Graph Neural Networks

A Gentle Introduction

April 2022

Main Sources

A Gentle Introduction to Graph Neural Networks
Benjamin Sanchez-Lengeling, Emily Reif, Adam Pearce, Alexander B. Wiltschko
Google Research, https://doi.org/10.23915/distill.00033

Graph Convolutional Neural Networks (GCNs) Made Simple
WelcomeAIOverloards
https://youtu.be/2KRAOZIULzw

Agenda

1. What kind of data is naturally phrased as a graph?
2. What makes graphs different from other types of data?
3. Building a GNN
4. Try it out! Build intuition and work with a real-world task

Graphs

They’re everywhere!

- Real world objects are often defined in terms of their connections to other things
- A set of objects, and the connections between them, are expressed as a graph
- Neural network that operates on graph data = graph neural network (GNN)

Graphs

- A graph represents the relations ***(edges)*** between a collection of entities ***(nodes)***

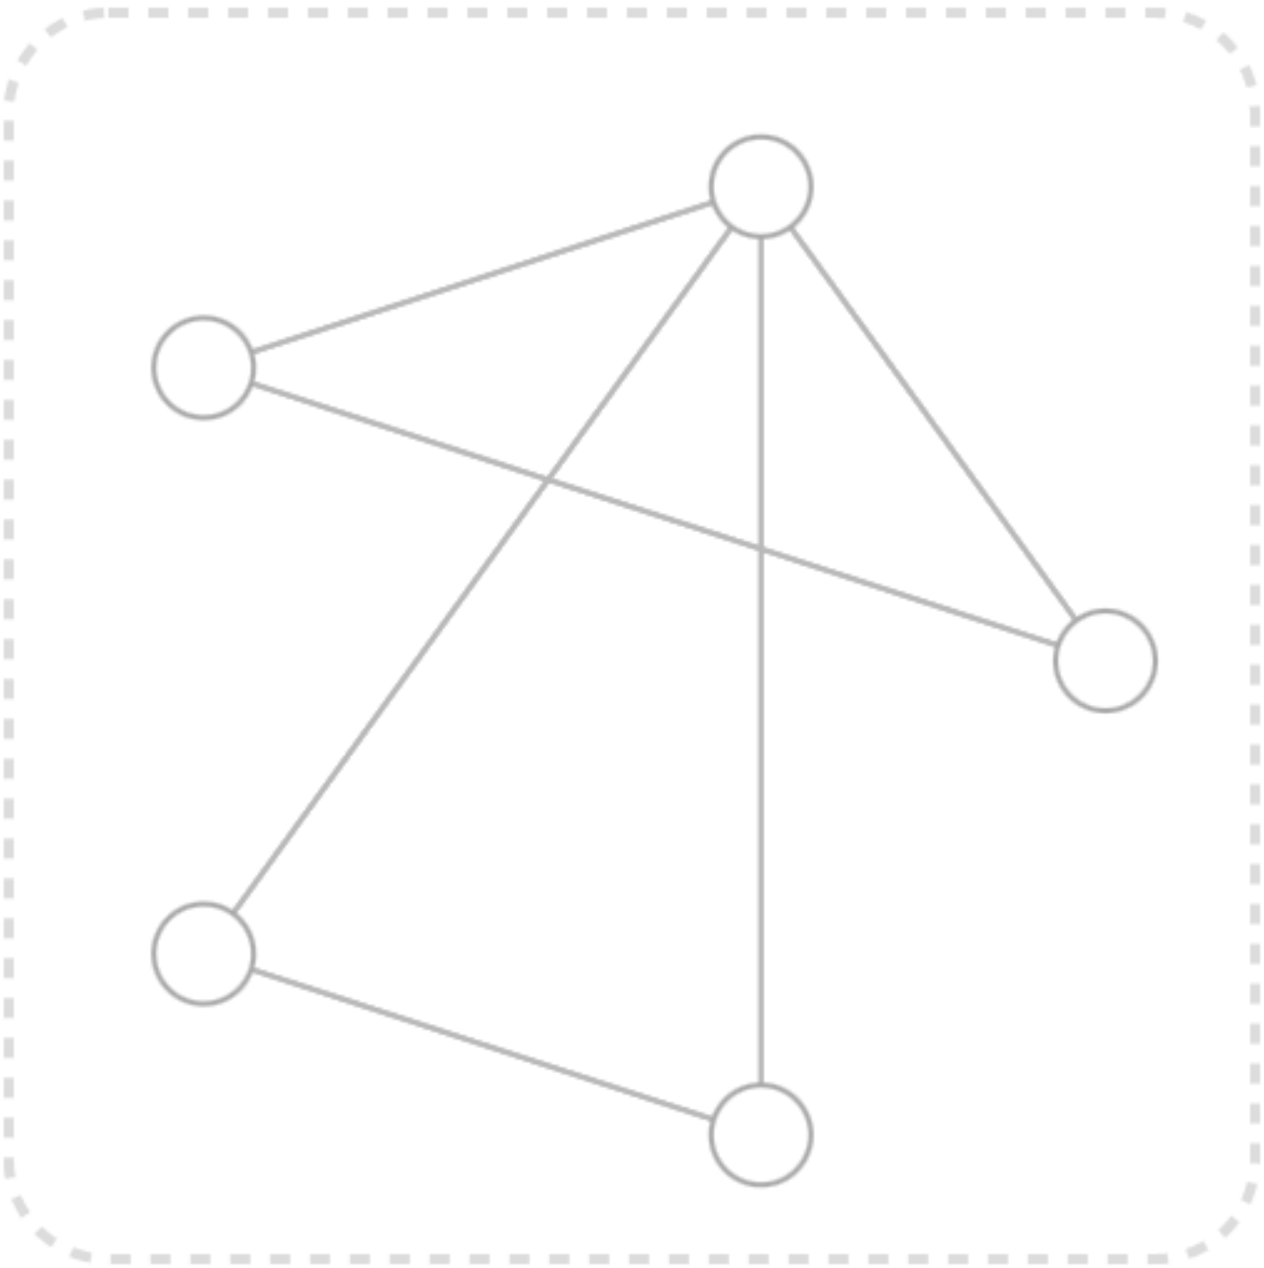

V Vertex (or node) attributes

e.g., node identity, number of neighbors

E Edge (or link) attributes and directions

e.g., edge identity, edge weight

U Global (or master node) attributes

e.g., number of nodes, longest path

Graphs

- We can associate directionality to edges

- An edge has a source node v_{src} and a destination node v_{dst}
- Note: having a single undirected edge is the same as having one directed edge from v_{src} to v_{dst}, and another directed edge from v_{dst} to v_{src}

Images as Graphs

- Visualize the connectivity of a graph through its adjacency matrix

- Order the nodes and fill a matrix of $(n_{nodes} \times n_{nodes})$ with an entry if two nodes share an edge

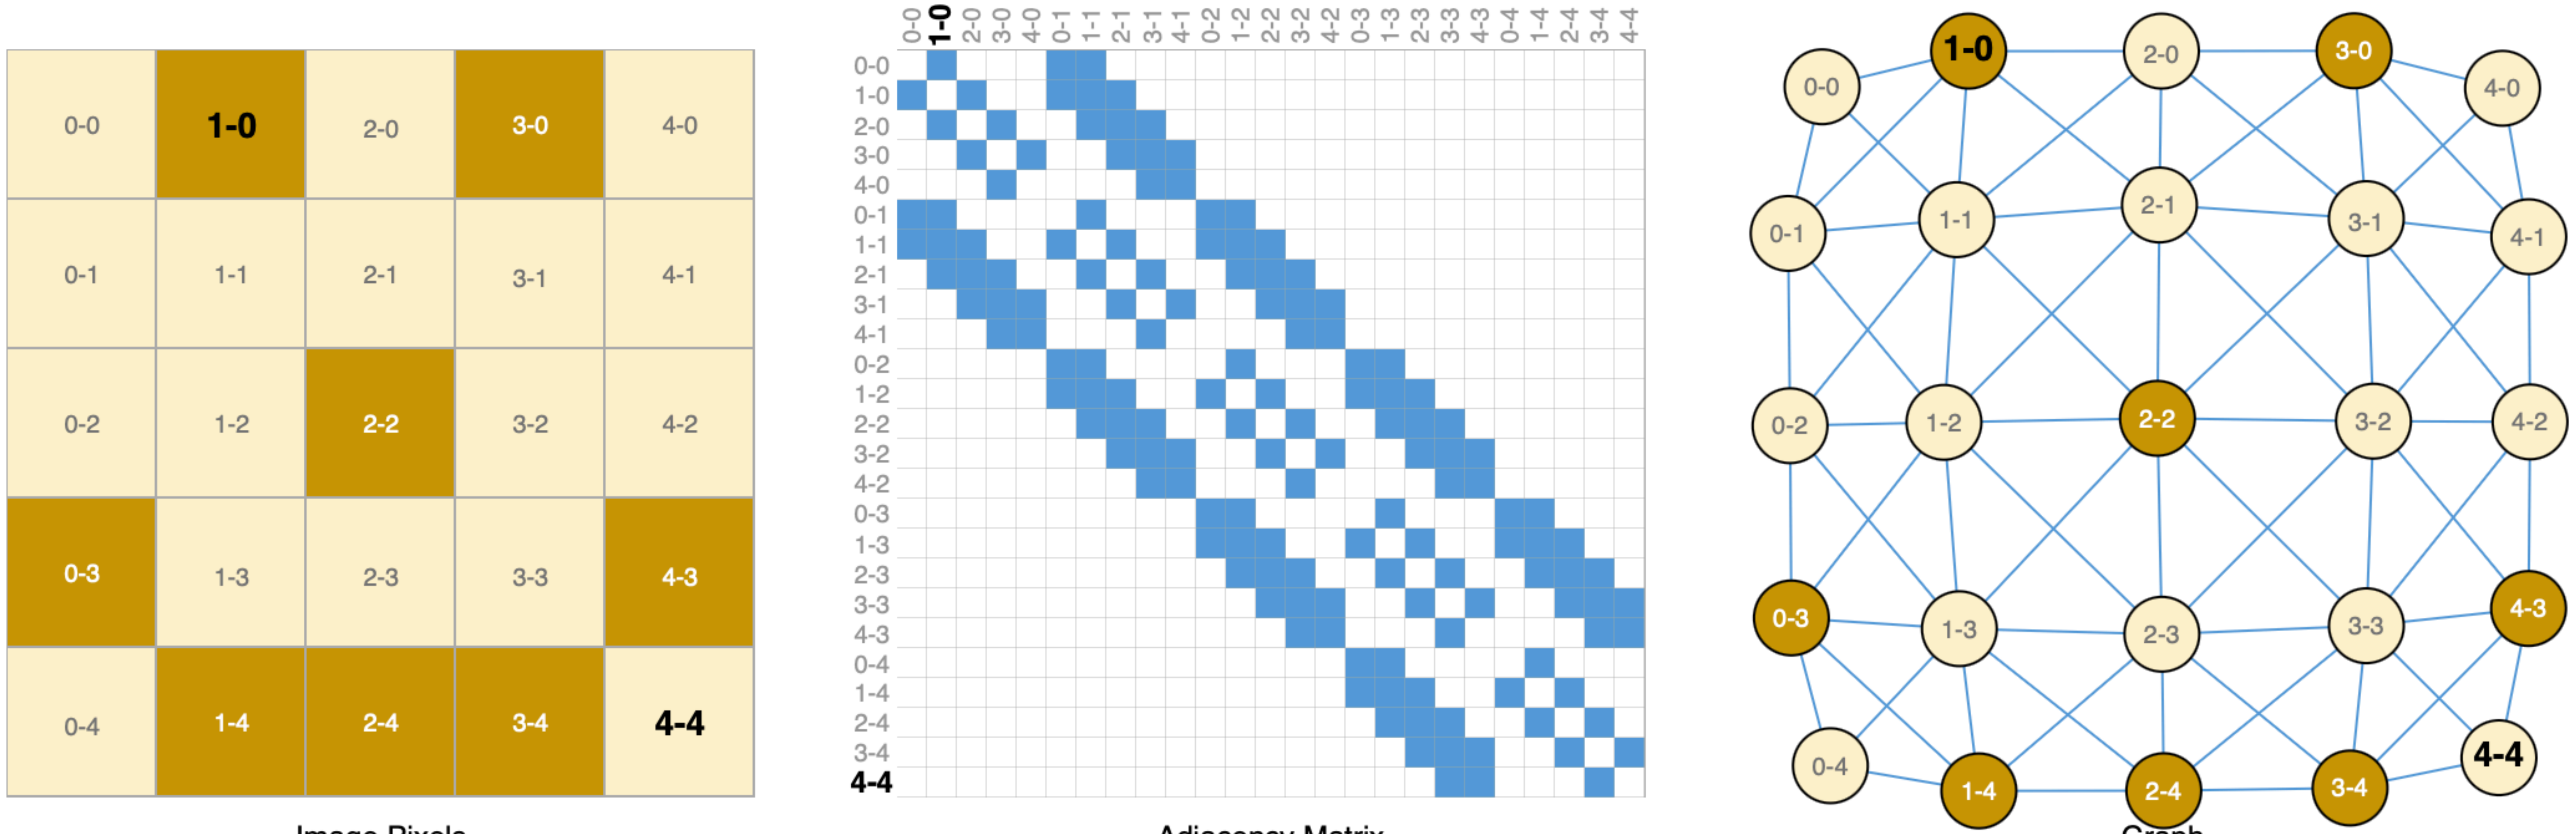

Image Pixels — Adjacency Matrix — Graph

Text as Graphs

- Digitize text by associating indices to each character, word, or token, and representing text as a sequence of these indices
- These graph representations are somewhat redundant (images and text have very regular structure)

