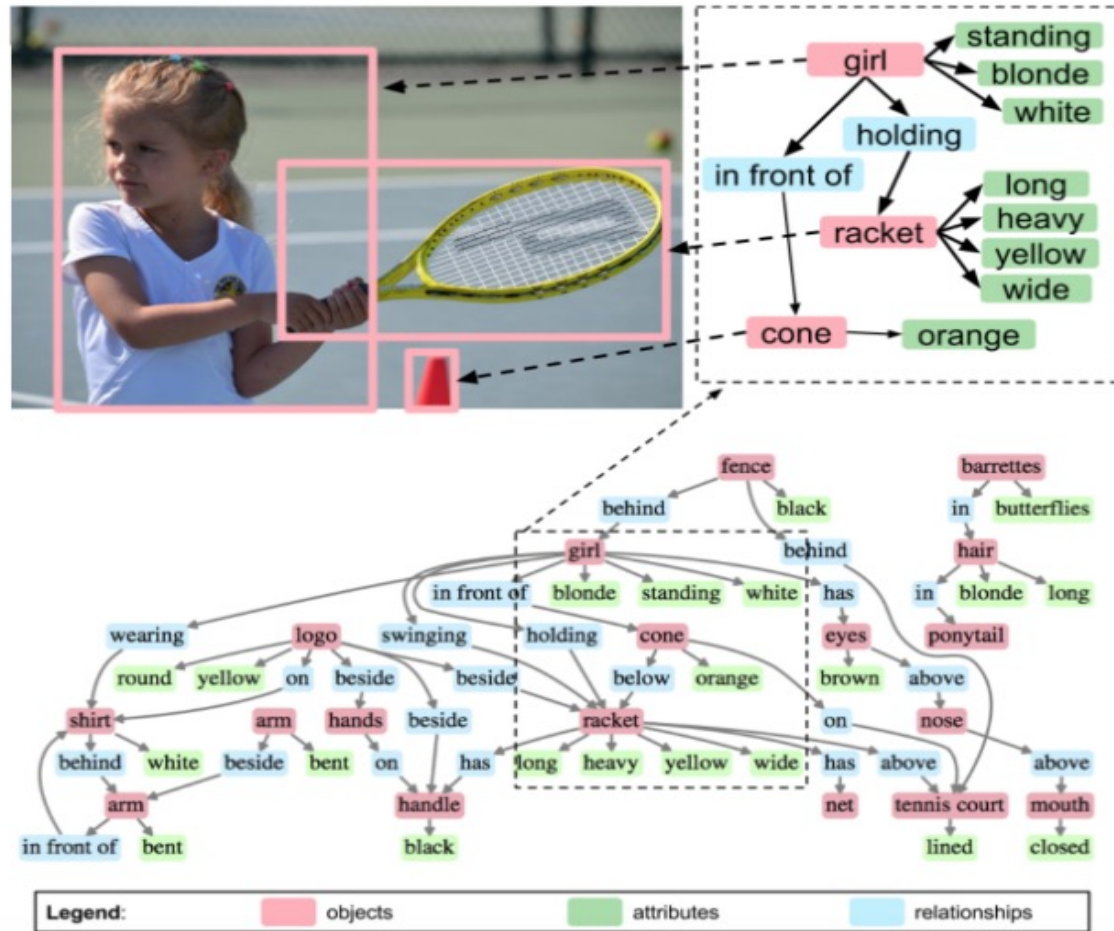


# Road Network




Traffic networks: An excerpt of the London Tube of Zone 1, showing different lines.

# Computer Vision: Scene Graphs








ML algorithms make assumptions about the data, e.g., the data points are

**independent and identically distributed (i.i.d.):**

- independence: no need to model the **dependencies**,
- identical distribution: **generalization** guarantees possible to new/unseen data points.

These assumptions are unrealistic in the context of graphs.




Simple, undirected, unweighted graphs attributed with node features.

Set of vertices/nodes


Set of undirected edges

Node feature matrix, which stores a feature vector for each node  $u$ .



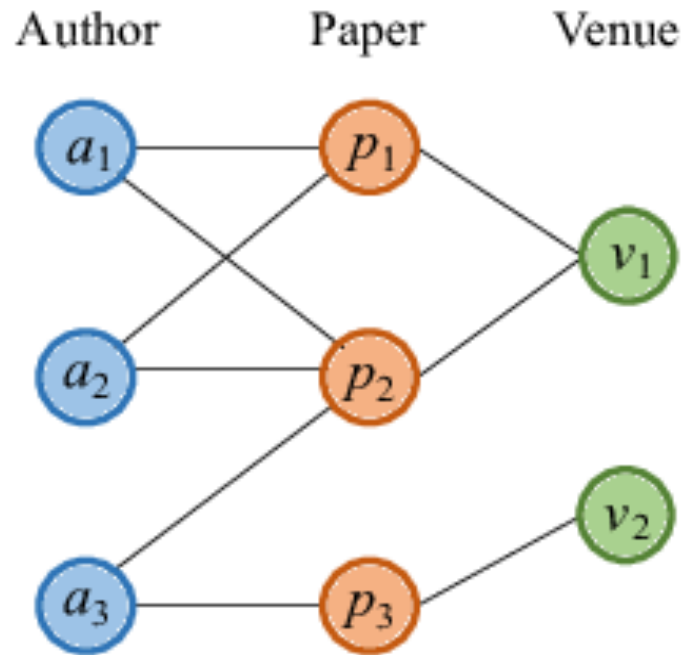
## Features:

- Node features can be domain-specific **attributes**, or **node degrees**, or simply **one-hot encodings**.
- We can extend the class of (attributed) graphs to include edge features .




Simple graph with loop

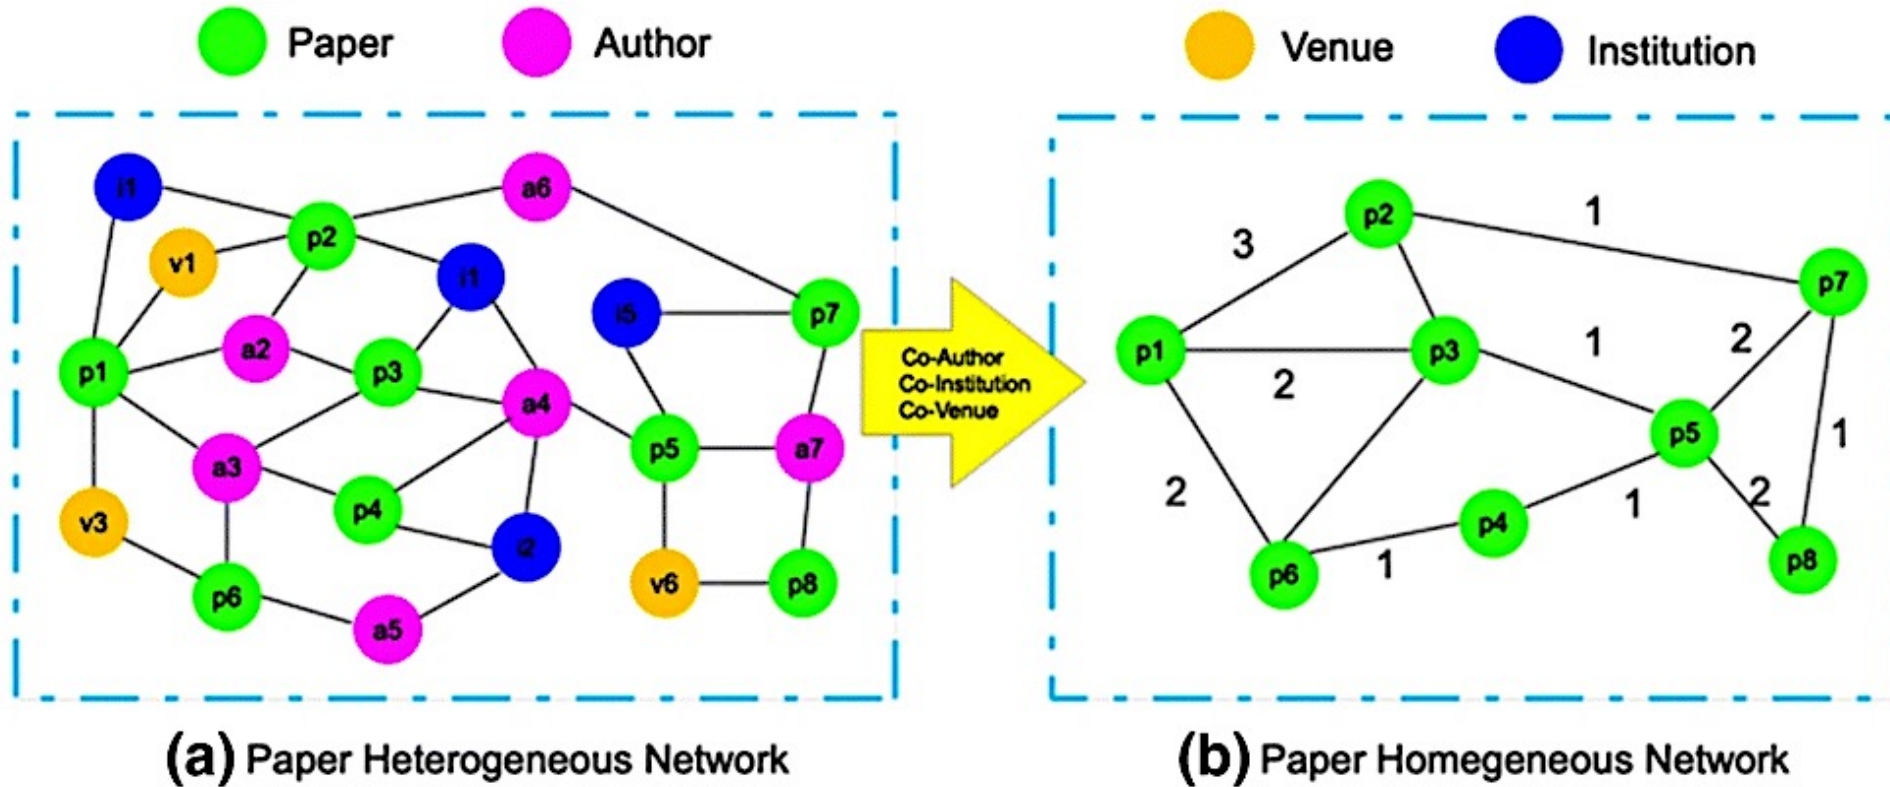
	A	B	C	D
A	0	1	1	1
B	1	0	0	0
C	1	0	1	1
D	1	0	1	0