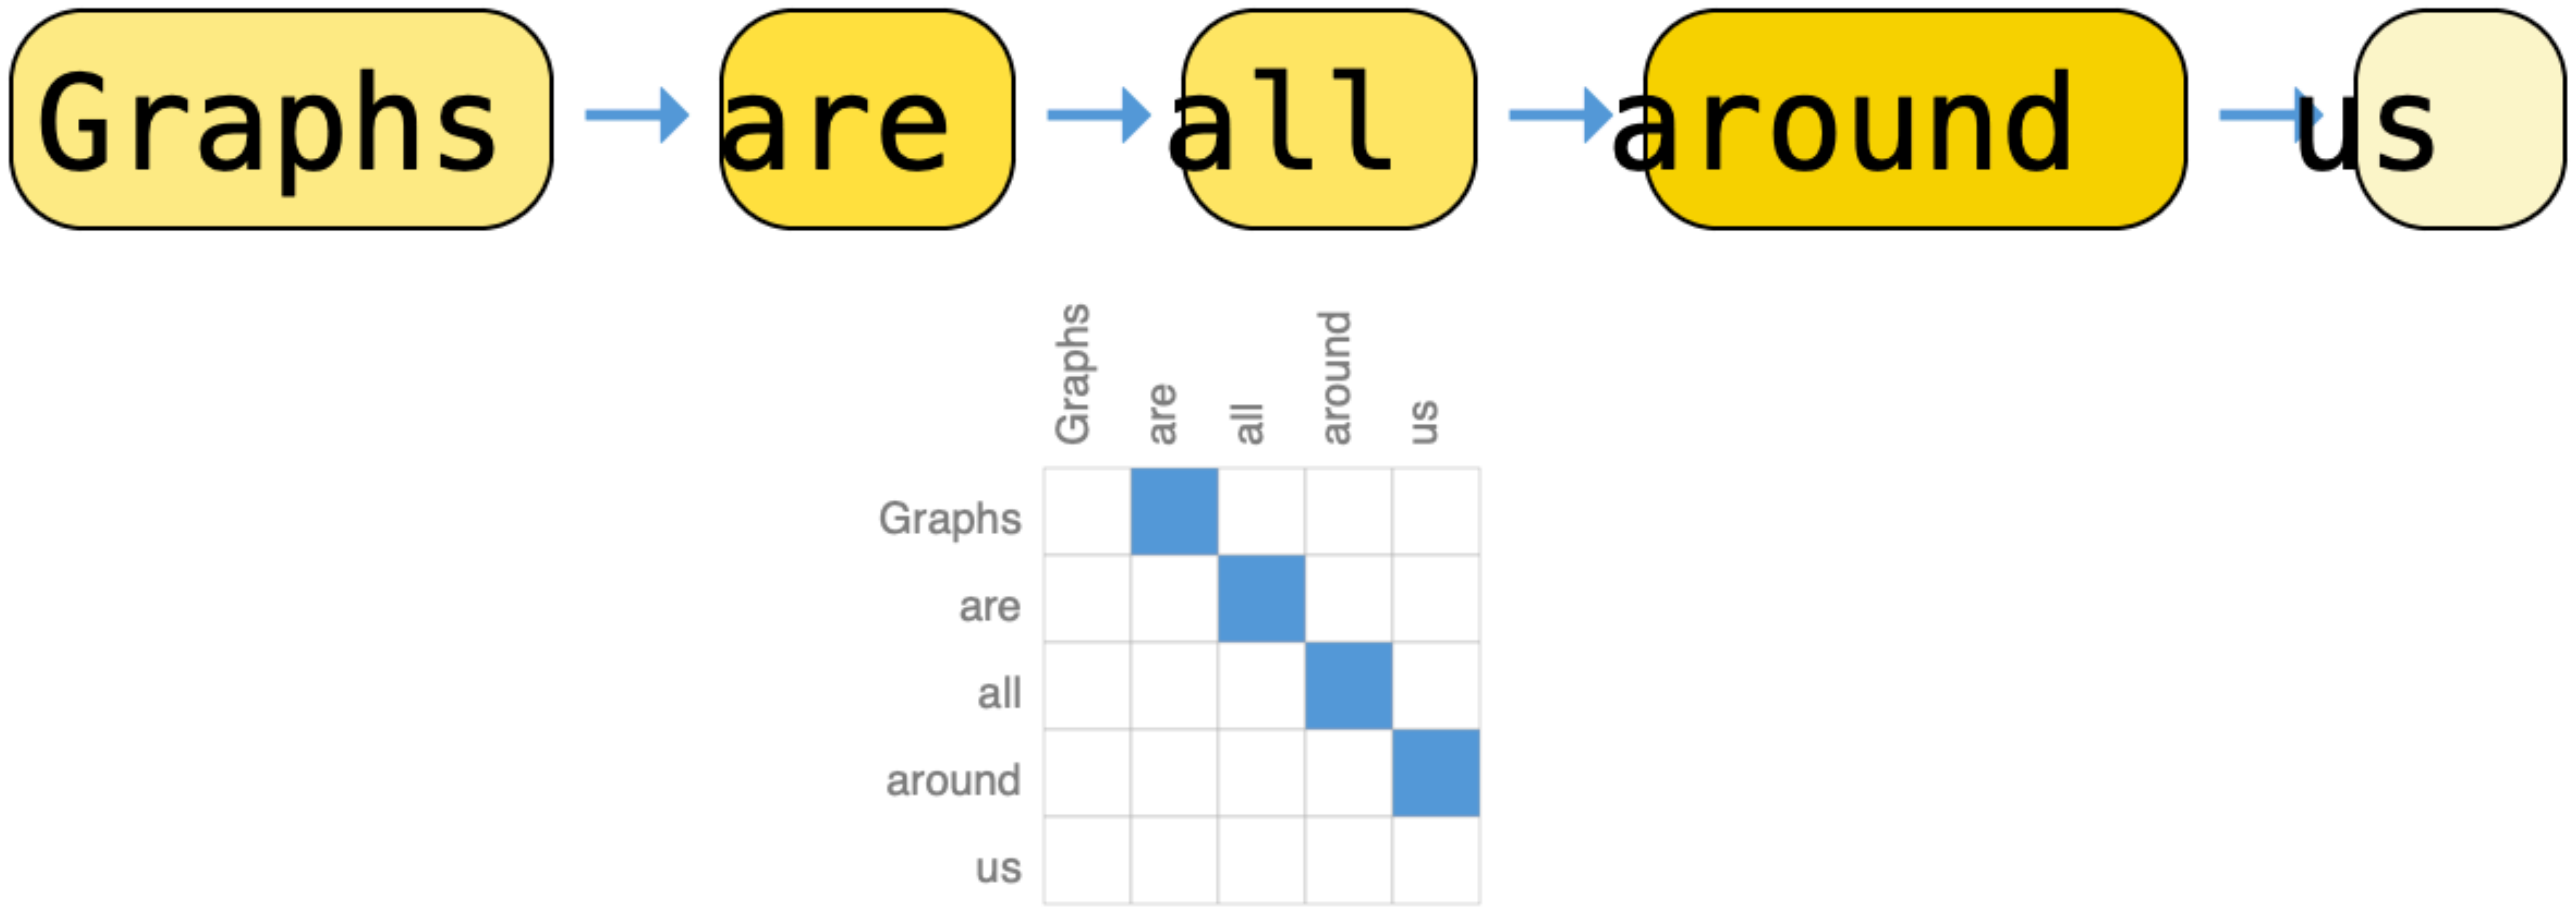

Graph-Valued Data

Molecules as graphs

- Nodes are atoms and edges are covalent bonds

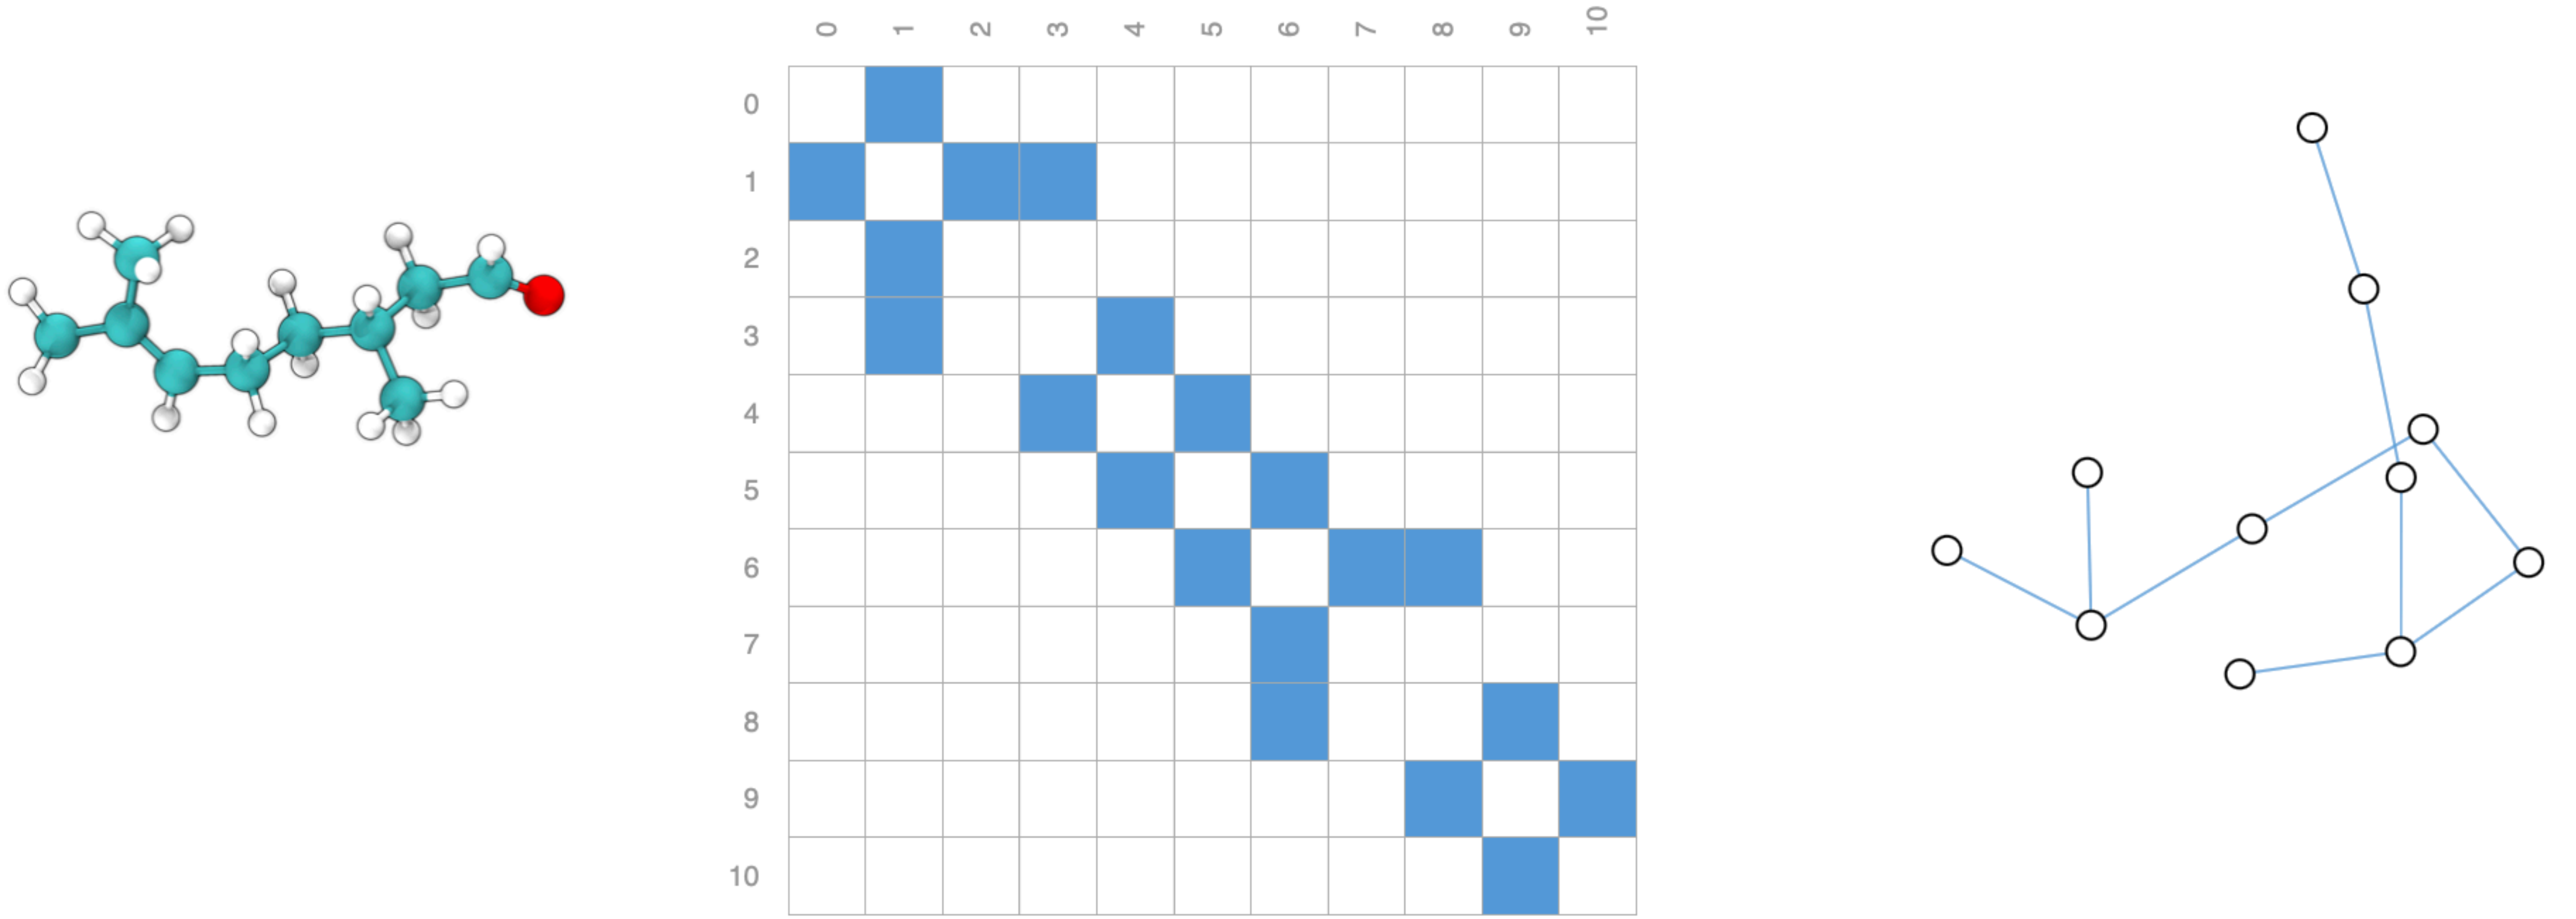

(Left) 3d representation of the Citronellal molecule (Center) Adjacency matrix of the bonds in the molecule (Right) Graph representation of the molecule