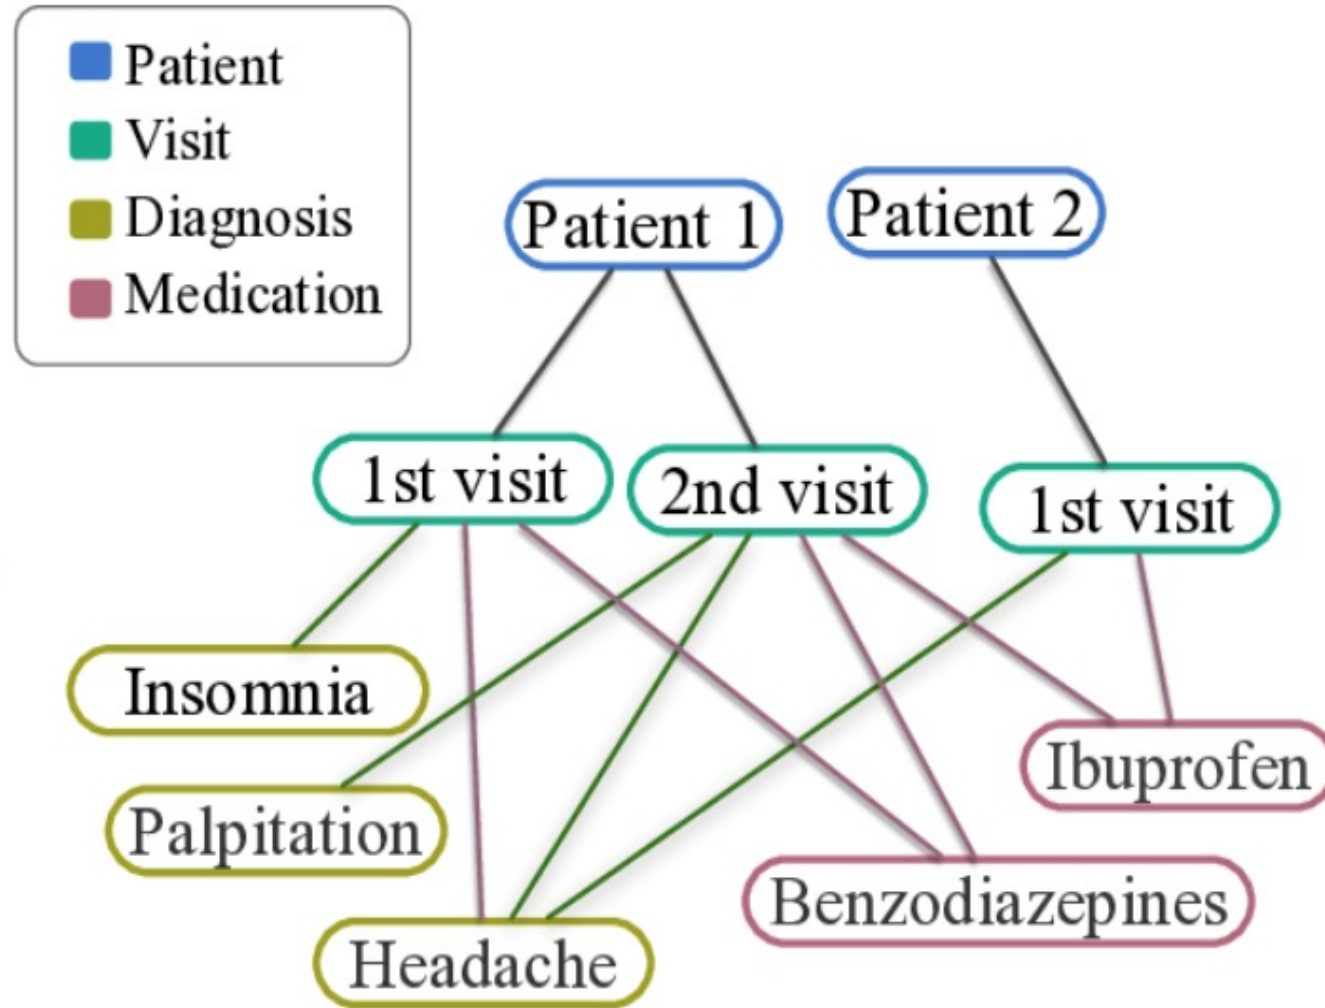


PAP homogenous subgraph





PA heterogenous subgraph






- A **graph neural network** is a neural model that we can apply directly to graphs without **prior knowledge** of every component within the graph. GNN provides a convenient way for node level, edge level and graph level prediction tasks.
- In GNNs, neighbours and connections define nodes. If we remove the neighbours and connections around a node, then the node will lose all its information. .







Graph.




Messages.




Propagatic








A graph (left) and an illustration of message passing on this graph with respect to the target node A for 3 iterations (right). Directed arrows depict the **messages**, and yellow boxes denote **aggregation**. At least 3 iterations are needed to get information from **all** nodes, i.e., F will not pass any messages to A with .





A graph (left) and an illustration of message passing on this graph with respect to the target node A for 3 iterations (right). Directed arrows depict the **messages**, and yellow boxes denote **aggregation**. At least 3 iterations are needed to get information from **all** nodes, i.e., F will not pass any messages to A with .





A graph (left) and an illustration of message passing on this graph with respect to the target node A for 3 iterations (right). Directed arrows depict the **messages**, and yellow boxes denote **aggregation**. At least 3 iterations are needed to get information from **all** nodes, i.e., F will not pass any messages to A with .



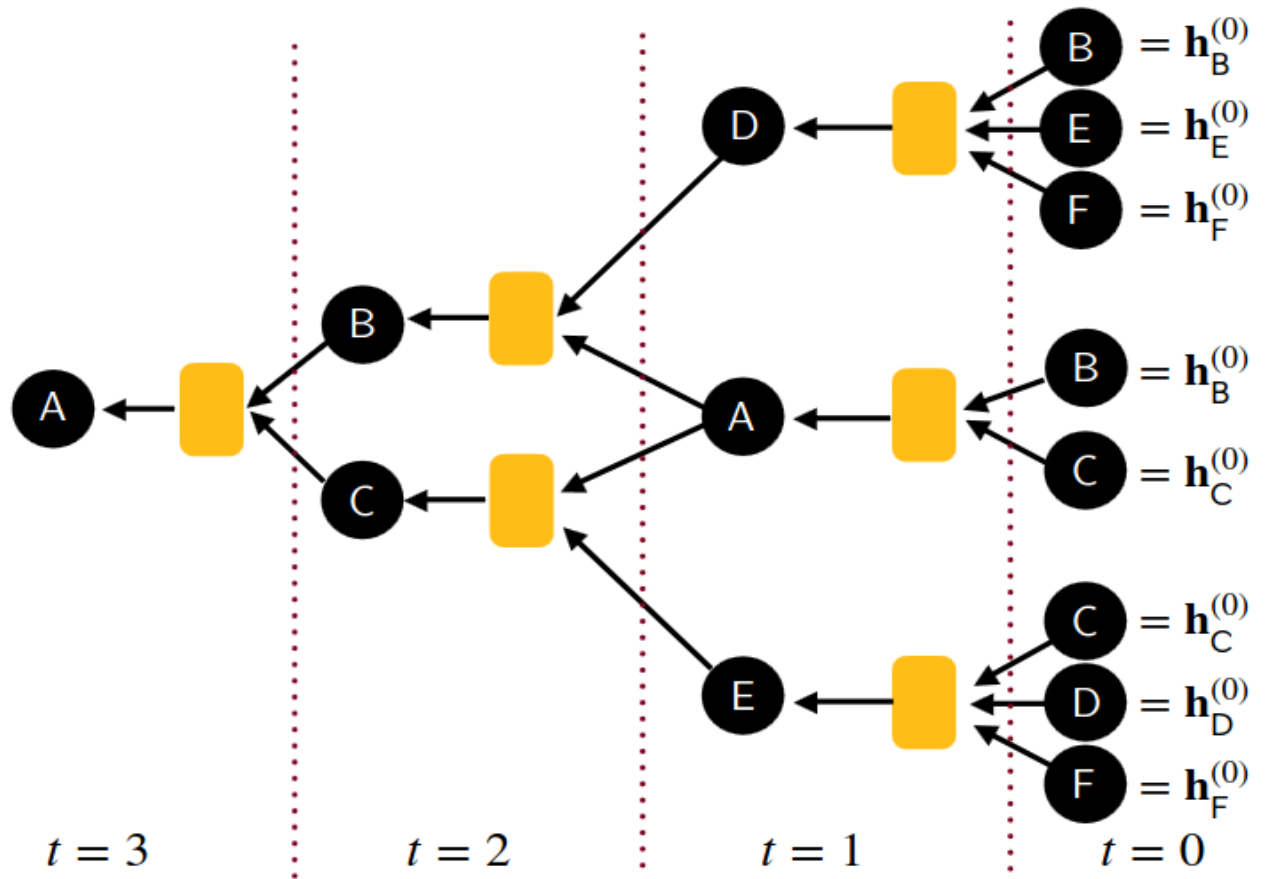

A graph (left) and an illustration of message passing on this graph with respect to the target node A for 3 iterations (right). Directed arrows depict the **messages**, and yellow boxes denote **aggregation**. At least 3 iterations are needed to get information from **all** nodes, i.e., F will not pass any messages to A with .