Graph-Valued Data

Molecules as graphs

- Nodes are atoms and edges are covalent bonds

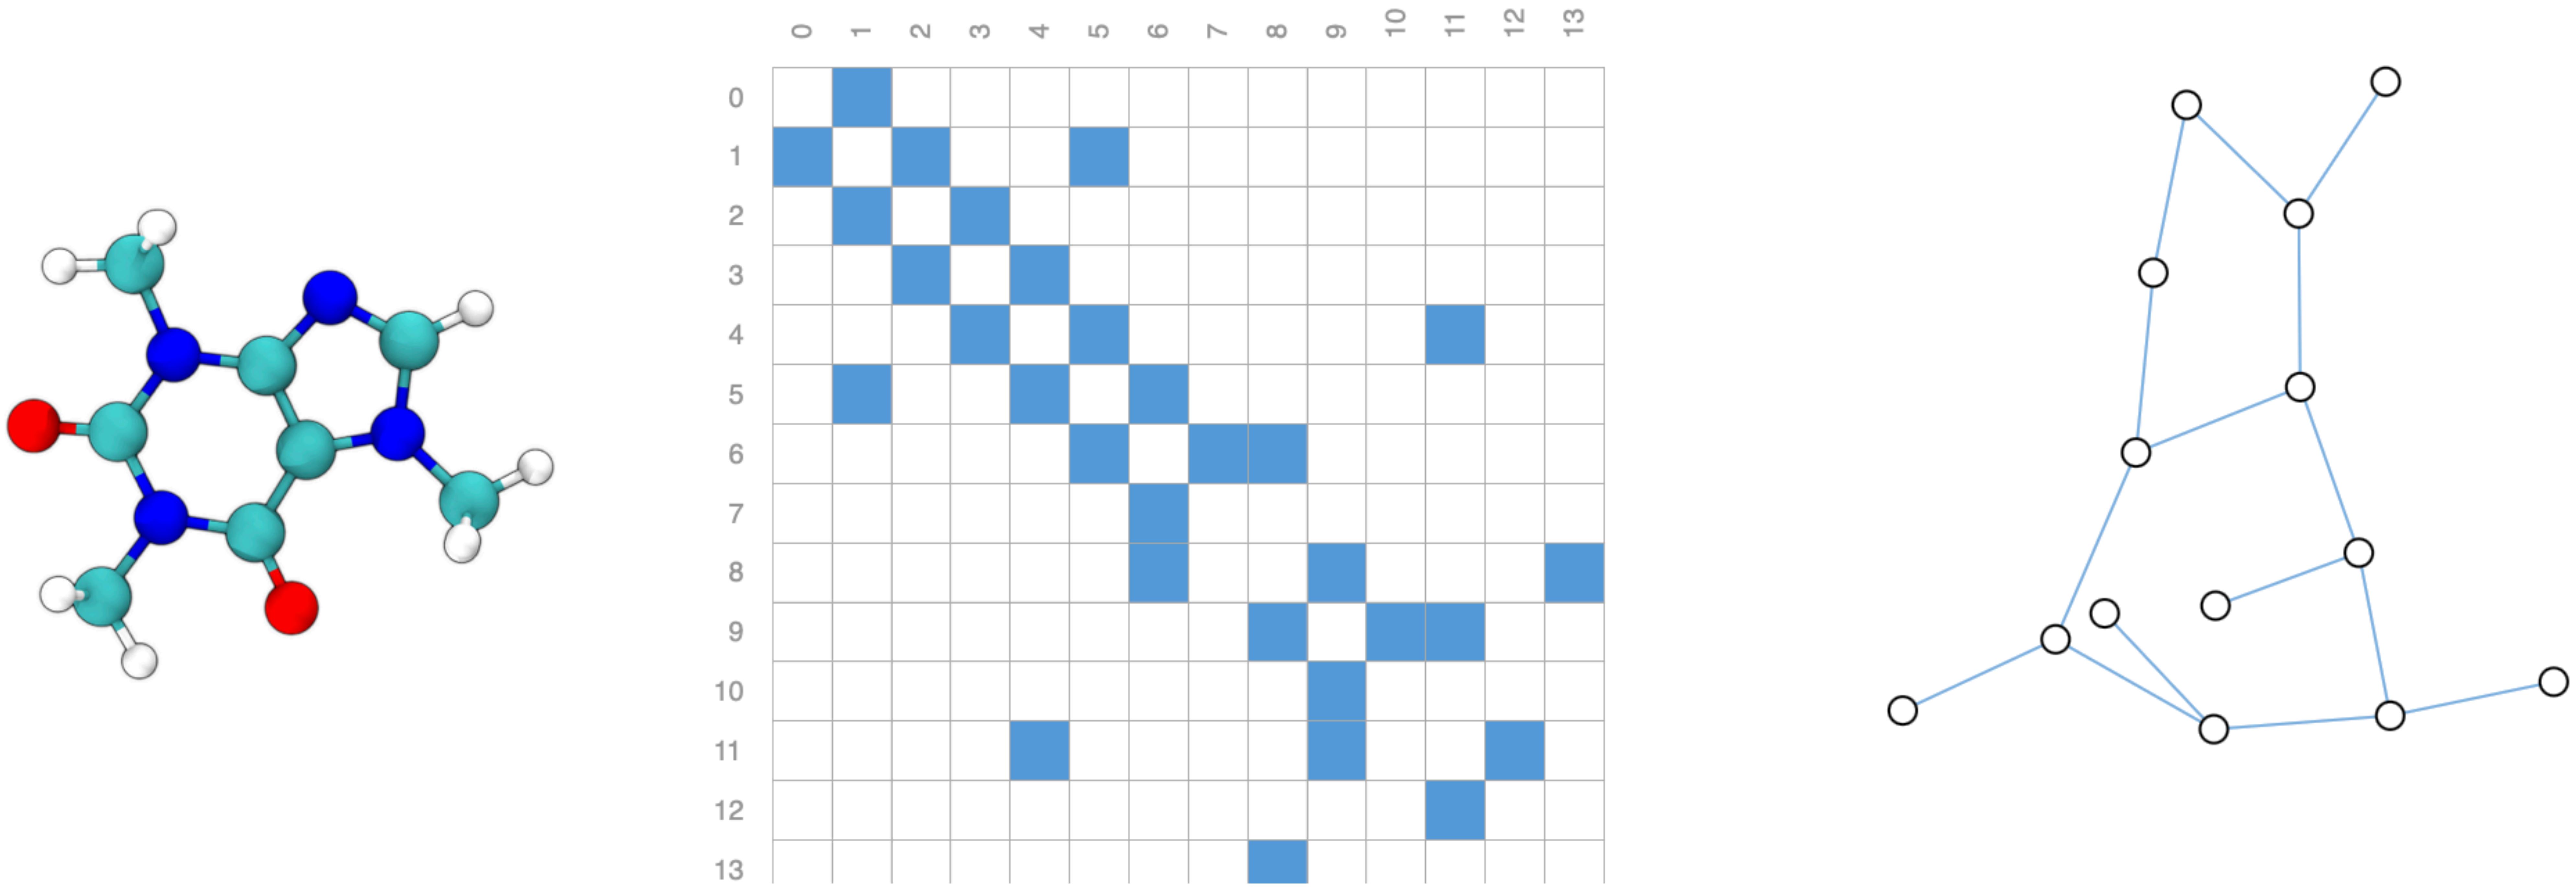

(Left) 3d representation of the Caffeine molecule (Center) Adjacency matrix of the bonds in the molecule (Right) Graph representation of the molecule

Graph-Valued Data

Social networks as graphs

- Represent groups of people by modeling individuals as nodes and their relationships as edges

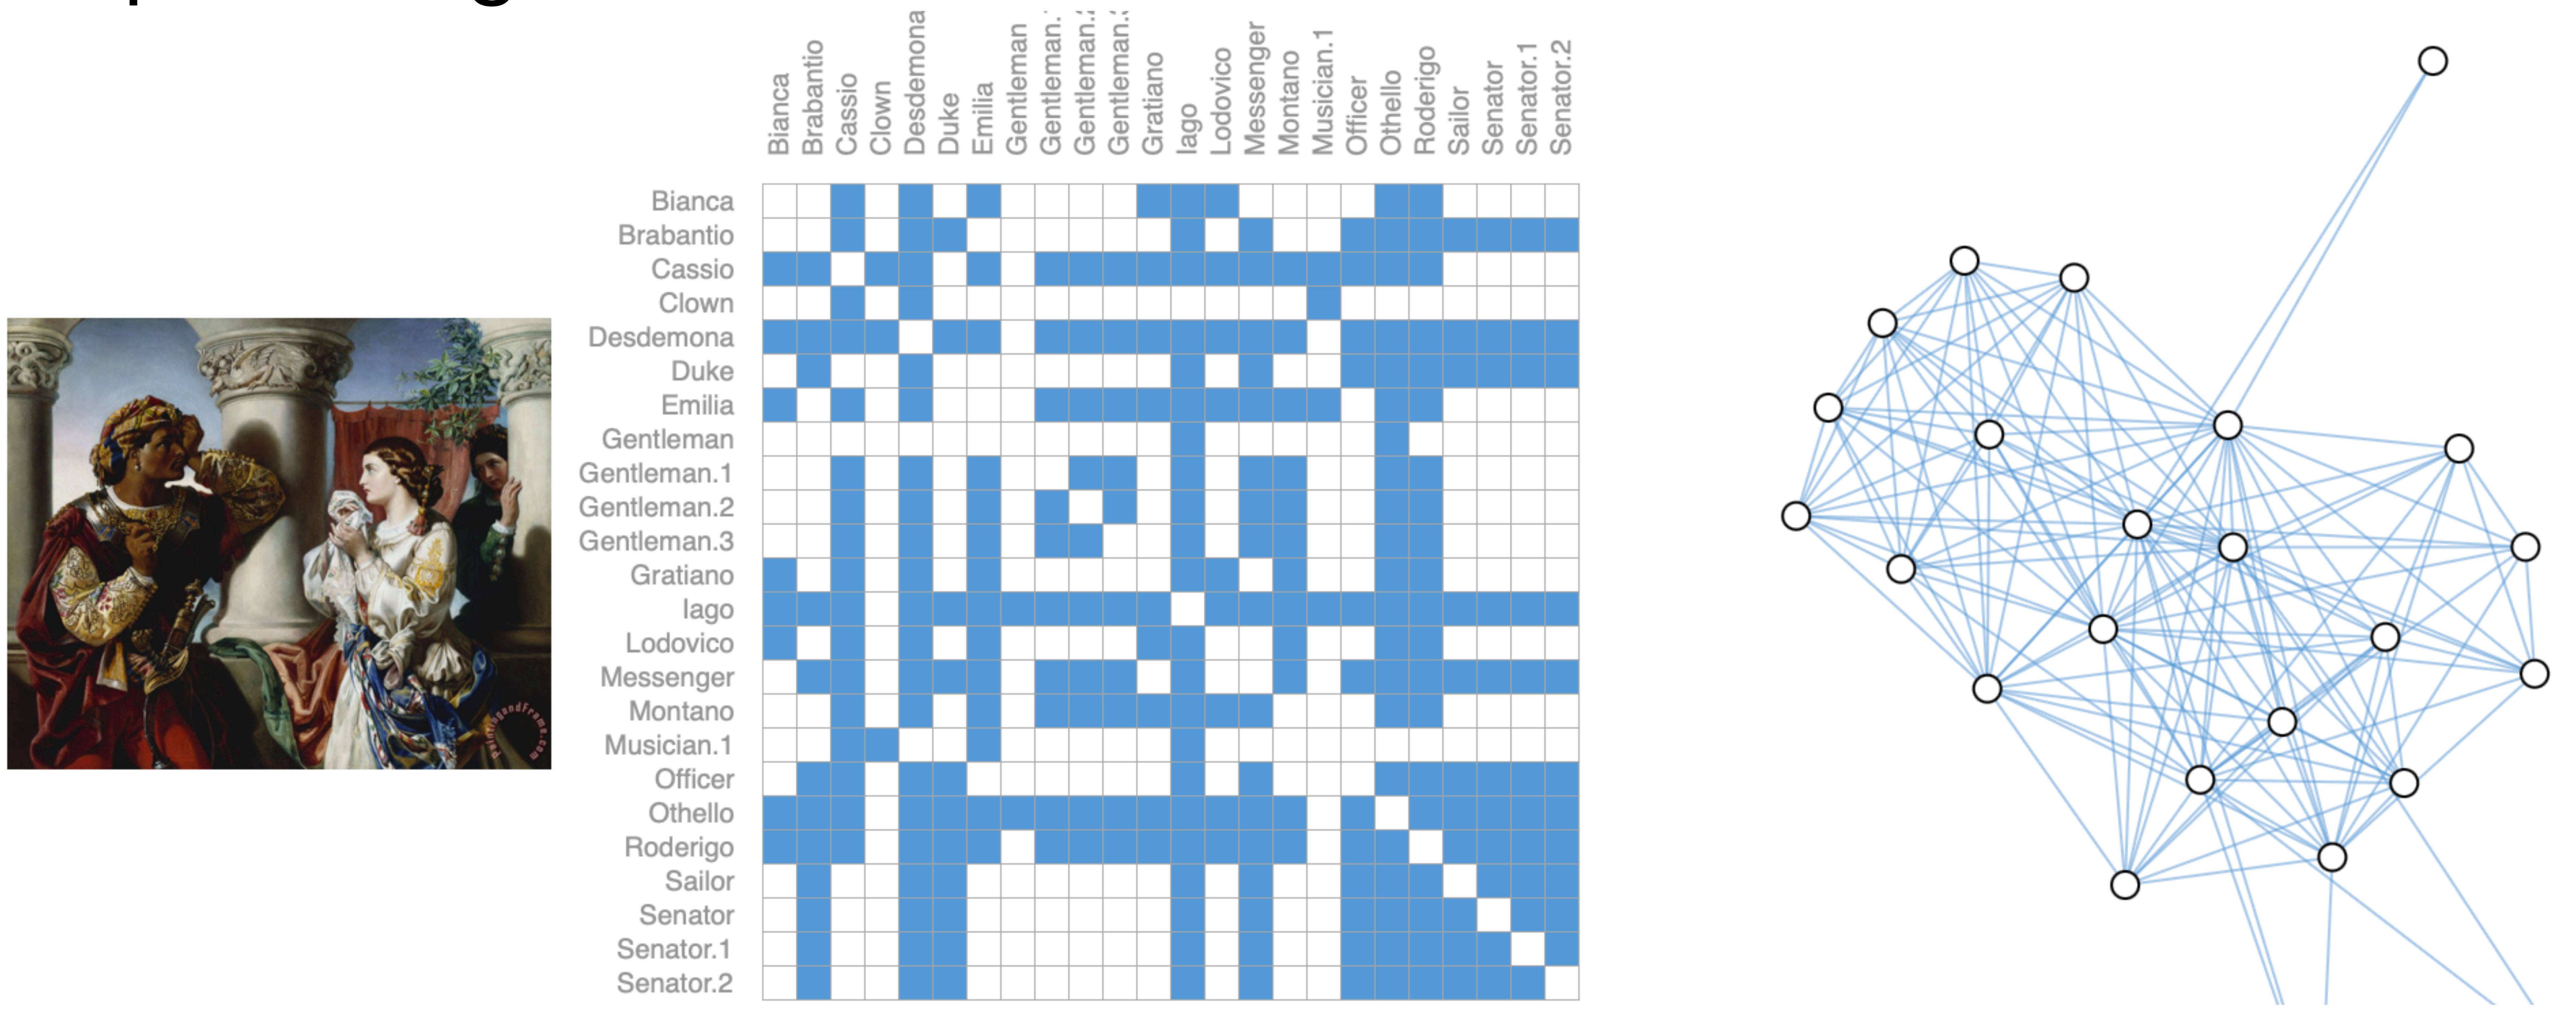

(Left) Image of the scene from the play “Othello” (Center) Adjacency matrix of the interaction between characters in the play (Right) Graph representation of these interactions

The Classes of Graph Prediction

What tasks do we want to perform on this data?

1. Graph-level task: predict a single property for the whole graph
2. Node-level task: predict some property for each node in a graph
3. Edge-level task: predict the property or presence of edges in a graph

Graph-Level Task

Predict the property of an entire graph

- Predict what the molecule smells like, or if it will bind to a receptor implicated in a disease, etc.

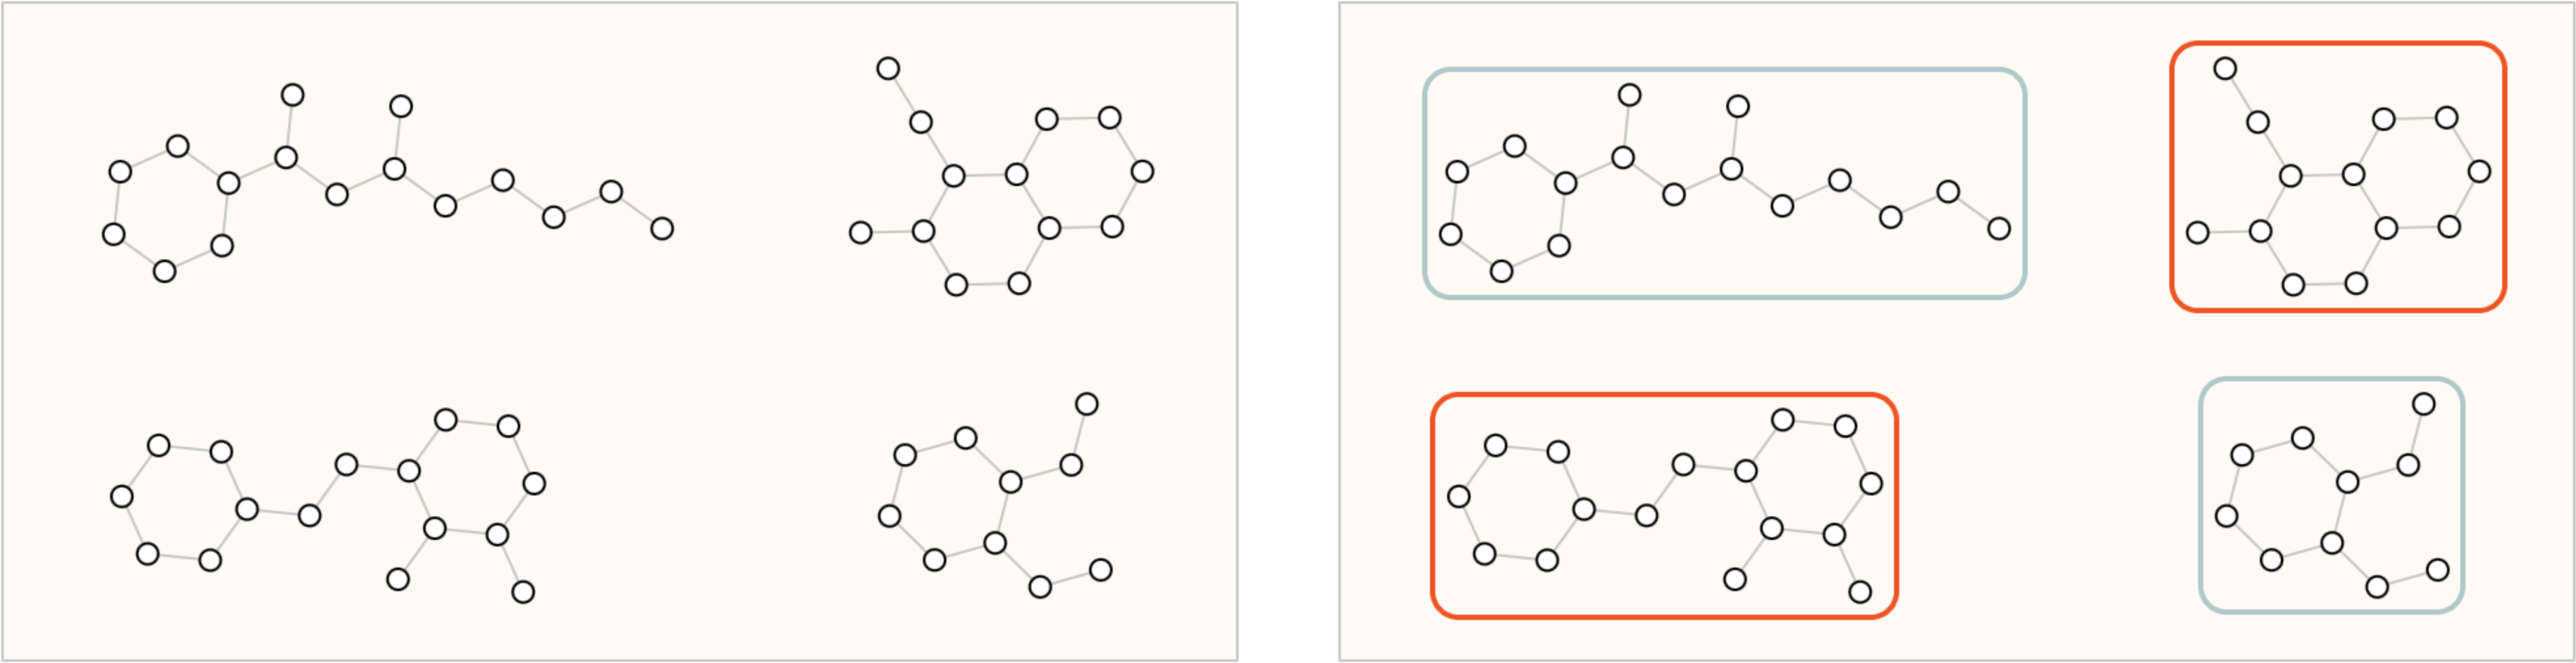

(Input) graphs (Output) labels for each graph (e.g., “does the graph contain two rings?”)

- With text, one could do sentiment analysis (identify mood or emotion of an entire sentence)

Node-Level Task

Predict identity/role of each node in a graph

- Example: Zach’s karate club

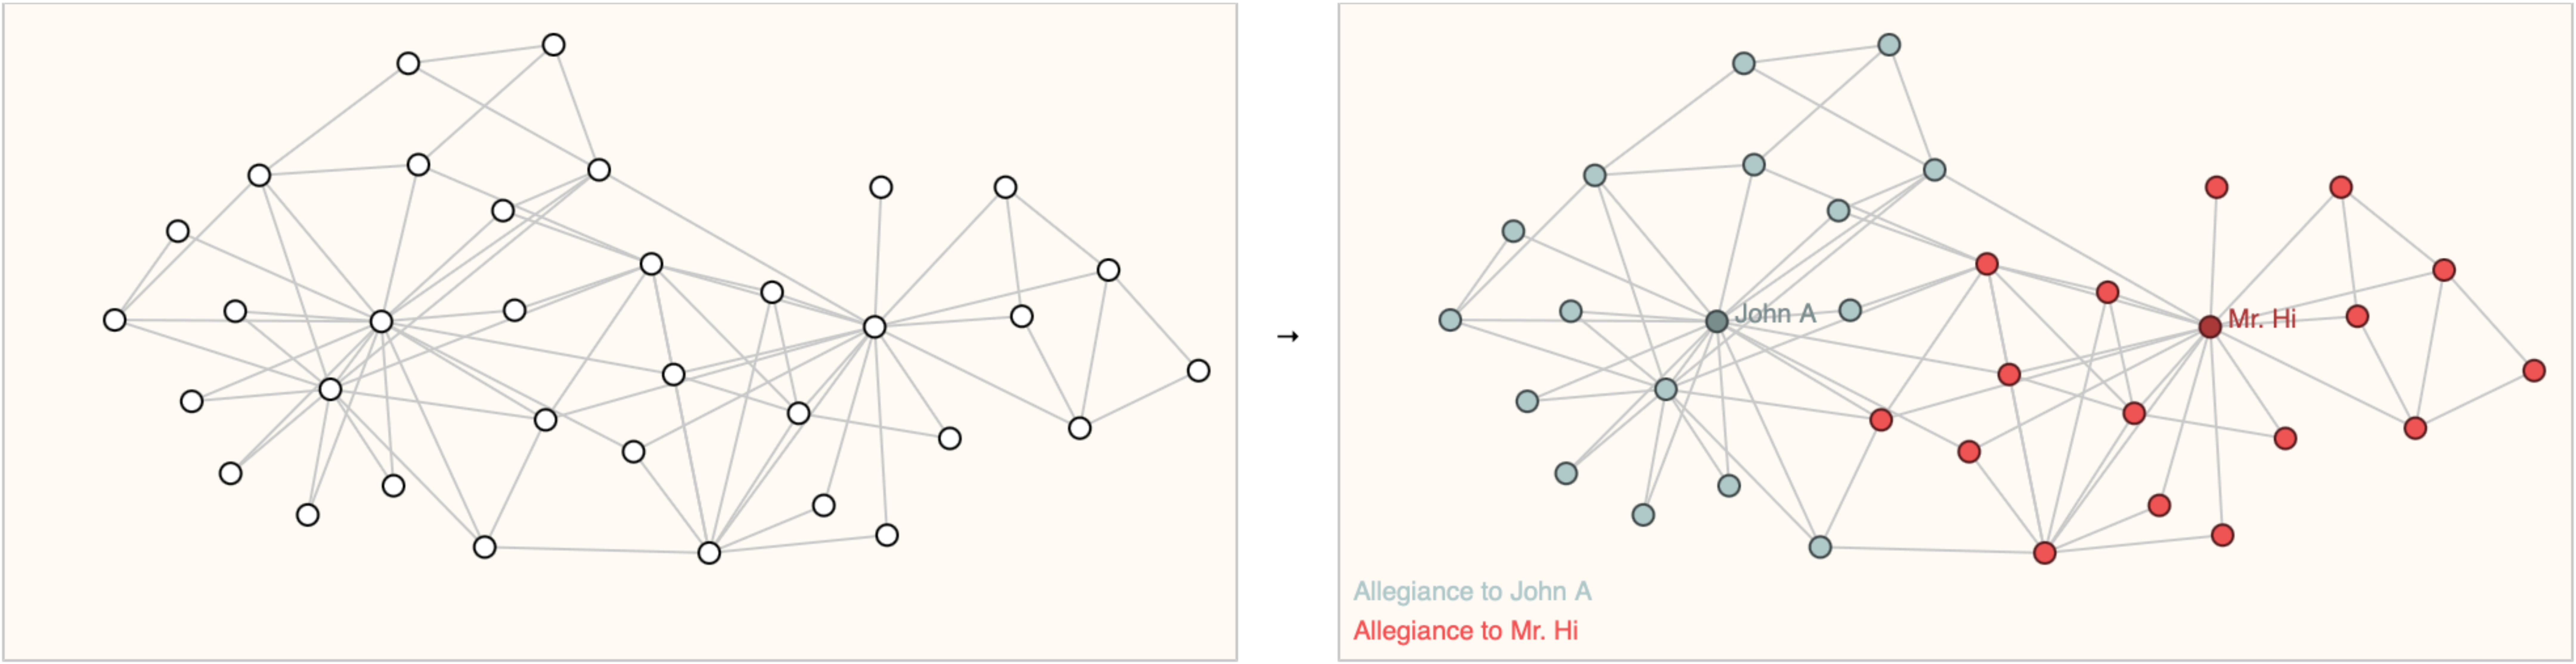

(Input) graph with unlabeled nodes (Output) graph node labels

Edge-Level Task

Predict property or presence of edges in a graph

- Example: Image scene understanding

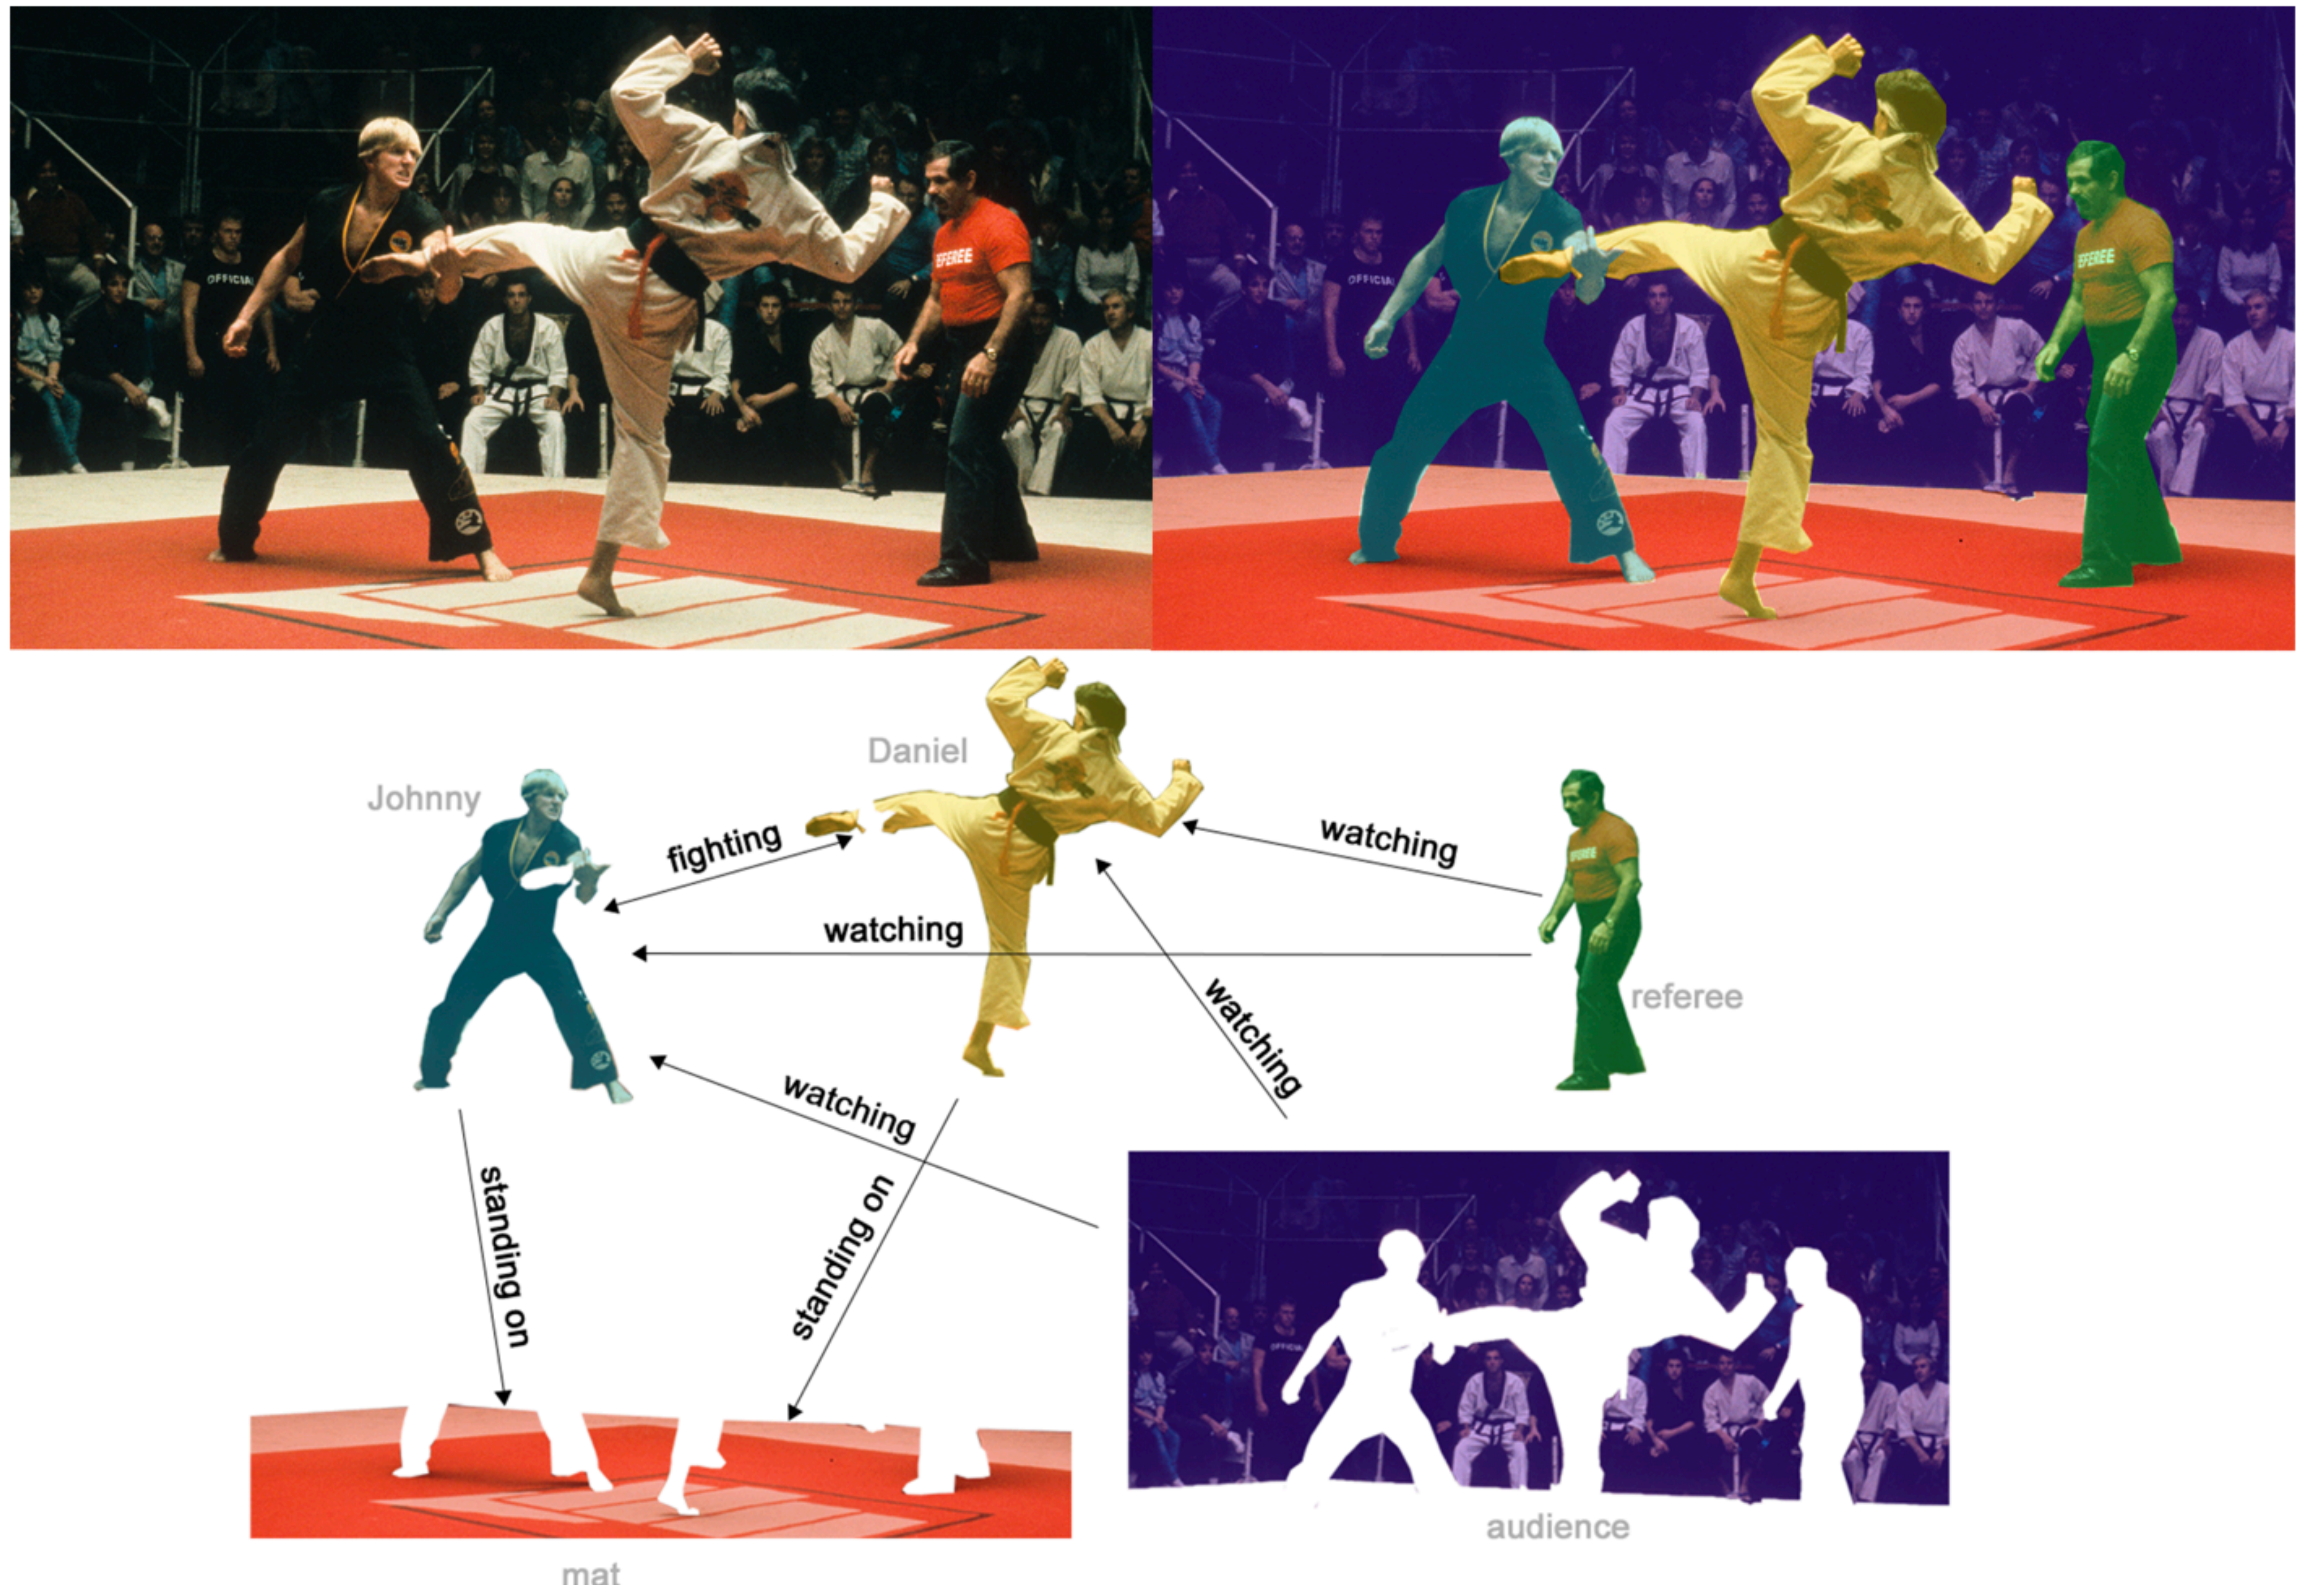

Image segmented into five entities: each fighter, referee, audience, and the mat. Also shows the relationships between these entities.