A graph (left) and an illustration of message passing on this graph with respect to the target node A for 3 iterations (right). Directed arrows depict the **messages**, and yellow boxes denote **aggregation**. At least 3 iterations are needed to get information from **all** nodes, i.e., F will not pass any messages to A with .



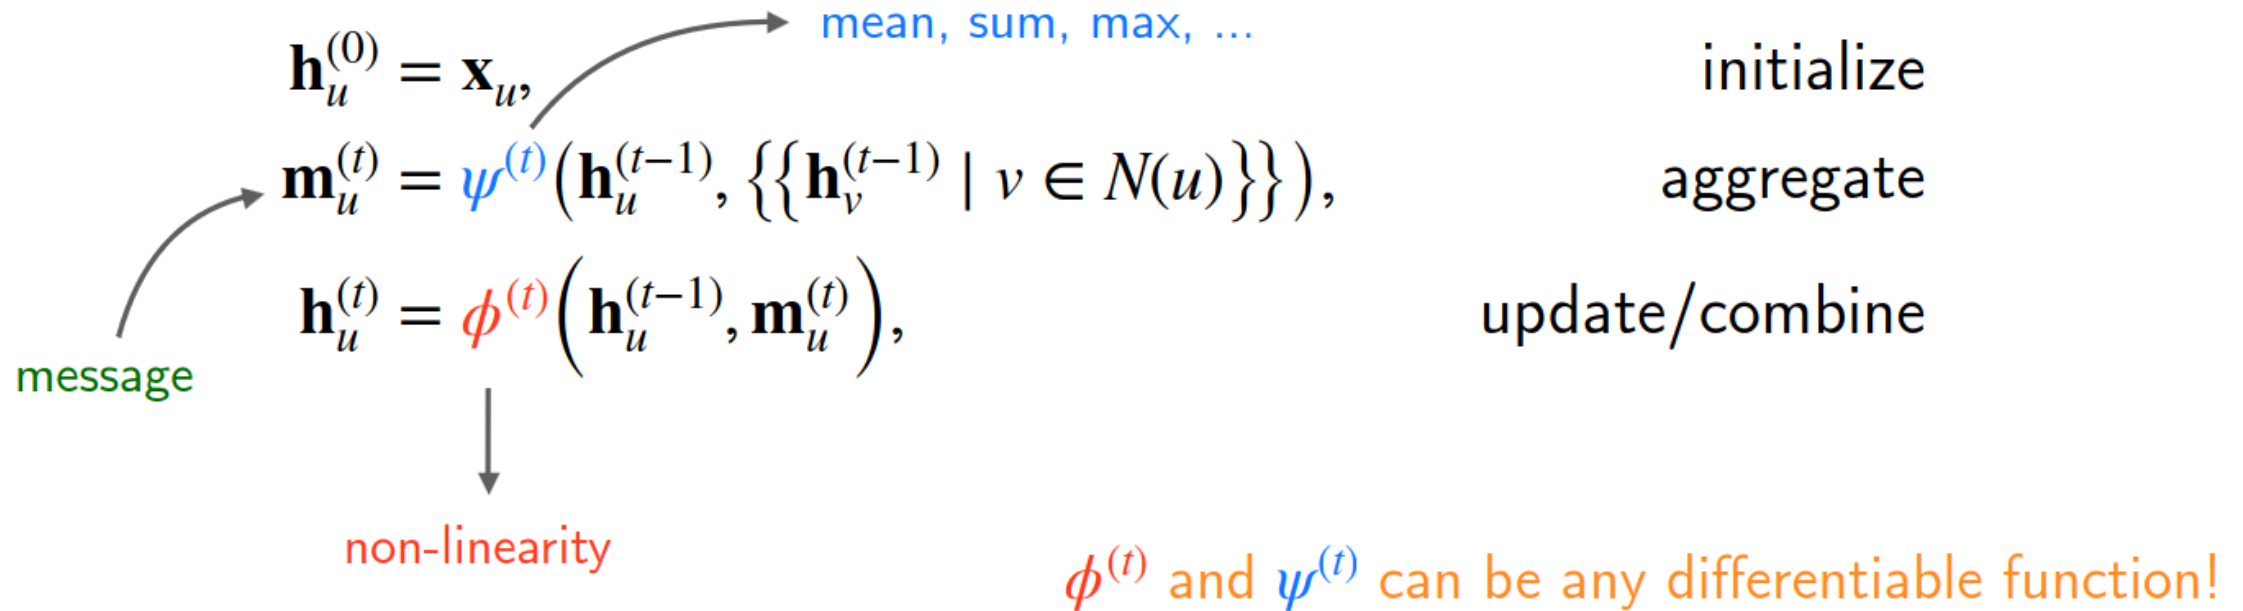
A graph (left) and an illustration of message passing on this graph with respect to the target node A for 3 iterations (right). Directed arrows depict the **messages**, and yellow boxes denote **aggregation**. At least 3 iterations are needed to get information from **all** nodes, i.e., F will not pass any messages to A with .




A graph (left) and an illustration of message passing on this graph with respect to the target node A for 3 iterations (right). Directed arrows depict the **messages**, and yellow boxes denote **aggregation**. At least 3 iterations are needed to get information from **all** nodes, i.e., F will not pass any messages to A with .



Given a graph  $G = (V, E, \mathbf{X})$ , an MPNN iteratively computes  $\mathbf{h}_u^{(t)}$  for every node  $u \in V$ :



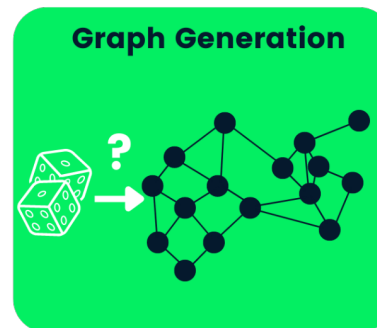




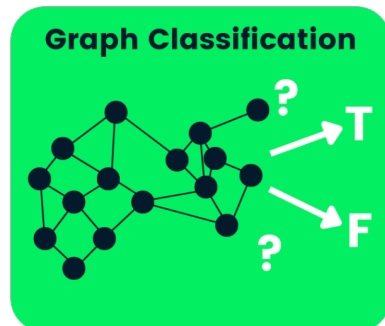


You may encounter variations, where a message computation function `msg` is defined w.r.t the source node:

$$\mathbf{h}_u^{(t)} = \phi^{(t)} \left( \mathbf{h}_u^{(t-1)}, \psi^{(t)} \left( \left\{ \left\{ \text{msg}(\mathbf{h}_u^{(t-1)}, \mathbf{h}_v^{(t-1)}) \mid v \in N(u) \right\} \right\} \right) \right),$$


**Remark:** The function `msg` typically depends on the neighborhood - hard to decouple `msg` from  $\psi^{(t)}$ . Following a common convention, we view the message computation as part of aggregation.

# GNN Tasks



**Input:** a graph where a subset of the nodes are labelled with a class.

**Task:** predict the labels of the remaining nodes, i.e., test nodes in the graph





### Example:

- Citeseer is a citation network, where nodes represent **papers**, and edges denote **citation links**, and a subset of the nodes are labelled with a **paper category** (e.g., AI, ML).
- Predict the category (or, categories) of the remaining papers.

**Input:** a graph

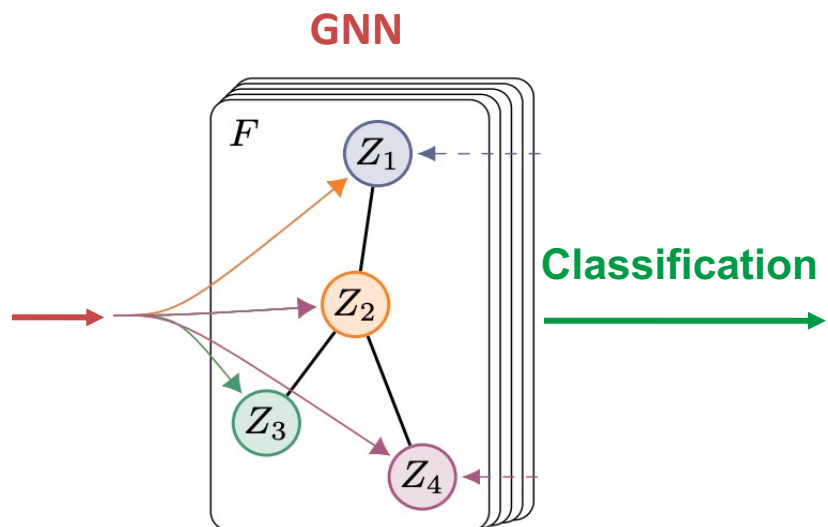

**Task:** predict new links





### Example:

- **OGBL-DDI** is a **drug-drug interaction network**: Nodes represent drugs and edges interactions between drugs.
- Predict drug-drug interactions: rank true drug interactions higher than non-interacting drug pairs, using, e.g., a pairwise decoder.


## homopolymers



Classification

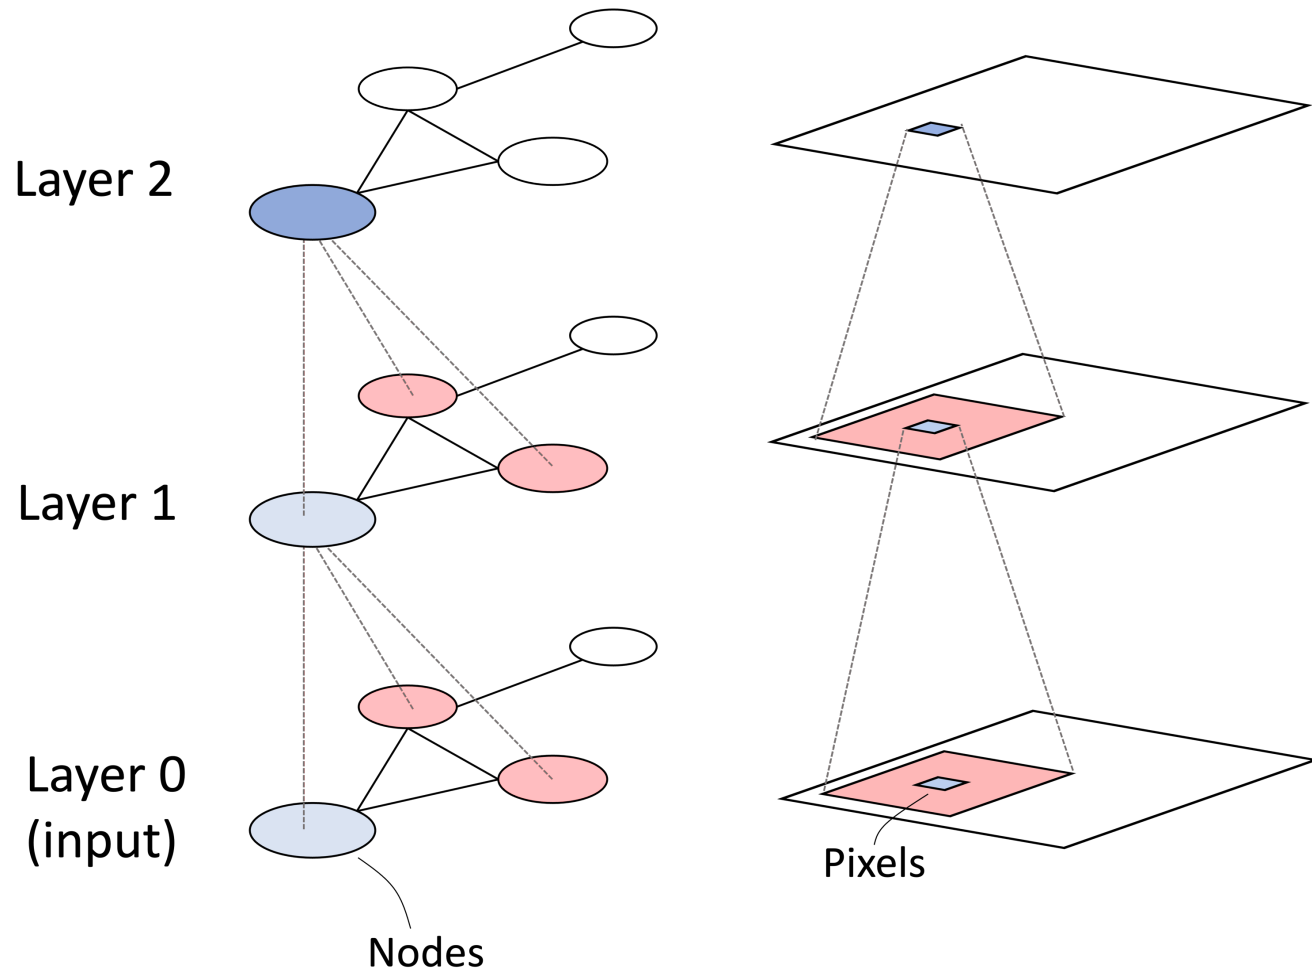



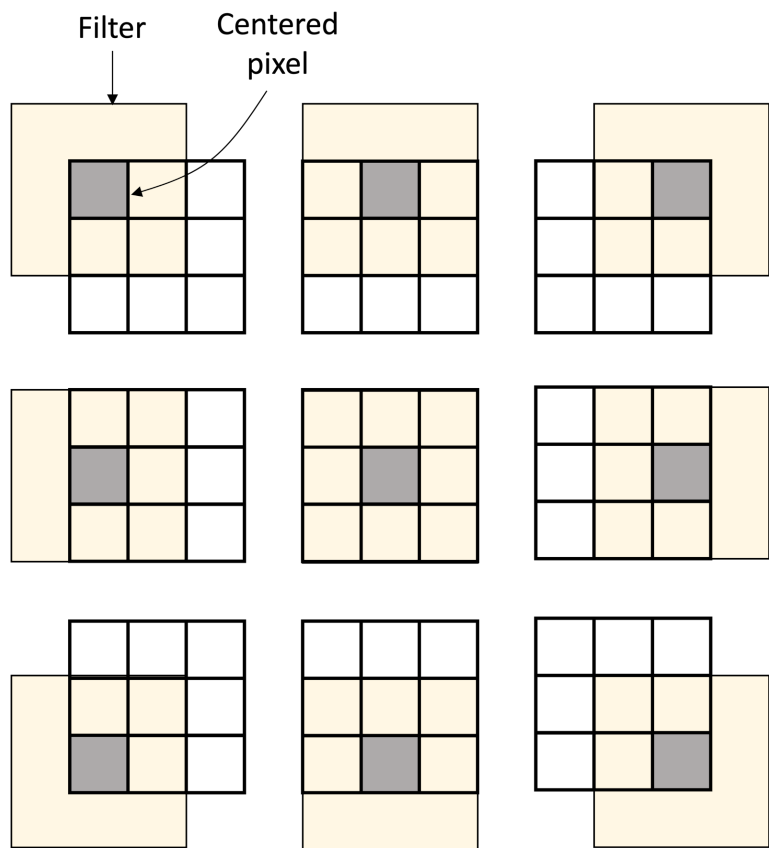
# Graph Convolutional Network

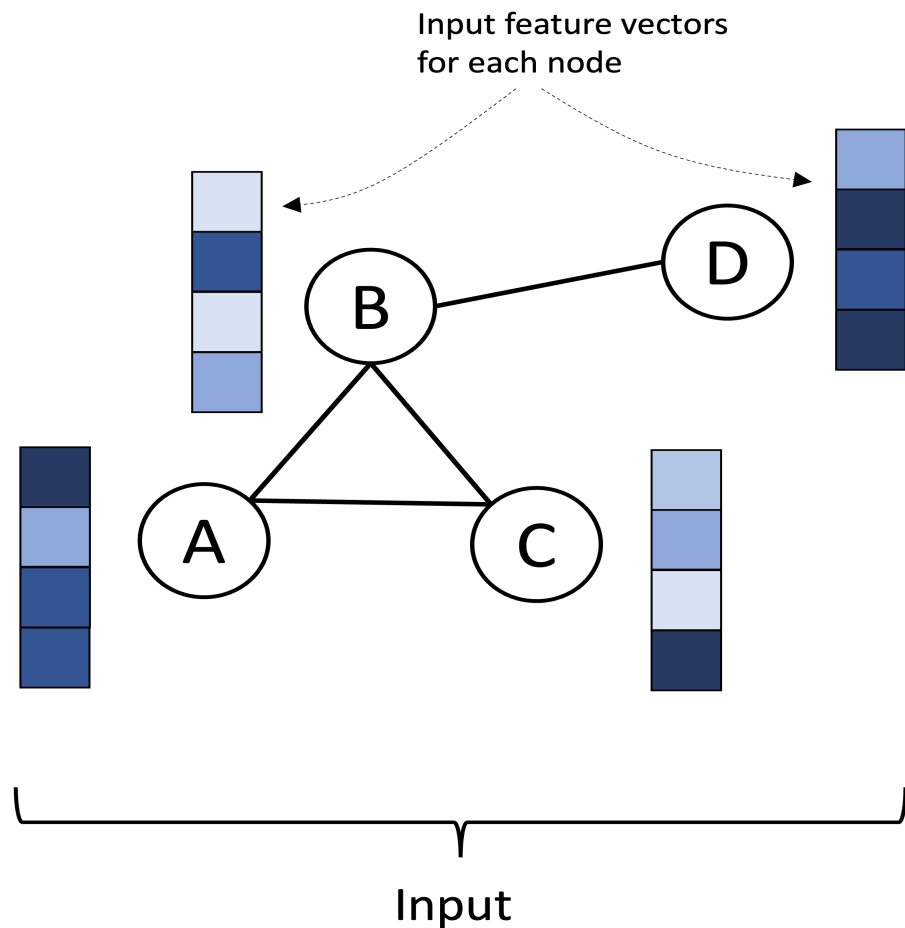


### GCN (Graph)



### CNN (Image)




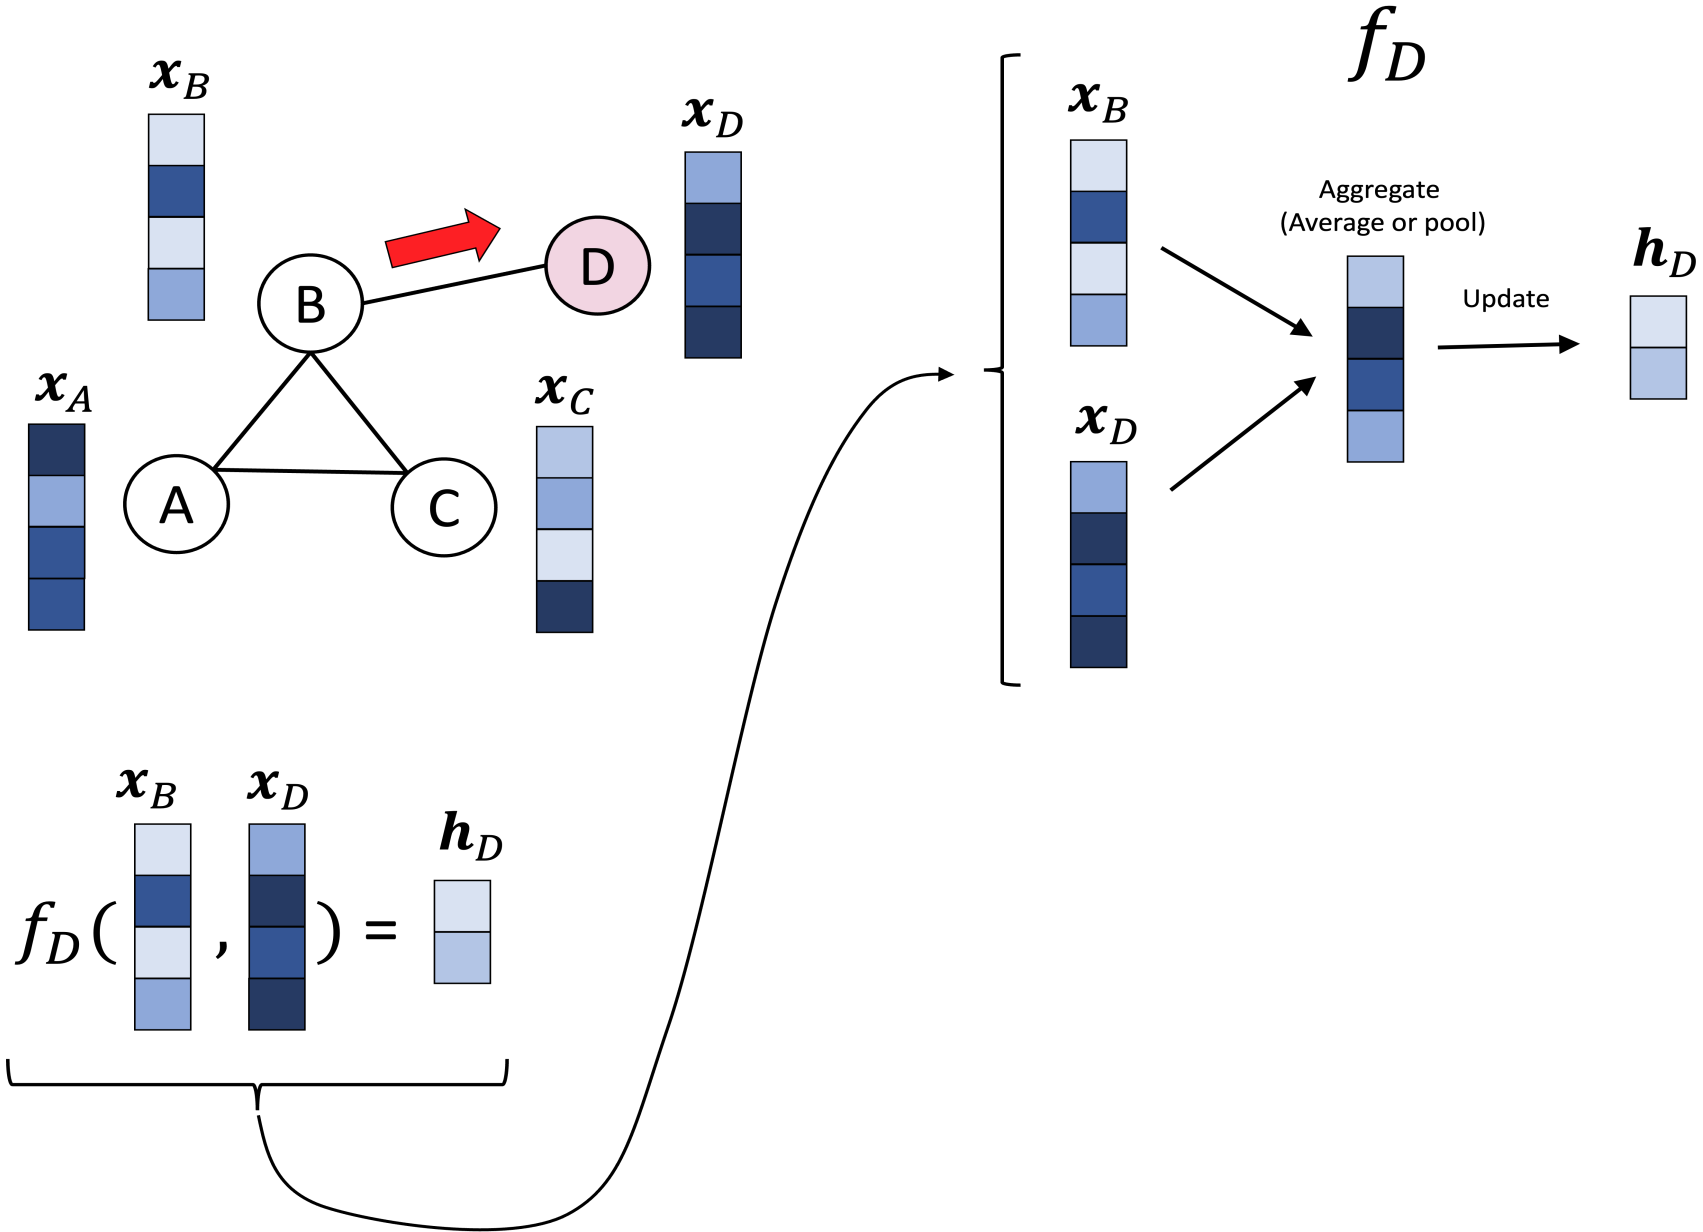




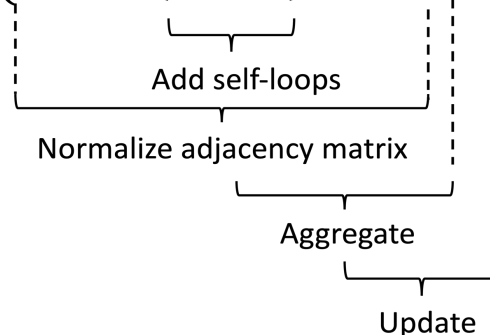
Graph convolutional layers







$$f(\mathbf{X}, \mathbf{A}) := \sigma(\mathbf{D}^{-1/2}(\mathbf{A} + \mathbf{I})\mathbf{D}^{-1/2}\mathbf{X}\mathbf{W})$$



$\mathbf{A} \in \mathbb{R}^{n \times n}$  := The adjacency matrix

$\mathbf{I} \in \mathbb{R}^{n \times n}$  := The identity matrix