Edge-Level Task

Predict property or presence of edges in a graph

- Example: Image scene understanding

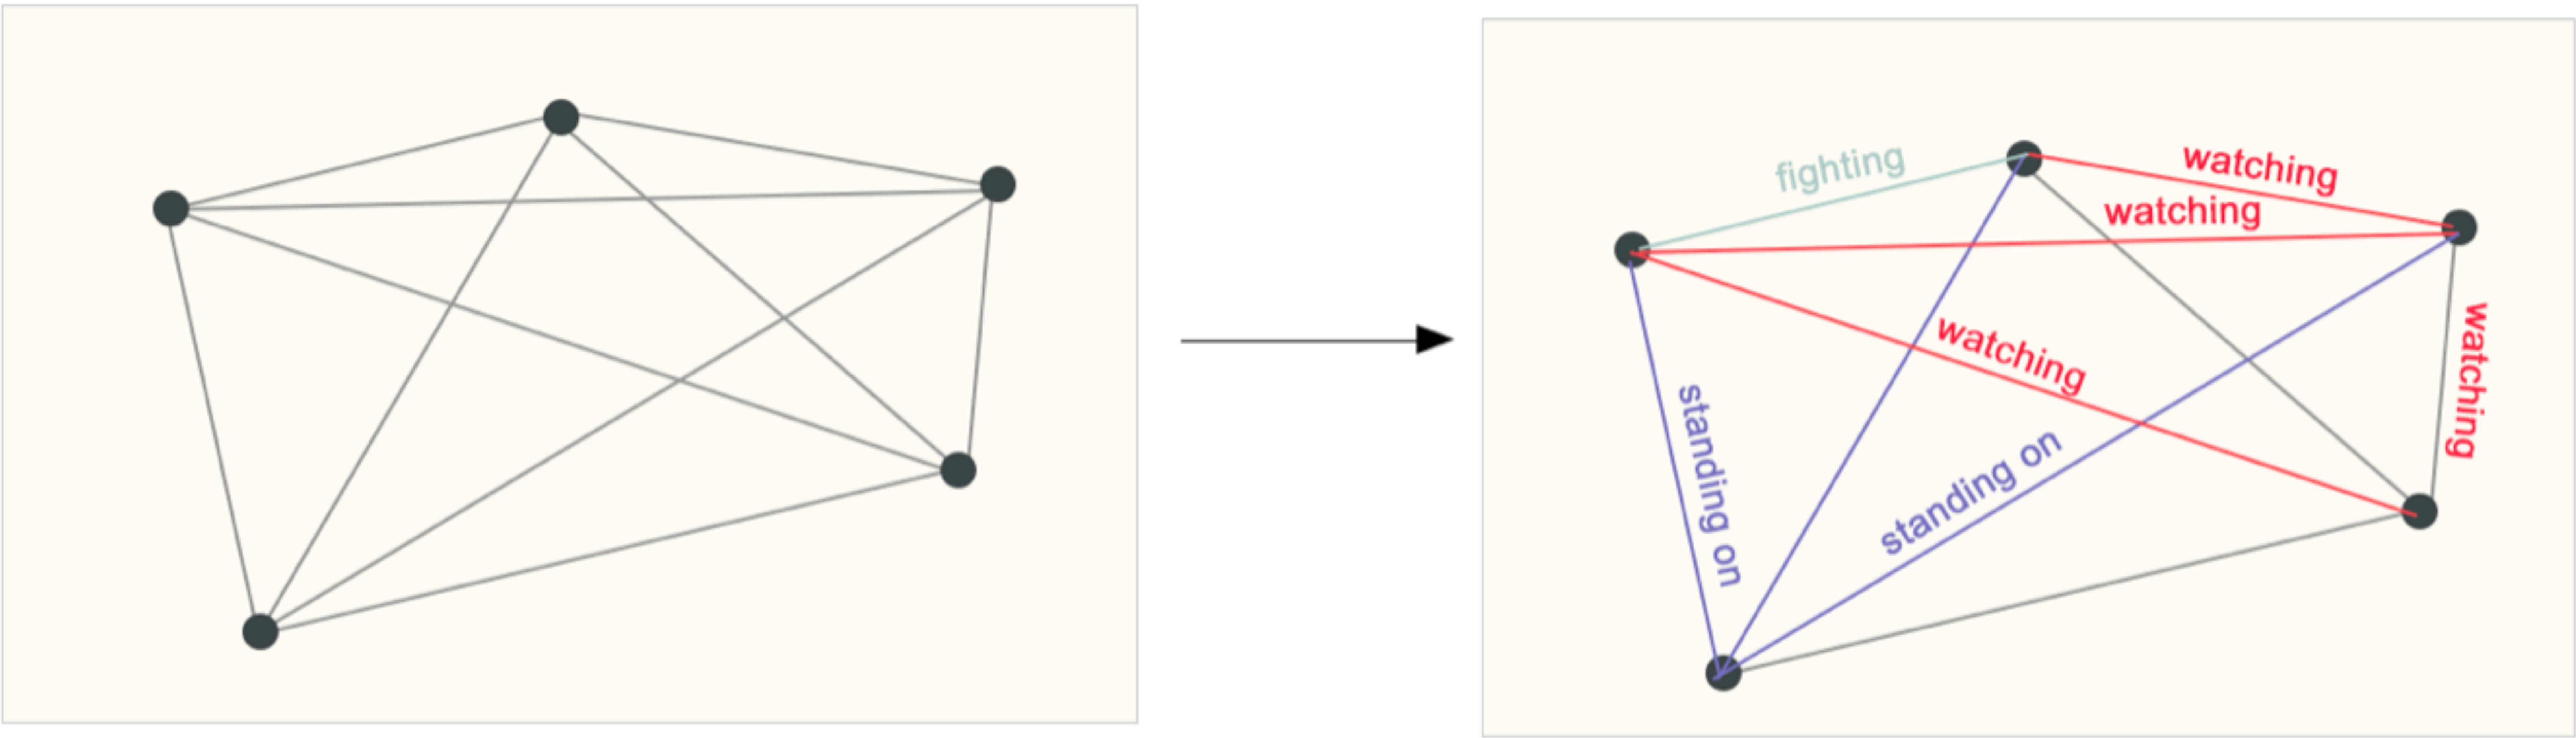

(Left) initial graph built from the previous visual scene (Right) example edge-labeling of the graph