$\mathbf{D} \in \mathbb{R}^{n \times n}$  := The degree matrix of  $\mathbf{A} + \mathbf{I}$

$\mathbf{X} \in \mathbb{R}^{n \times d}$  := The input data (i.e., the per-node feature vectors)

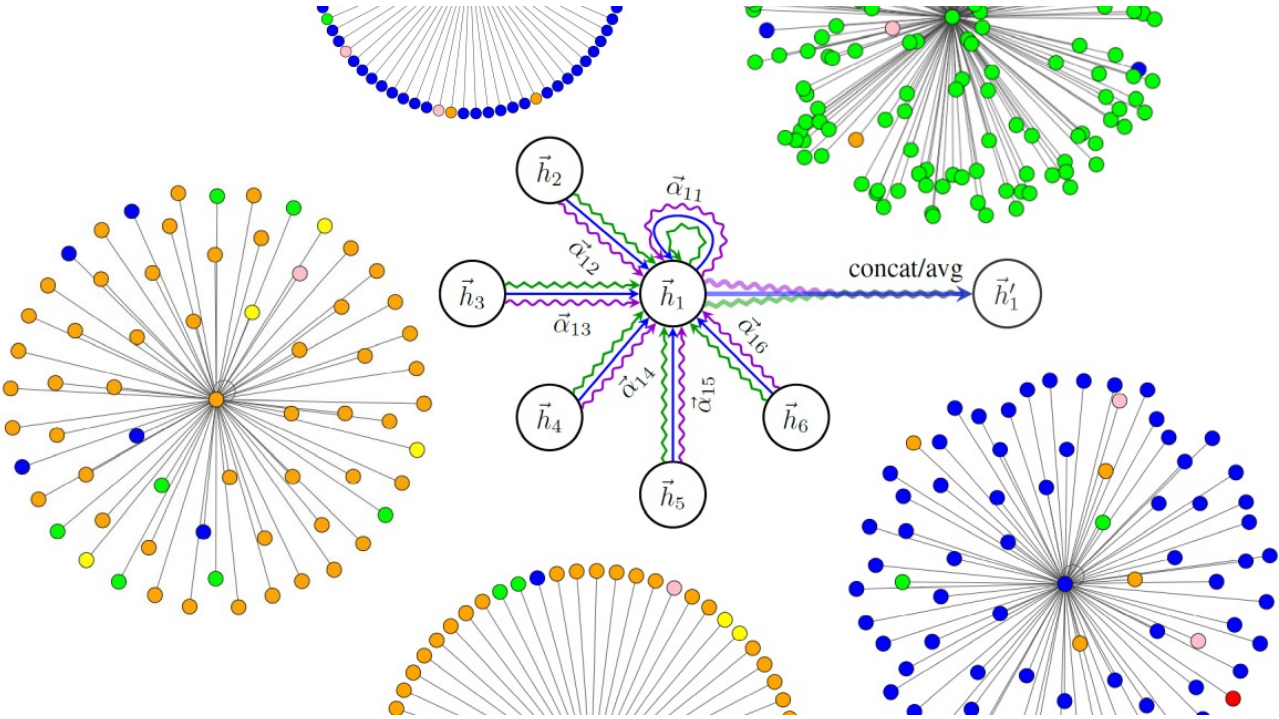
$\mathbf{W} \in \mathbb{R}^{d \times w}$  := The layer's weights


$\sigma(\cdot)$  := The activation function (e.g., ReLU)

$$\mathbf{D} := \begin{bmatrix} d_{1,1} & 0 & 0 & \dots & 0 \\ 0 & d_{2,2} & 0 & \dots & 0 \\ 0 & 0 & d_{3,3} & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & d_{n,n} \end{bmatrix}$$



# Graph Attention Network





$$h'_i = \sigma \left( \sum_{j \in \mathcal{N}(i)} W h_j \right)$$

1	1	1	0	0
1	1	1	1	1
1	1	1	0	0
0	1	0	1	1
0	1	0	1	1

Adj Matrix


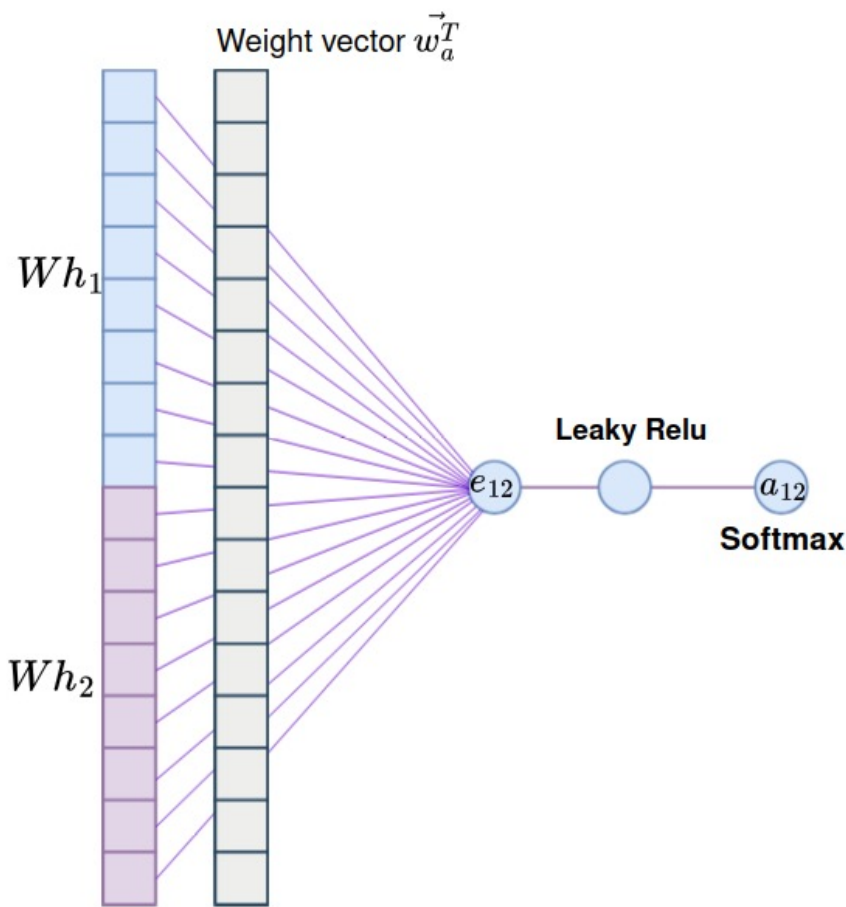
$h_1$							
$h_2$							
$h_3$							
$h_4$							
$h_5$							

Features per node


Learnable weight Matrix

$h'_1$							
$h'_2$							
$h'_3$							
$h'_4$							
$h'_5$							

Embeddings per node

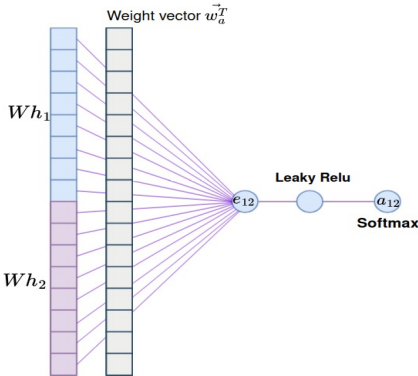


$$\alpha_{ij} = \frac{\exp(\text{LeakyReLU}(\vec{w}_a^T [Wh_i \parallel Wh_j]))}{\sum_{k \in \mathcal{N}(i)} \exp(\text{LeakyReLU}(\vec{w}_a^T [Wh_i \parallel Wh_k]))}$$