Representing Connectivity

It's the hard part

- Adjacency matrix

1. Sparse (space-inefficient)
2. Not permutation invariant

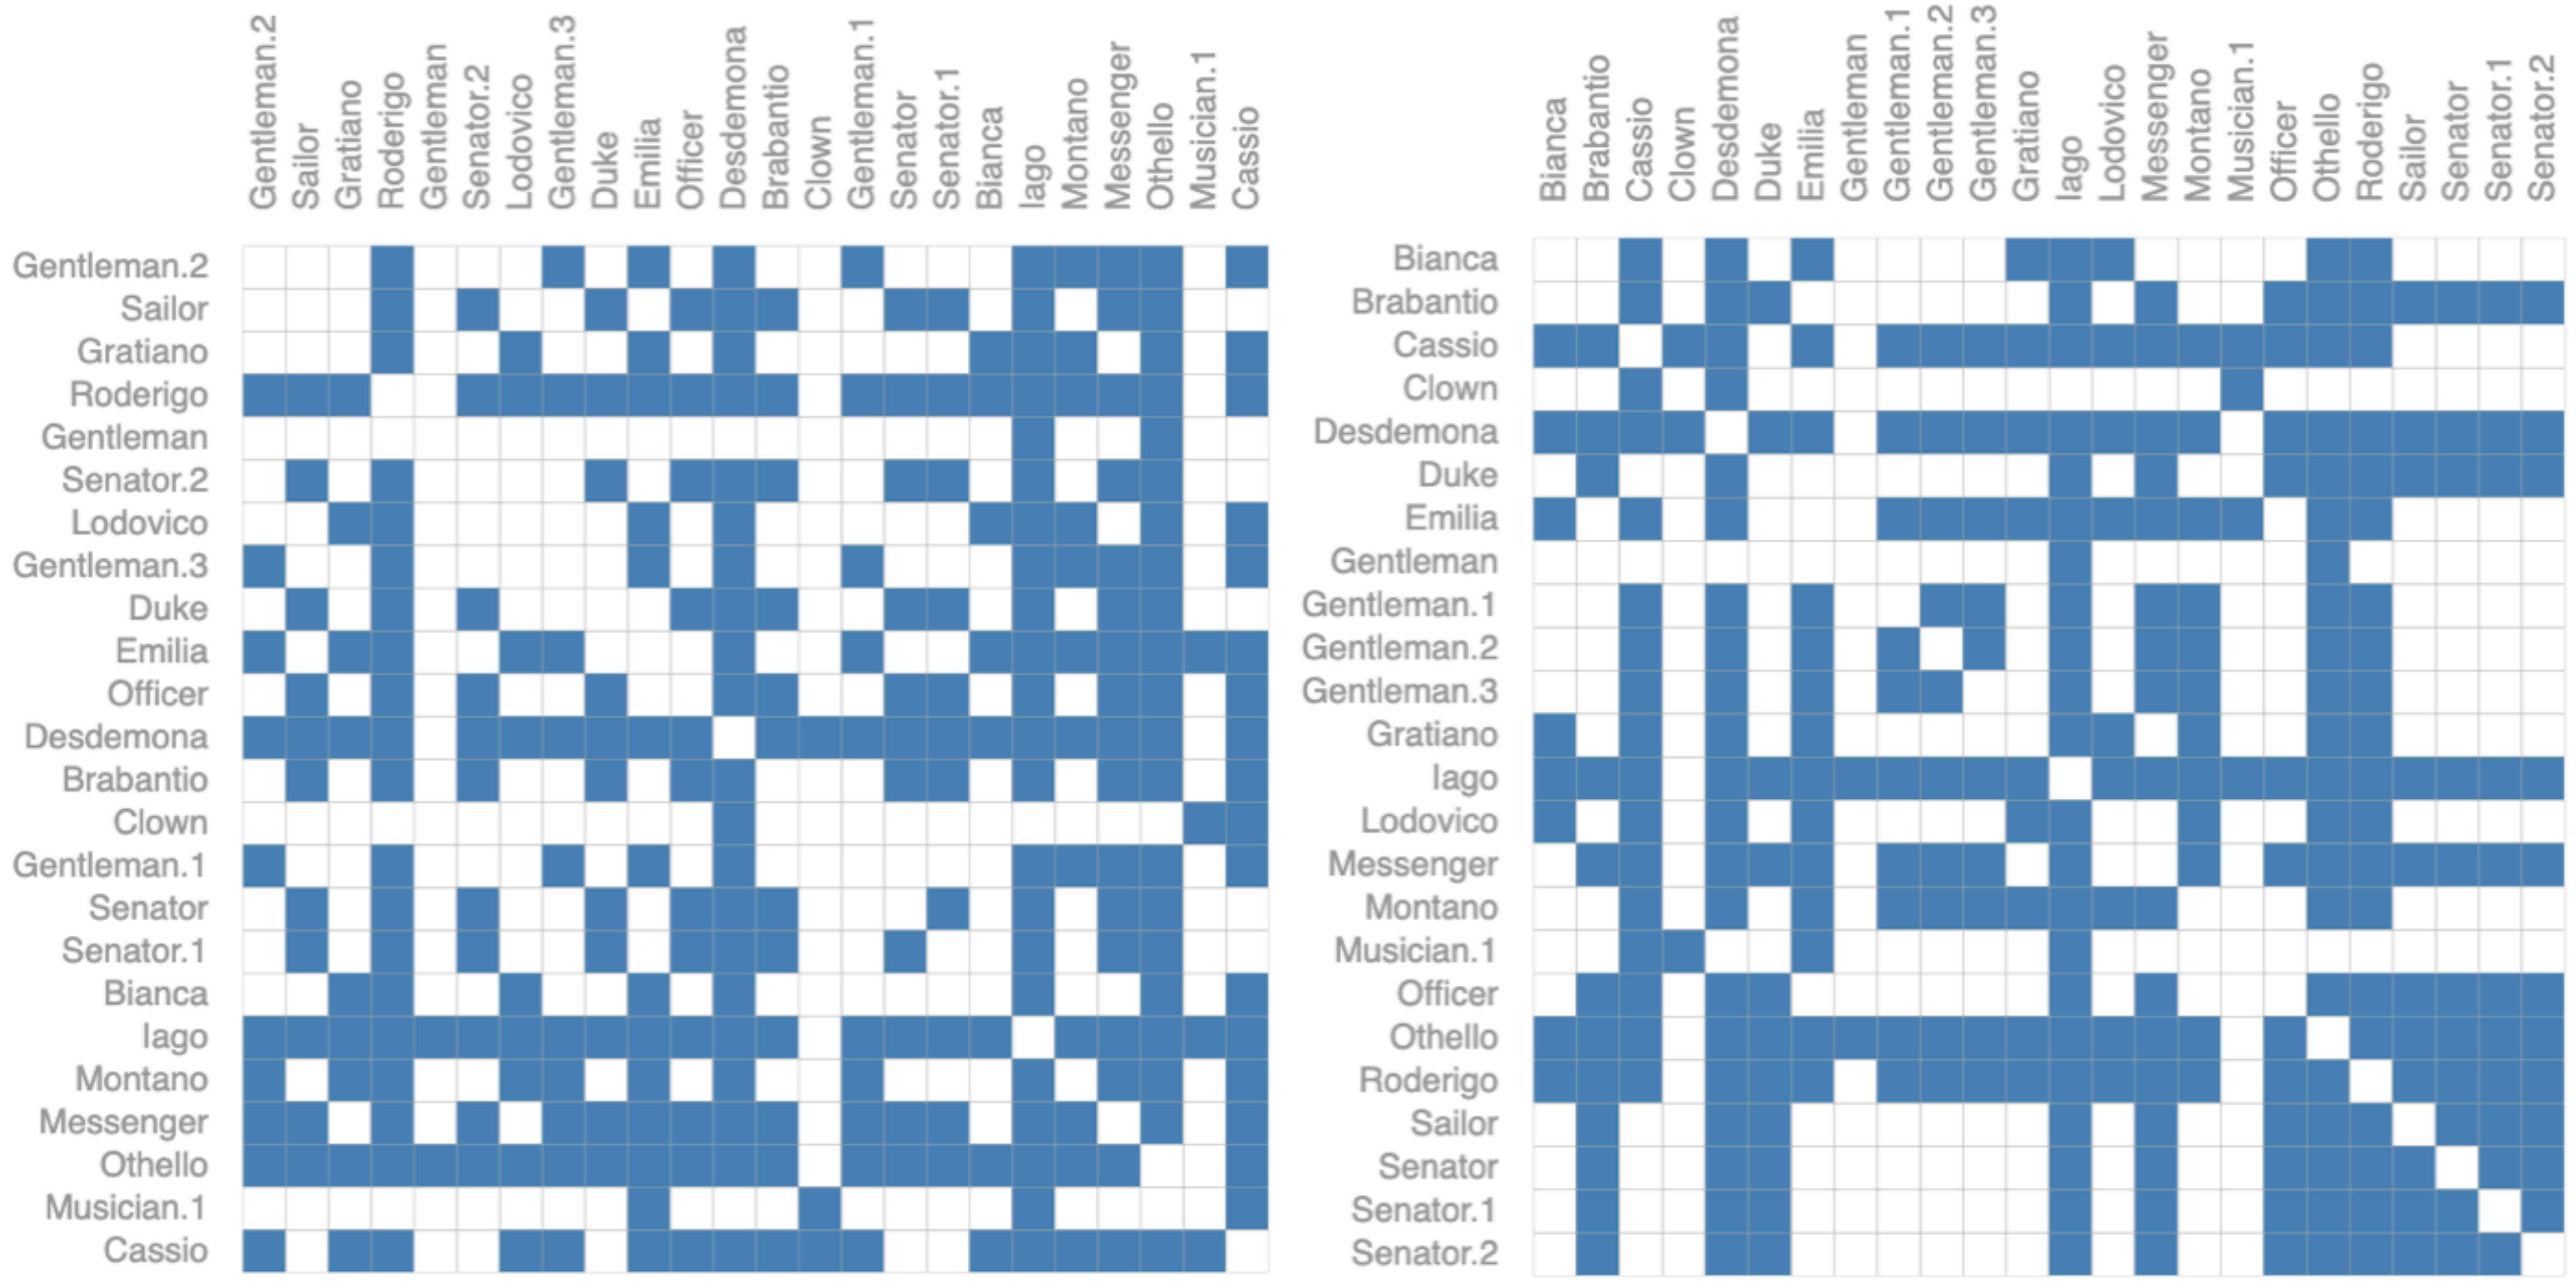

Two adjacency matrices representing the same graph

Representing Connectivity

It’s the hard part

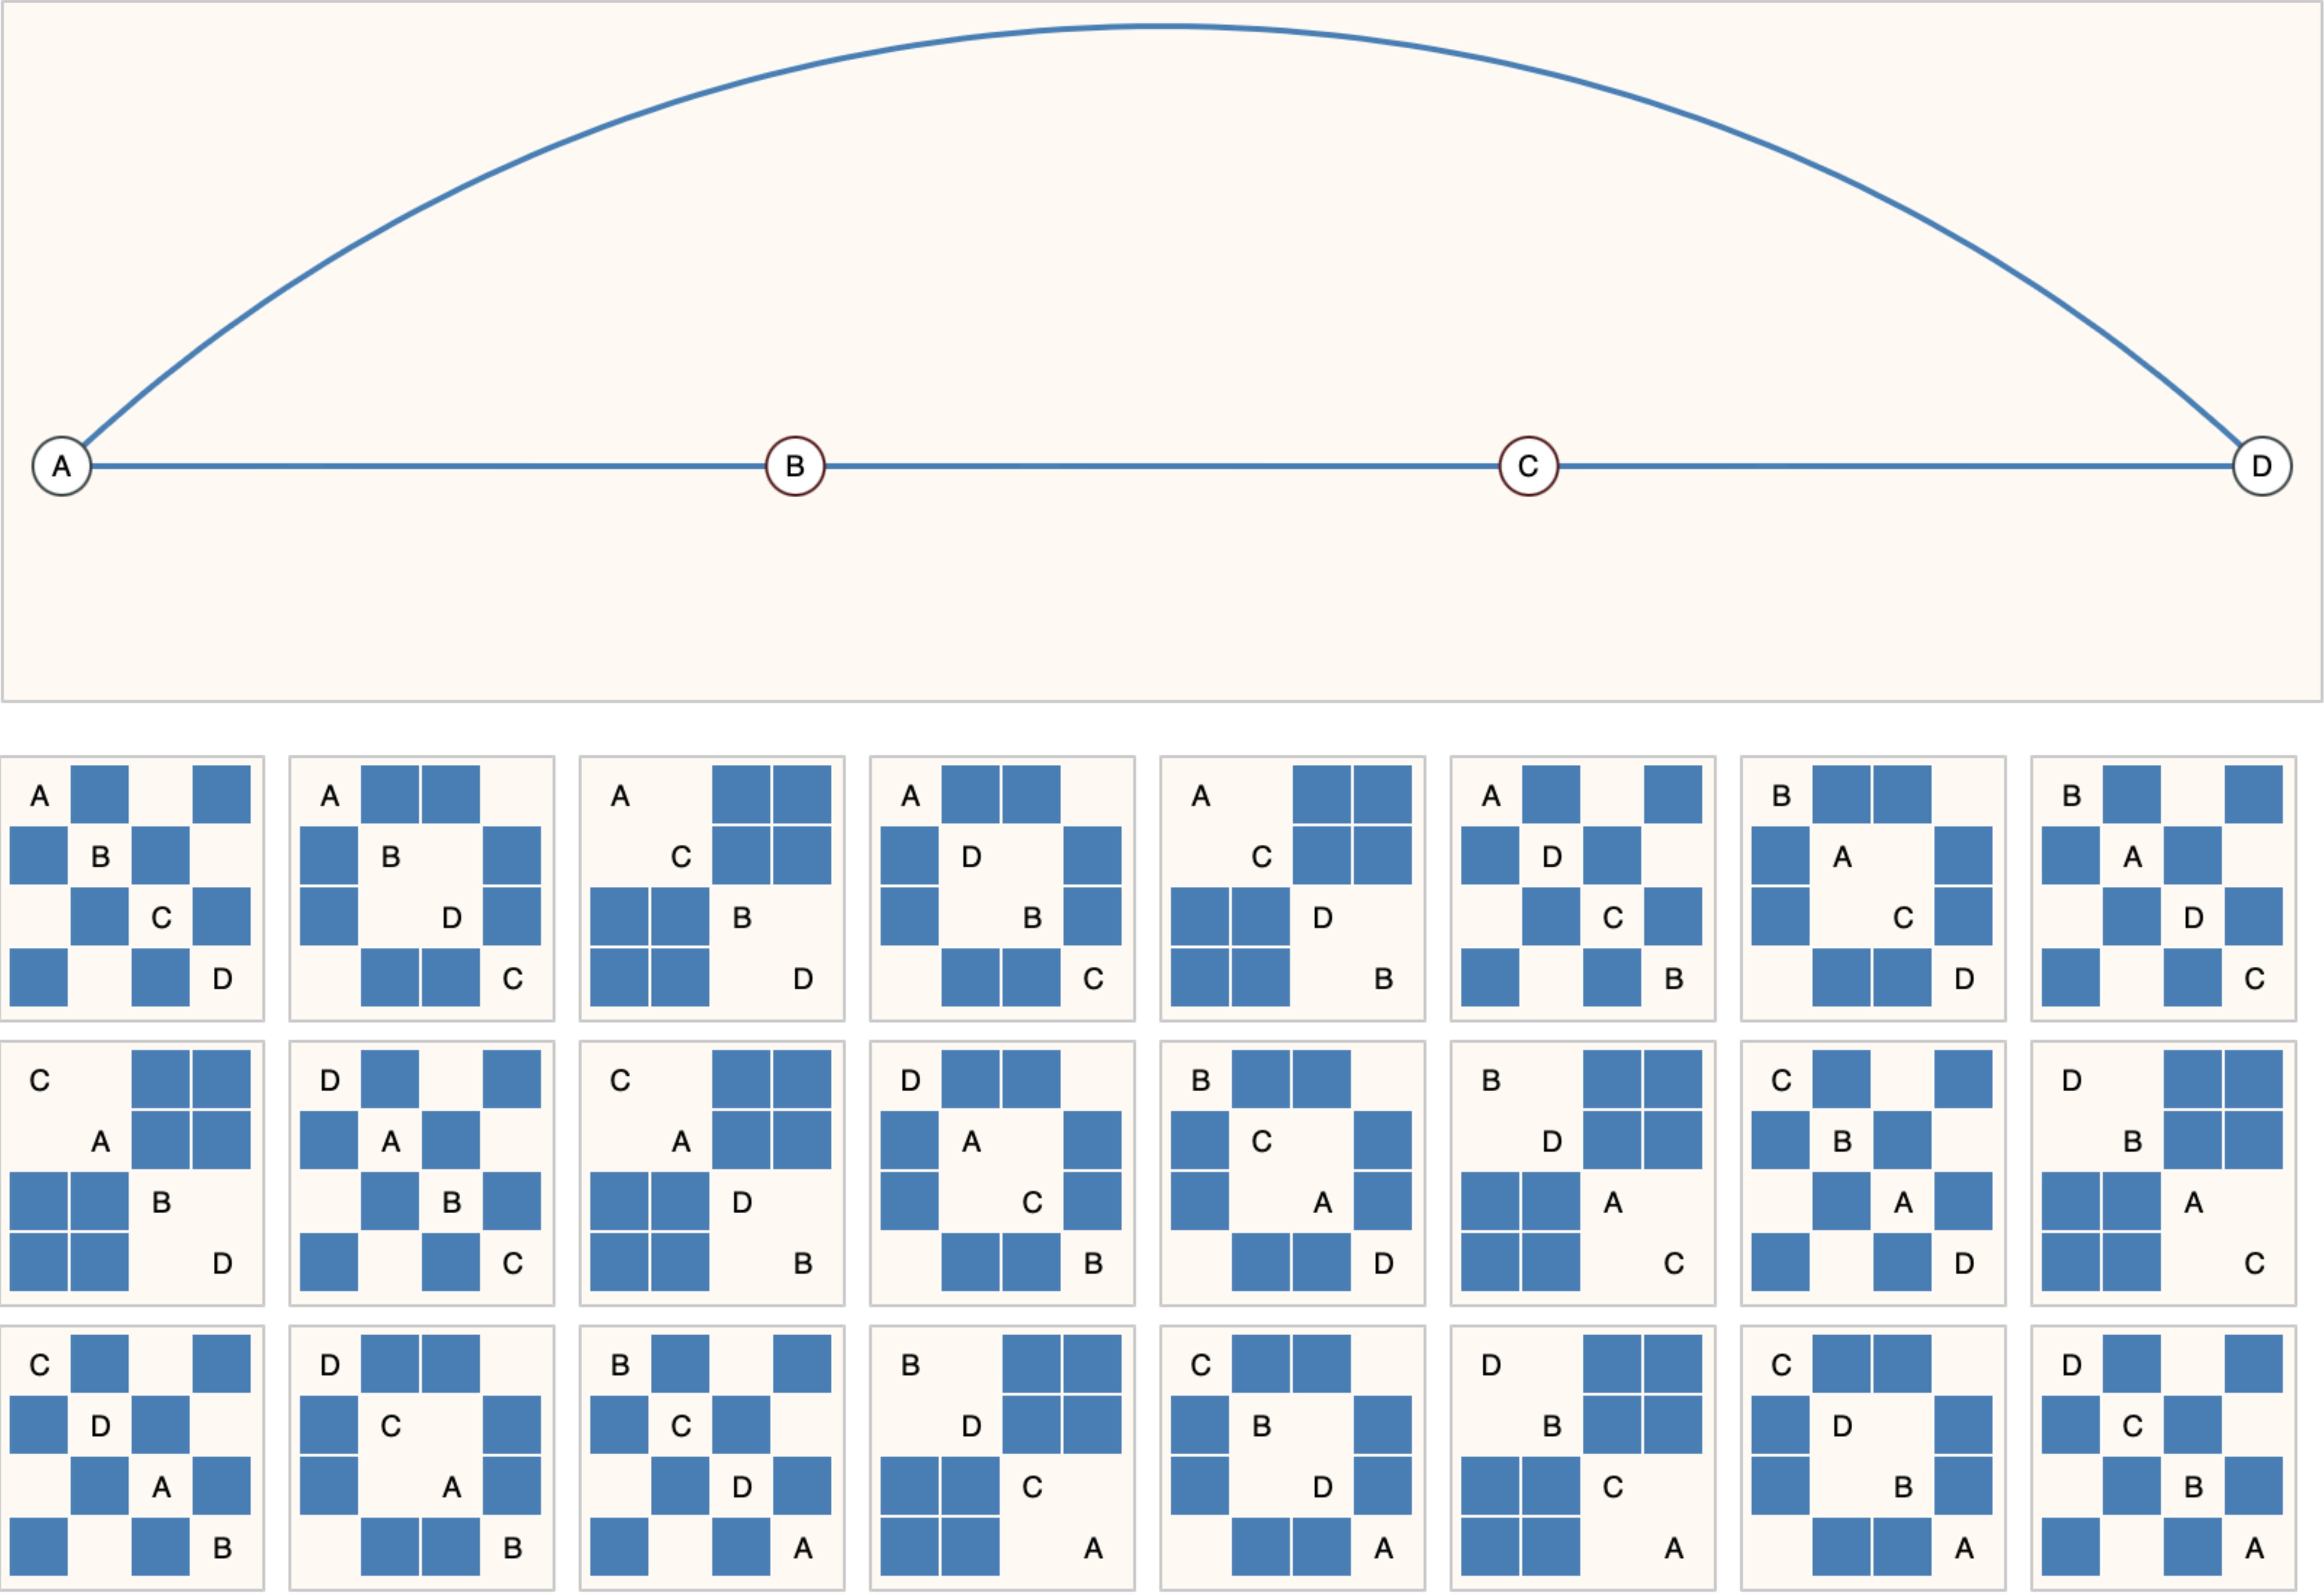

Every adjacency matrix that can describe a small graph of 4 nodes

Representing Connectivity

It’s the hard part

- Adjacency lists - describe connectivity of edge e_k between nodes n_i and n_j in the k^{th} entry

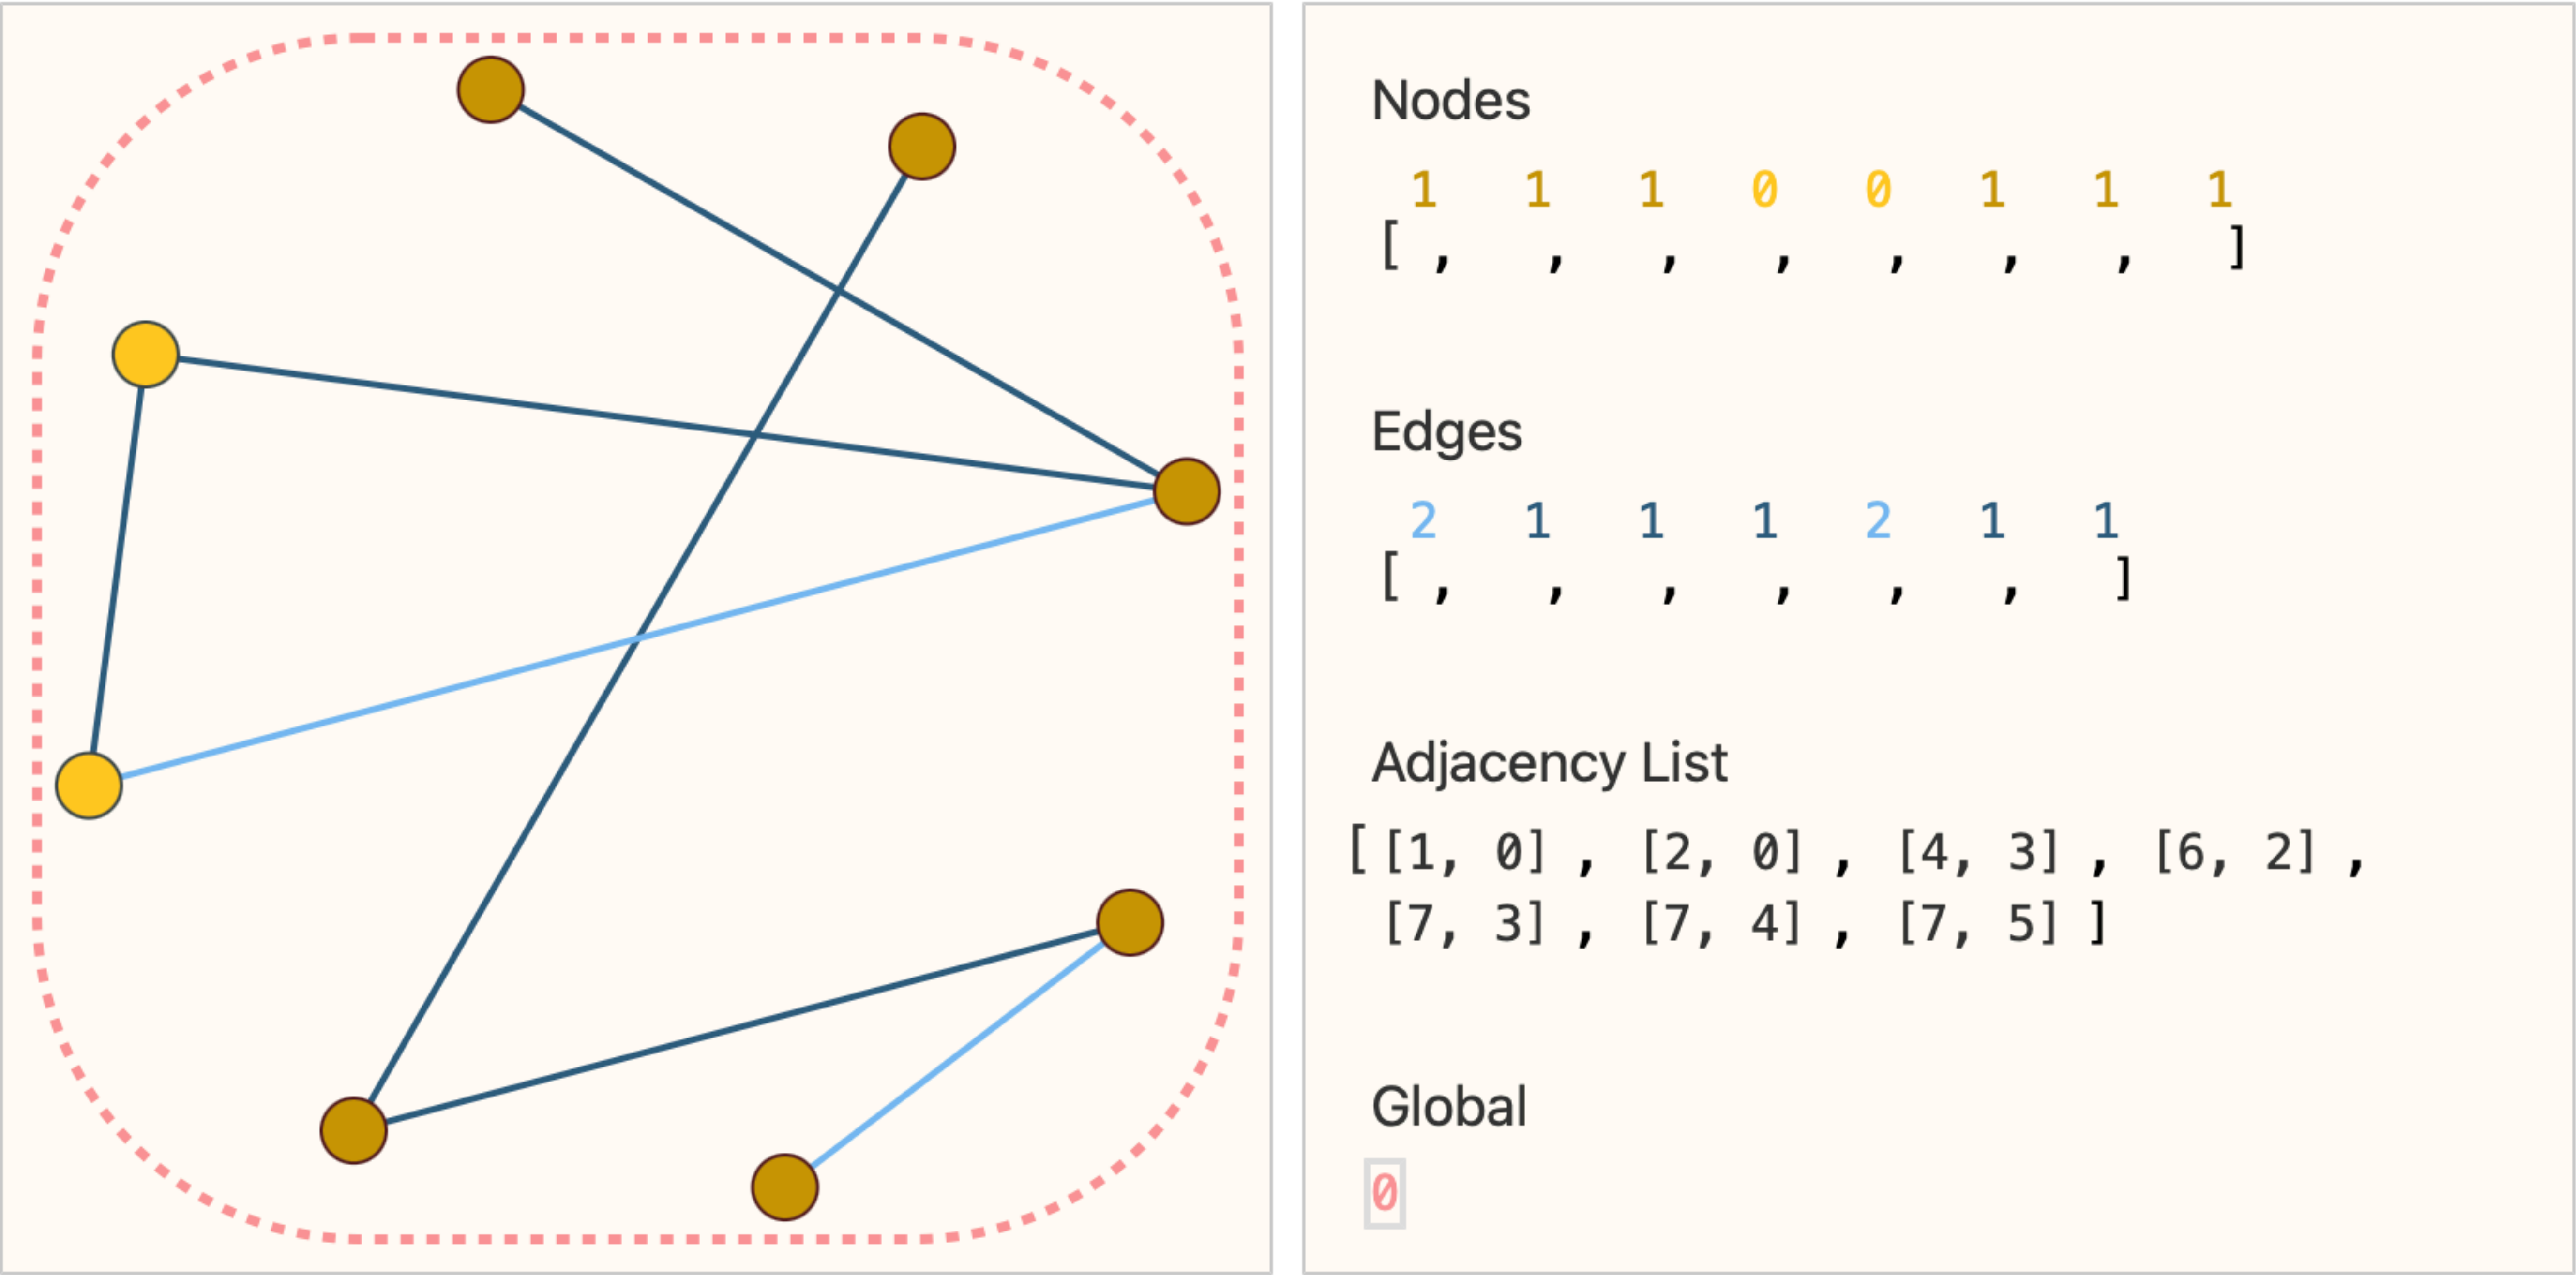

Build a Graph Neural Network

Using graph neural networks to solve graph prediction tasks

- These model types accept a graph as input, with information loaded into its nodes, edges, and global context
- Progressively transform embeddings without changing the connectivity of the input graph

Predictions by Pooling Information

How do we make predictions?