$$h'_i = \sigma \left( \sum_{j \in N(i)} a_{ij} W h_j \right)$$

1	1	1	0	0
1	1	1	1	1
1	1	1	0	0
0	1	0	1	1
0	1	0	1	1


Adj Matrix




Attention



Features per node

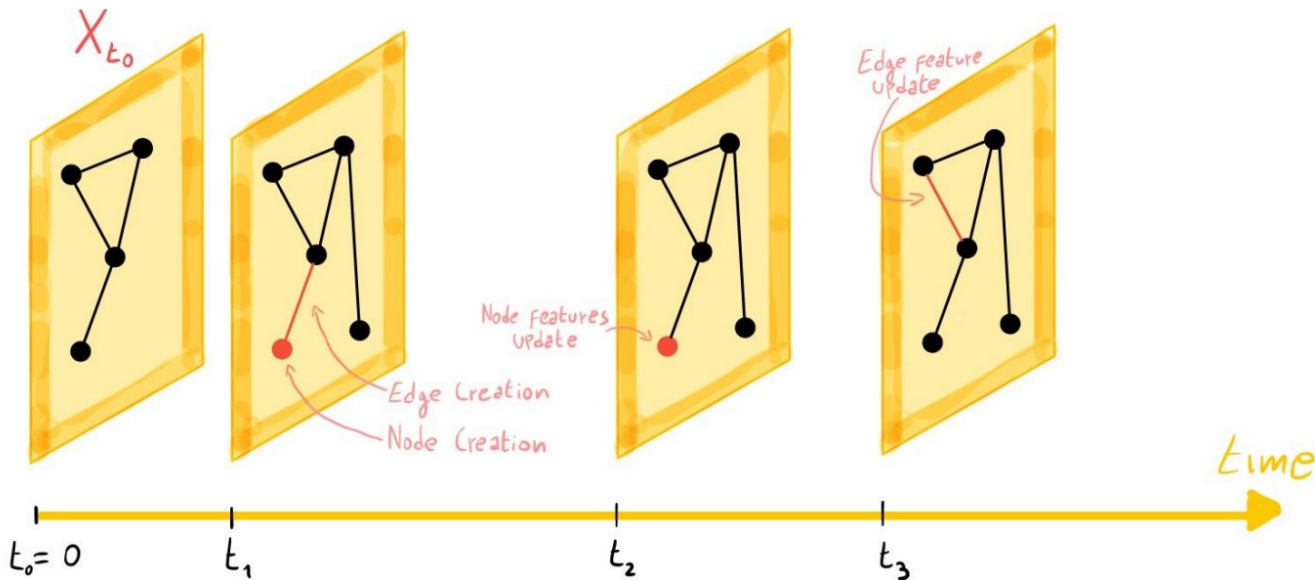



Learnable weight Matrix




Embeddings per node

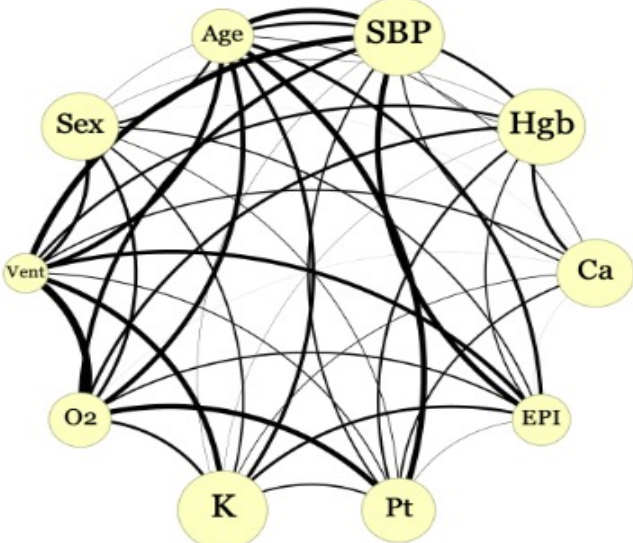
# Dynamic Graph Neural Network



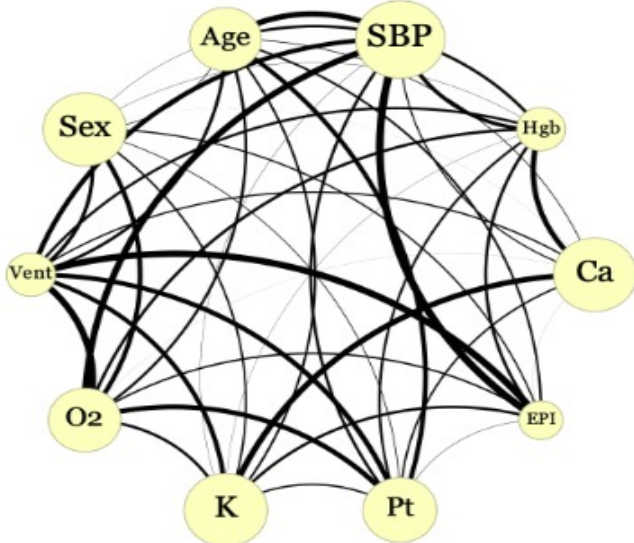




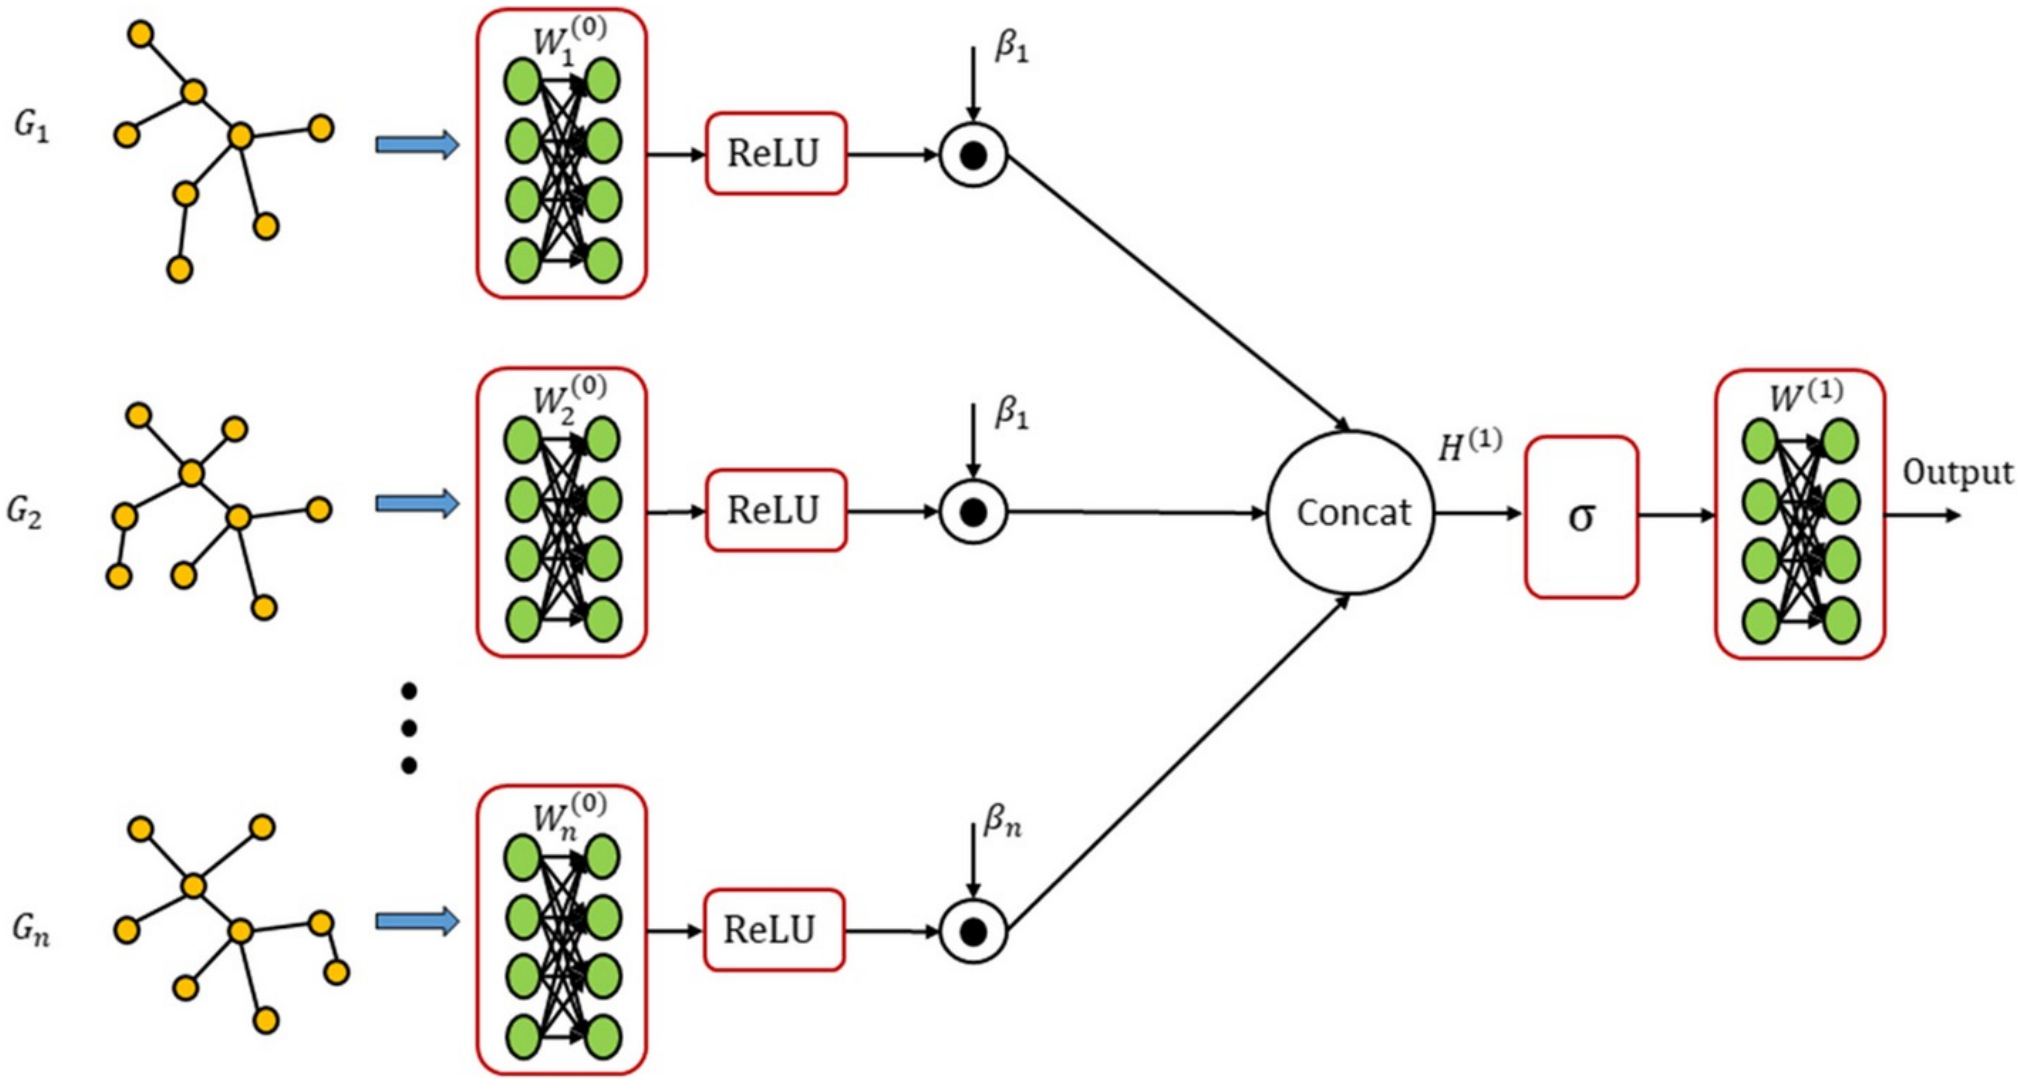
(a) t=1

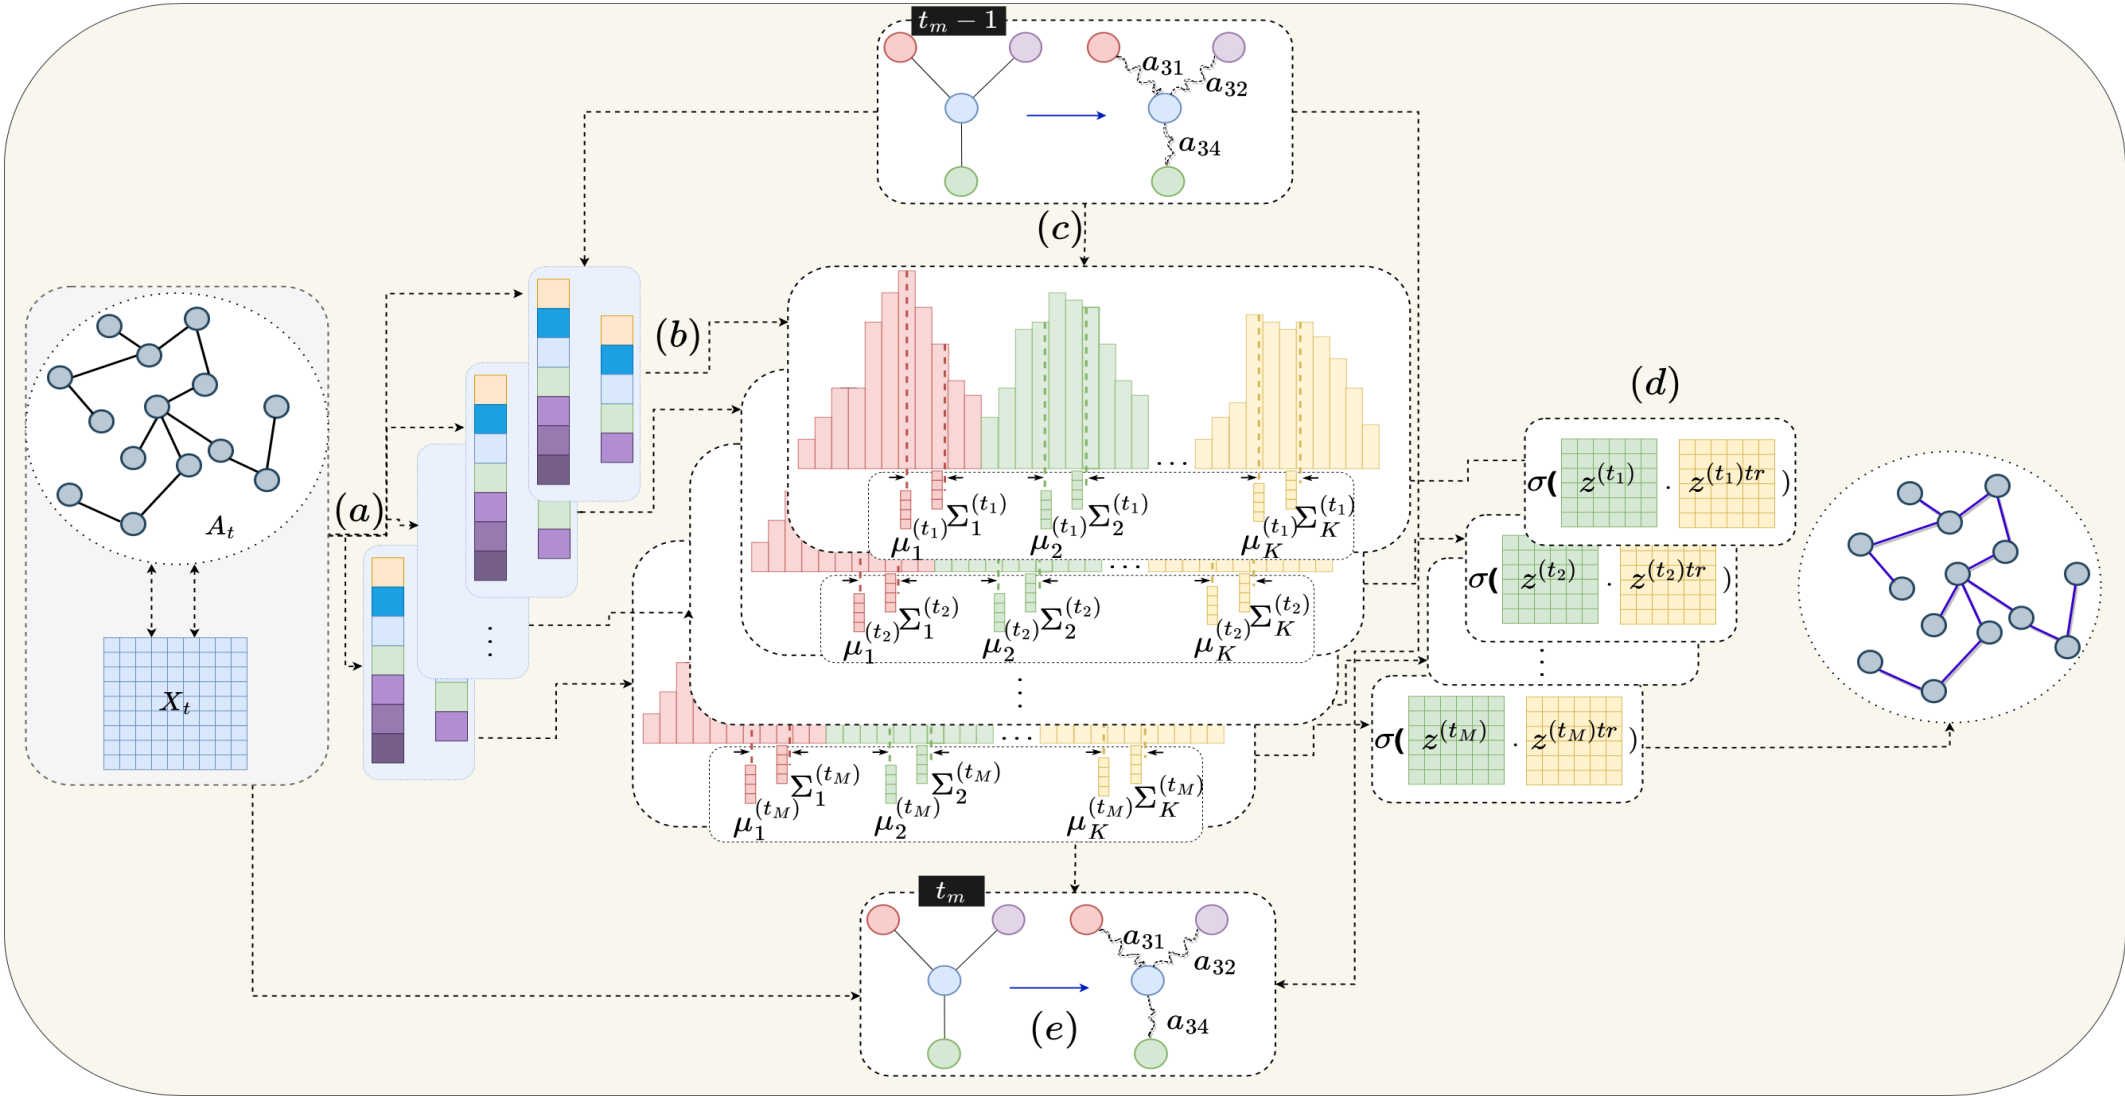


(b) t=4



(c) t=6

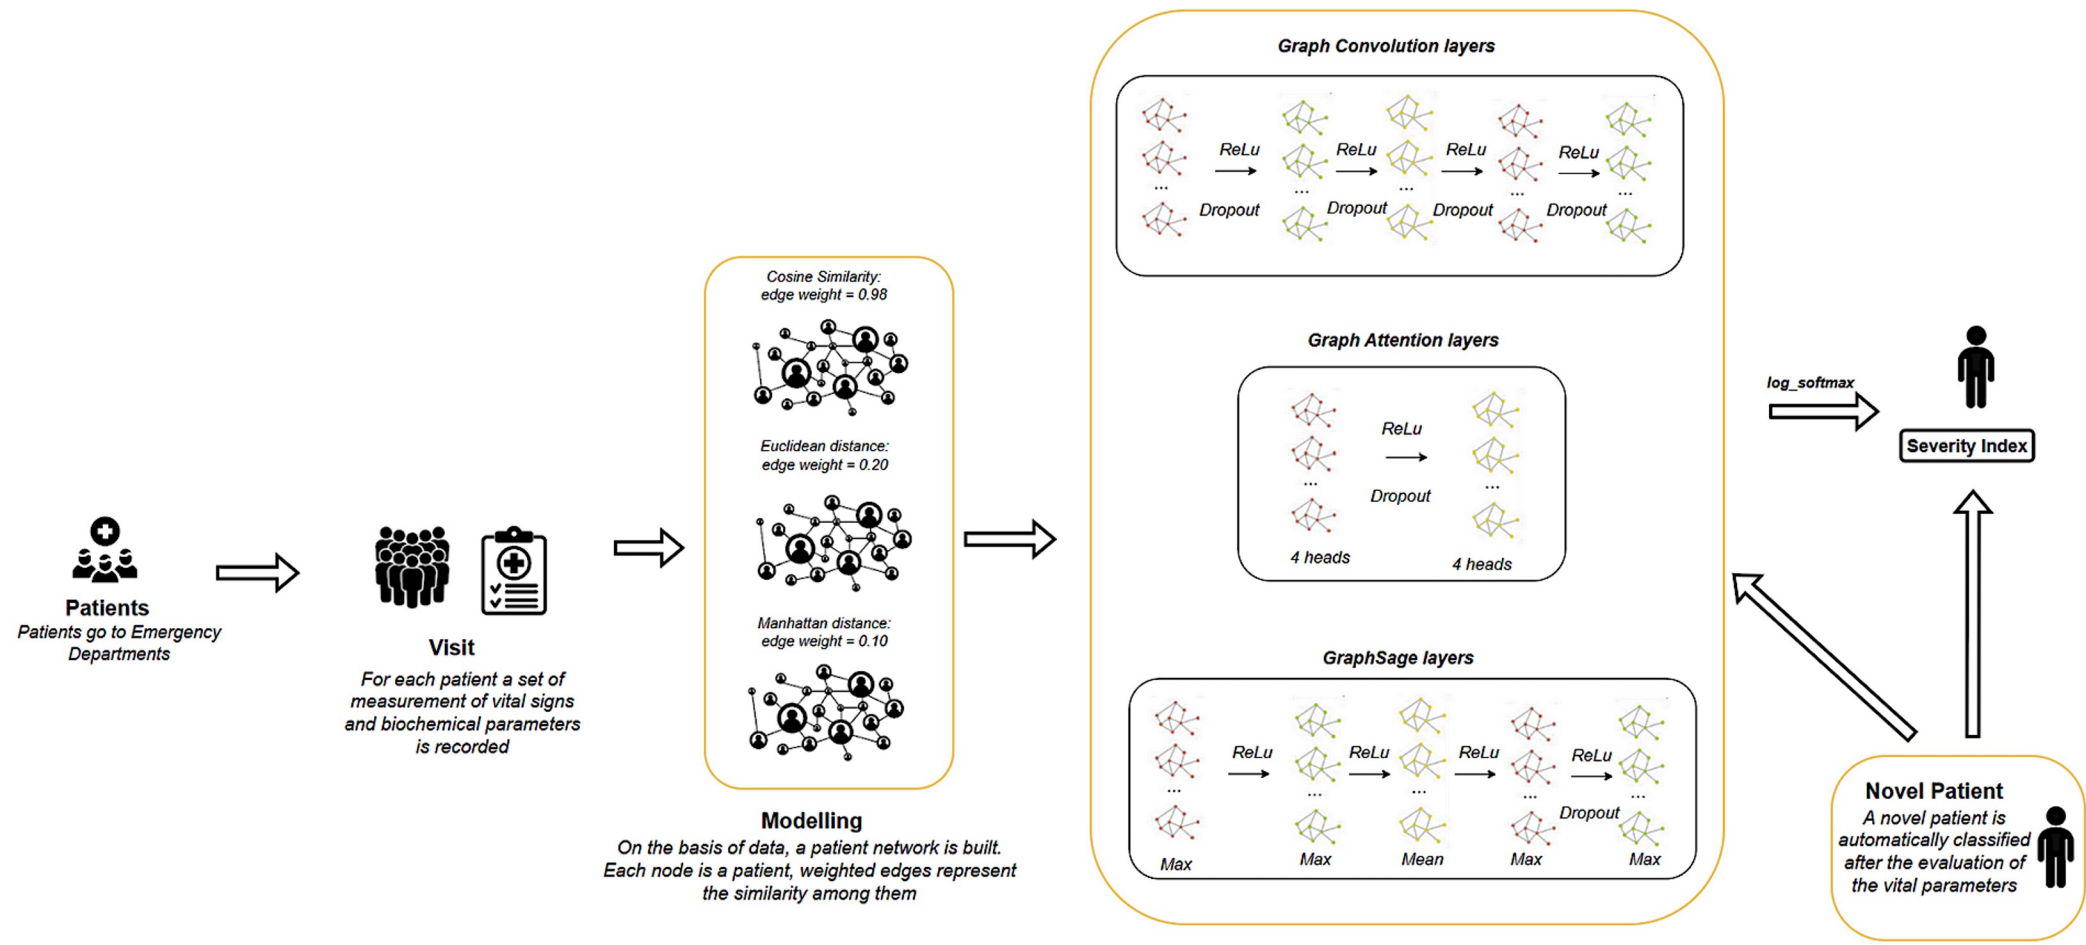





- Handling **Small Datasets**
- Robustness to **Missing Information**
- Effective with **Sparse** Data
- Adaptability to **Non-Euclidean** Data
- Capturing **Complex Interactions**
- Suitability for **Dynamic** and **Heterogeneous** Data


# Applications in Healthcare









(a) Brain MRI




(b) GNN modeling of the brain




fMRI



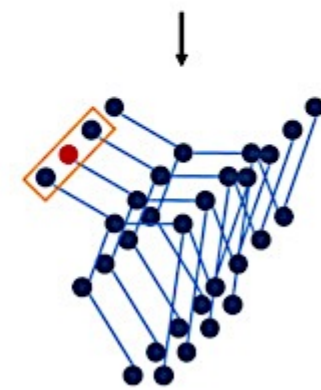
BOLD signals




Graph construction




Spatial convolution




Temporal convolution



Node projection



Graph pooling



Graph classification