- For each node embedding, apply a linear classifier

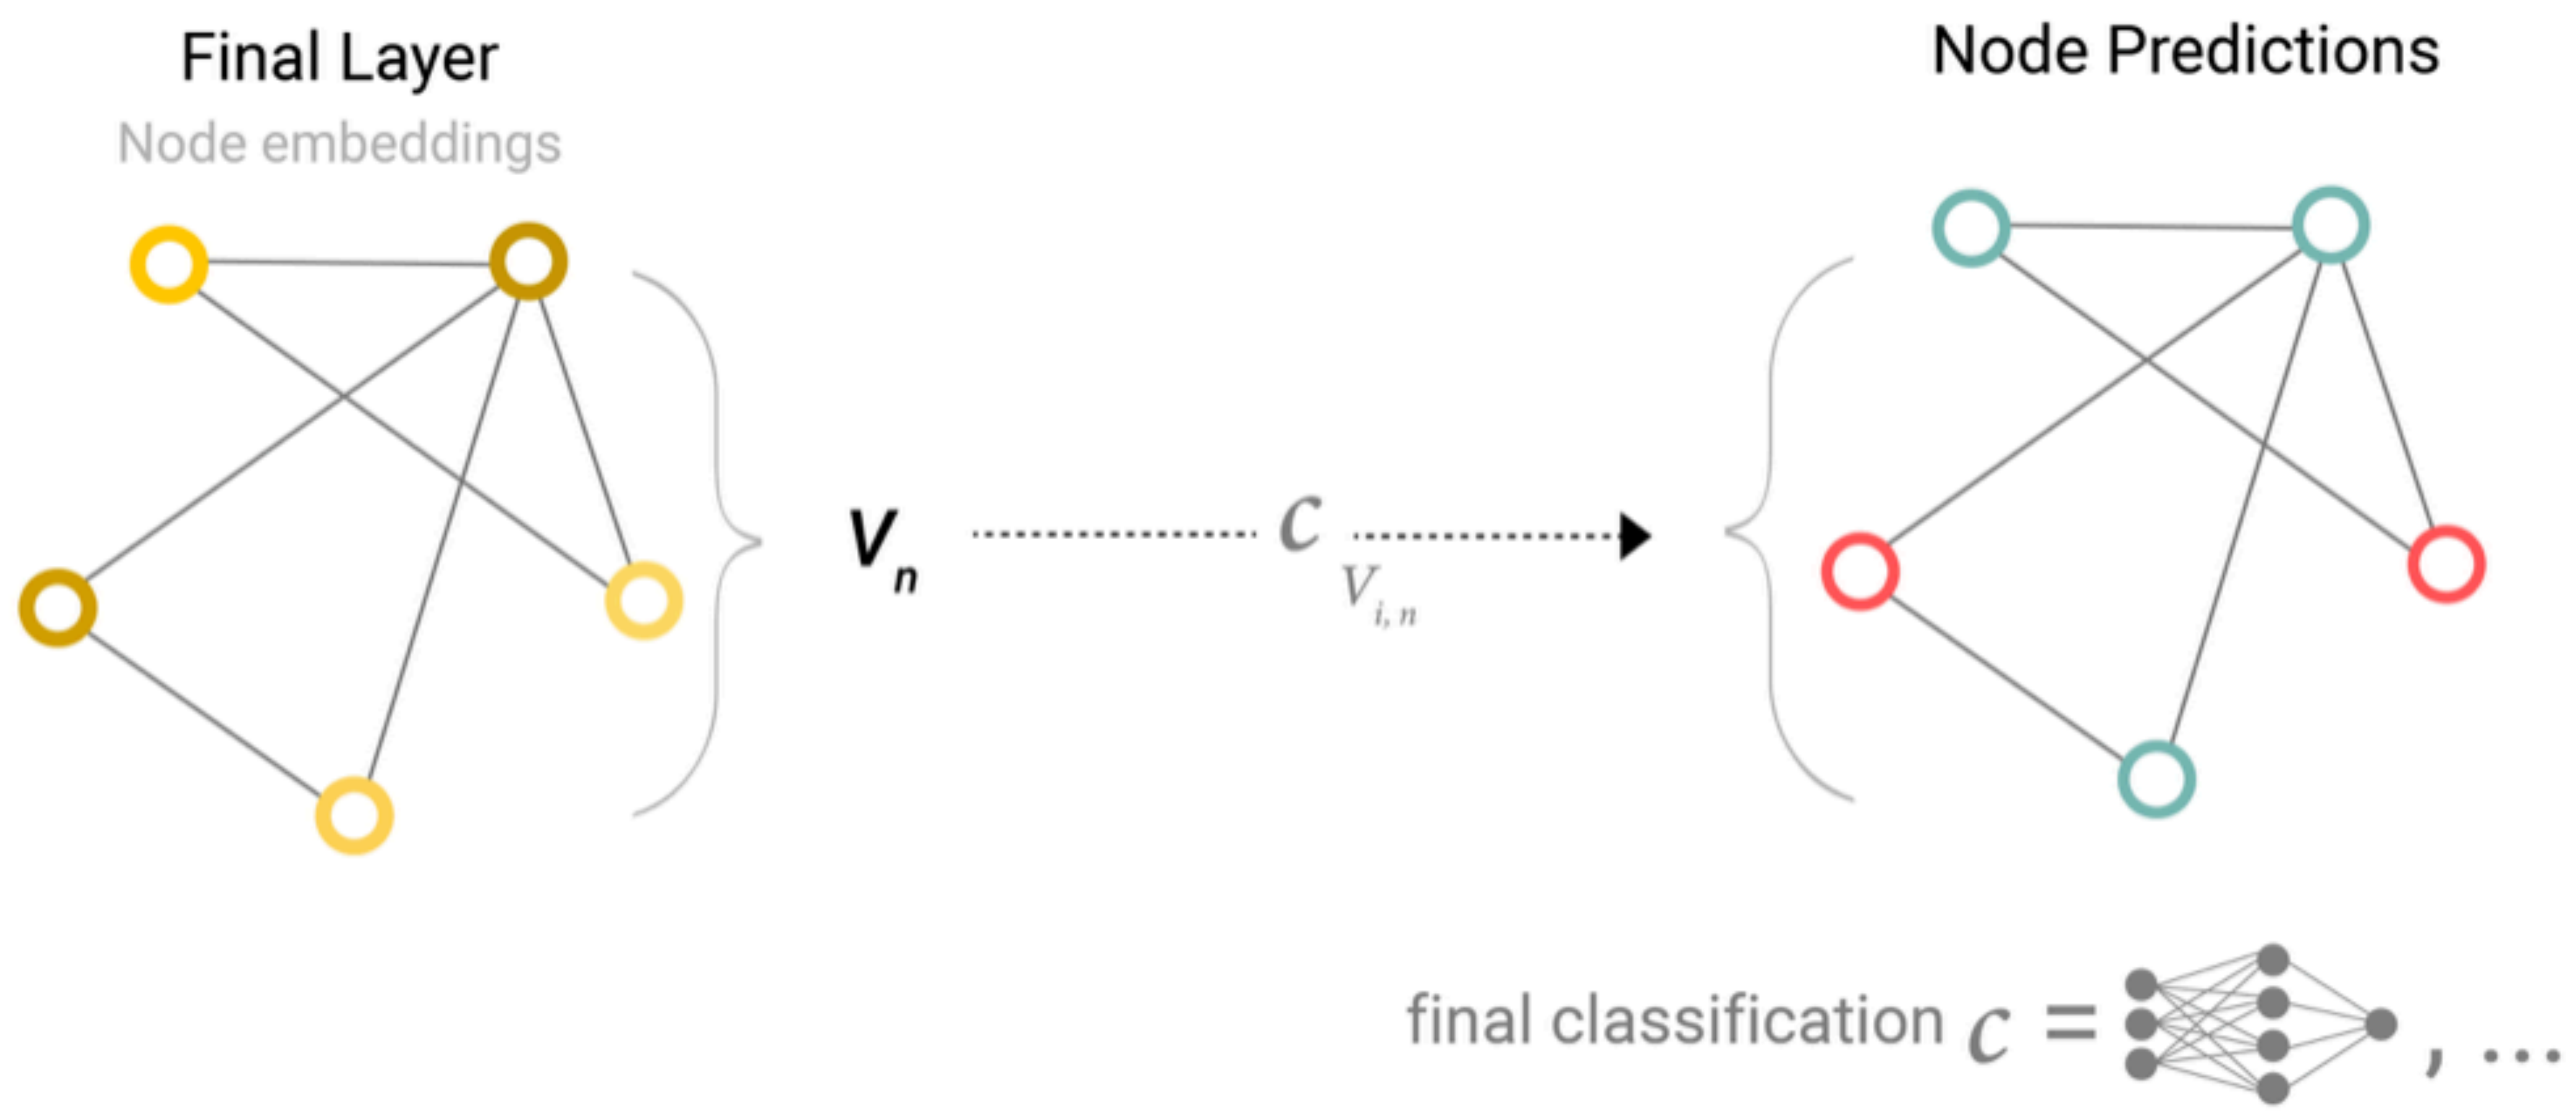

Predictions by Pooling Information

1. For each item to be pooled, *gather* each of their embeddings and concatenate them into a matrix

2. The gathered embeddings are then *aggregated*, usually via a sum operation

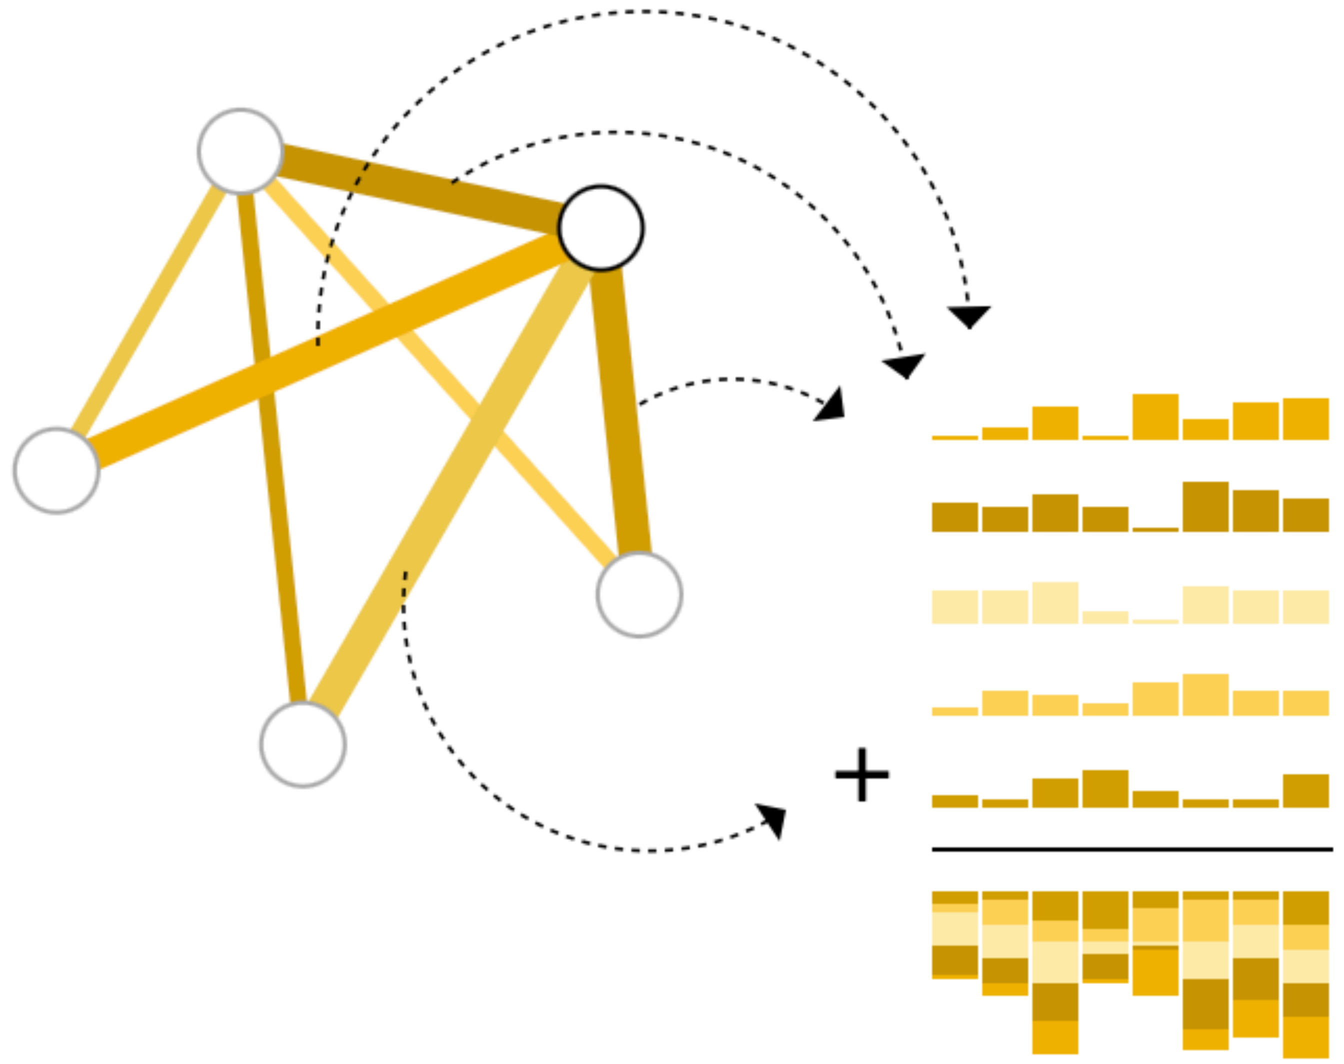

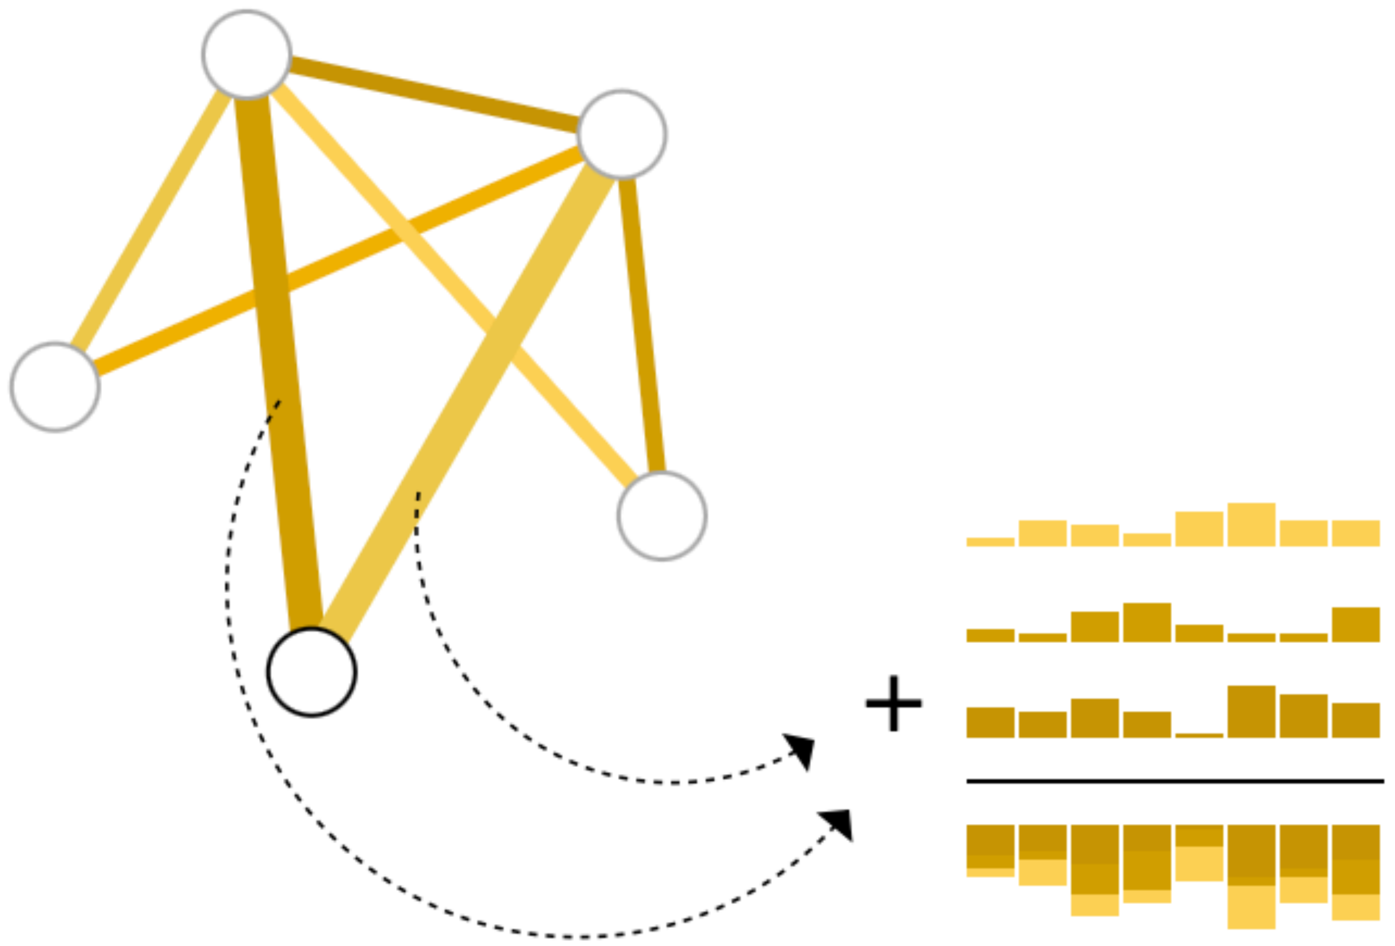

Predictions by Pooling Information

- If we only have edge-level features, use pooling to route information where it needs to go

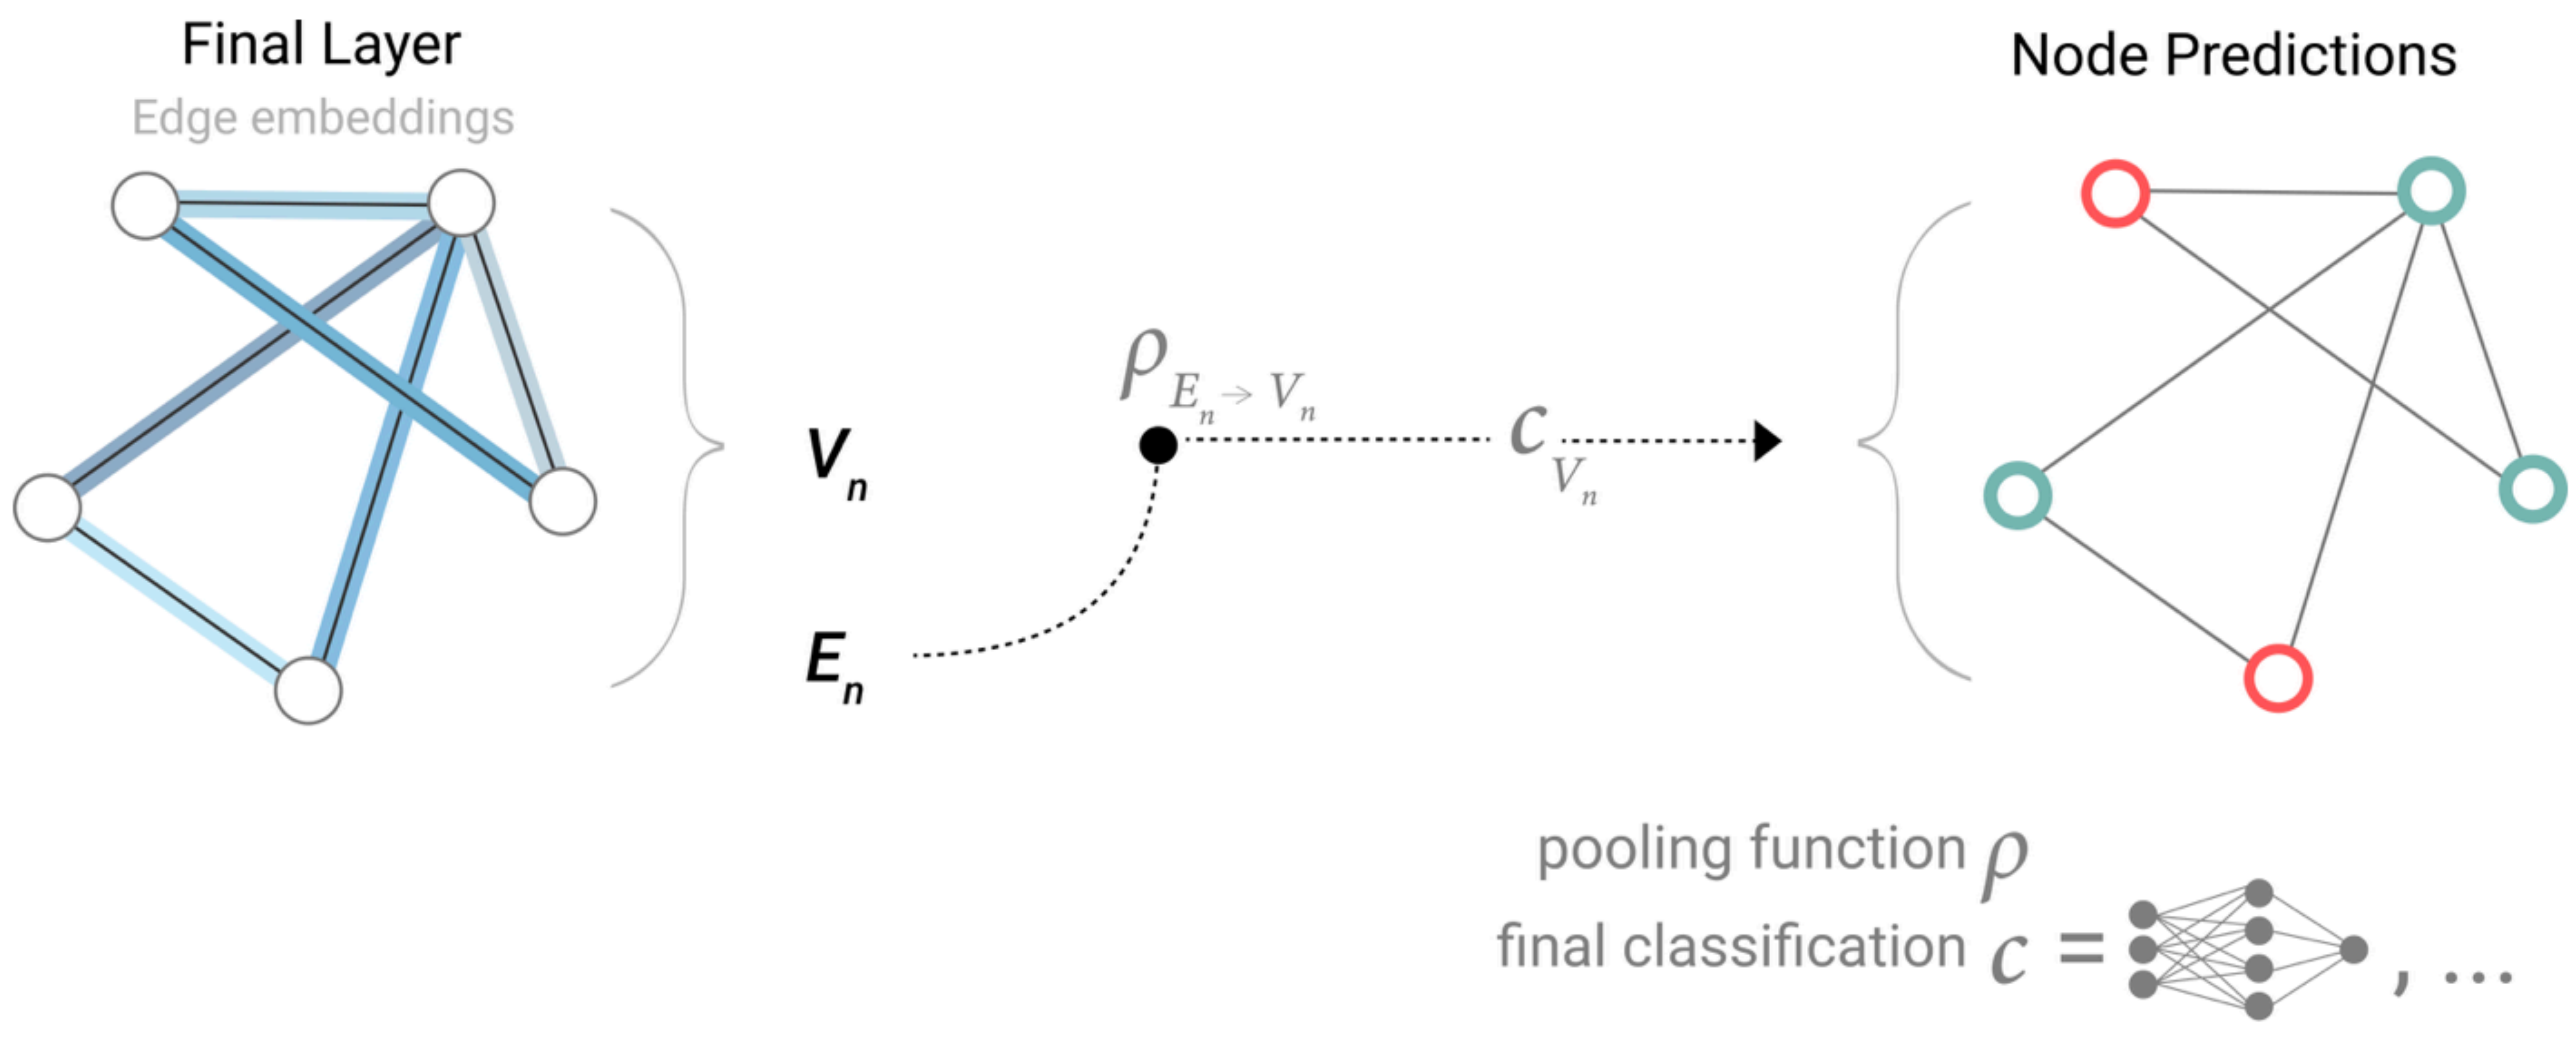

Predictions by Pooling Information

- If we only have node-level features and we are trying to predict edge-level information

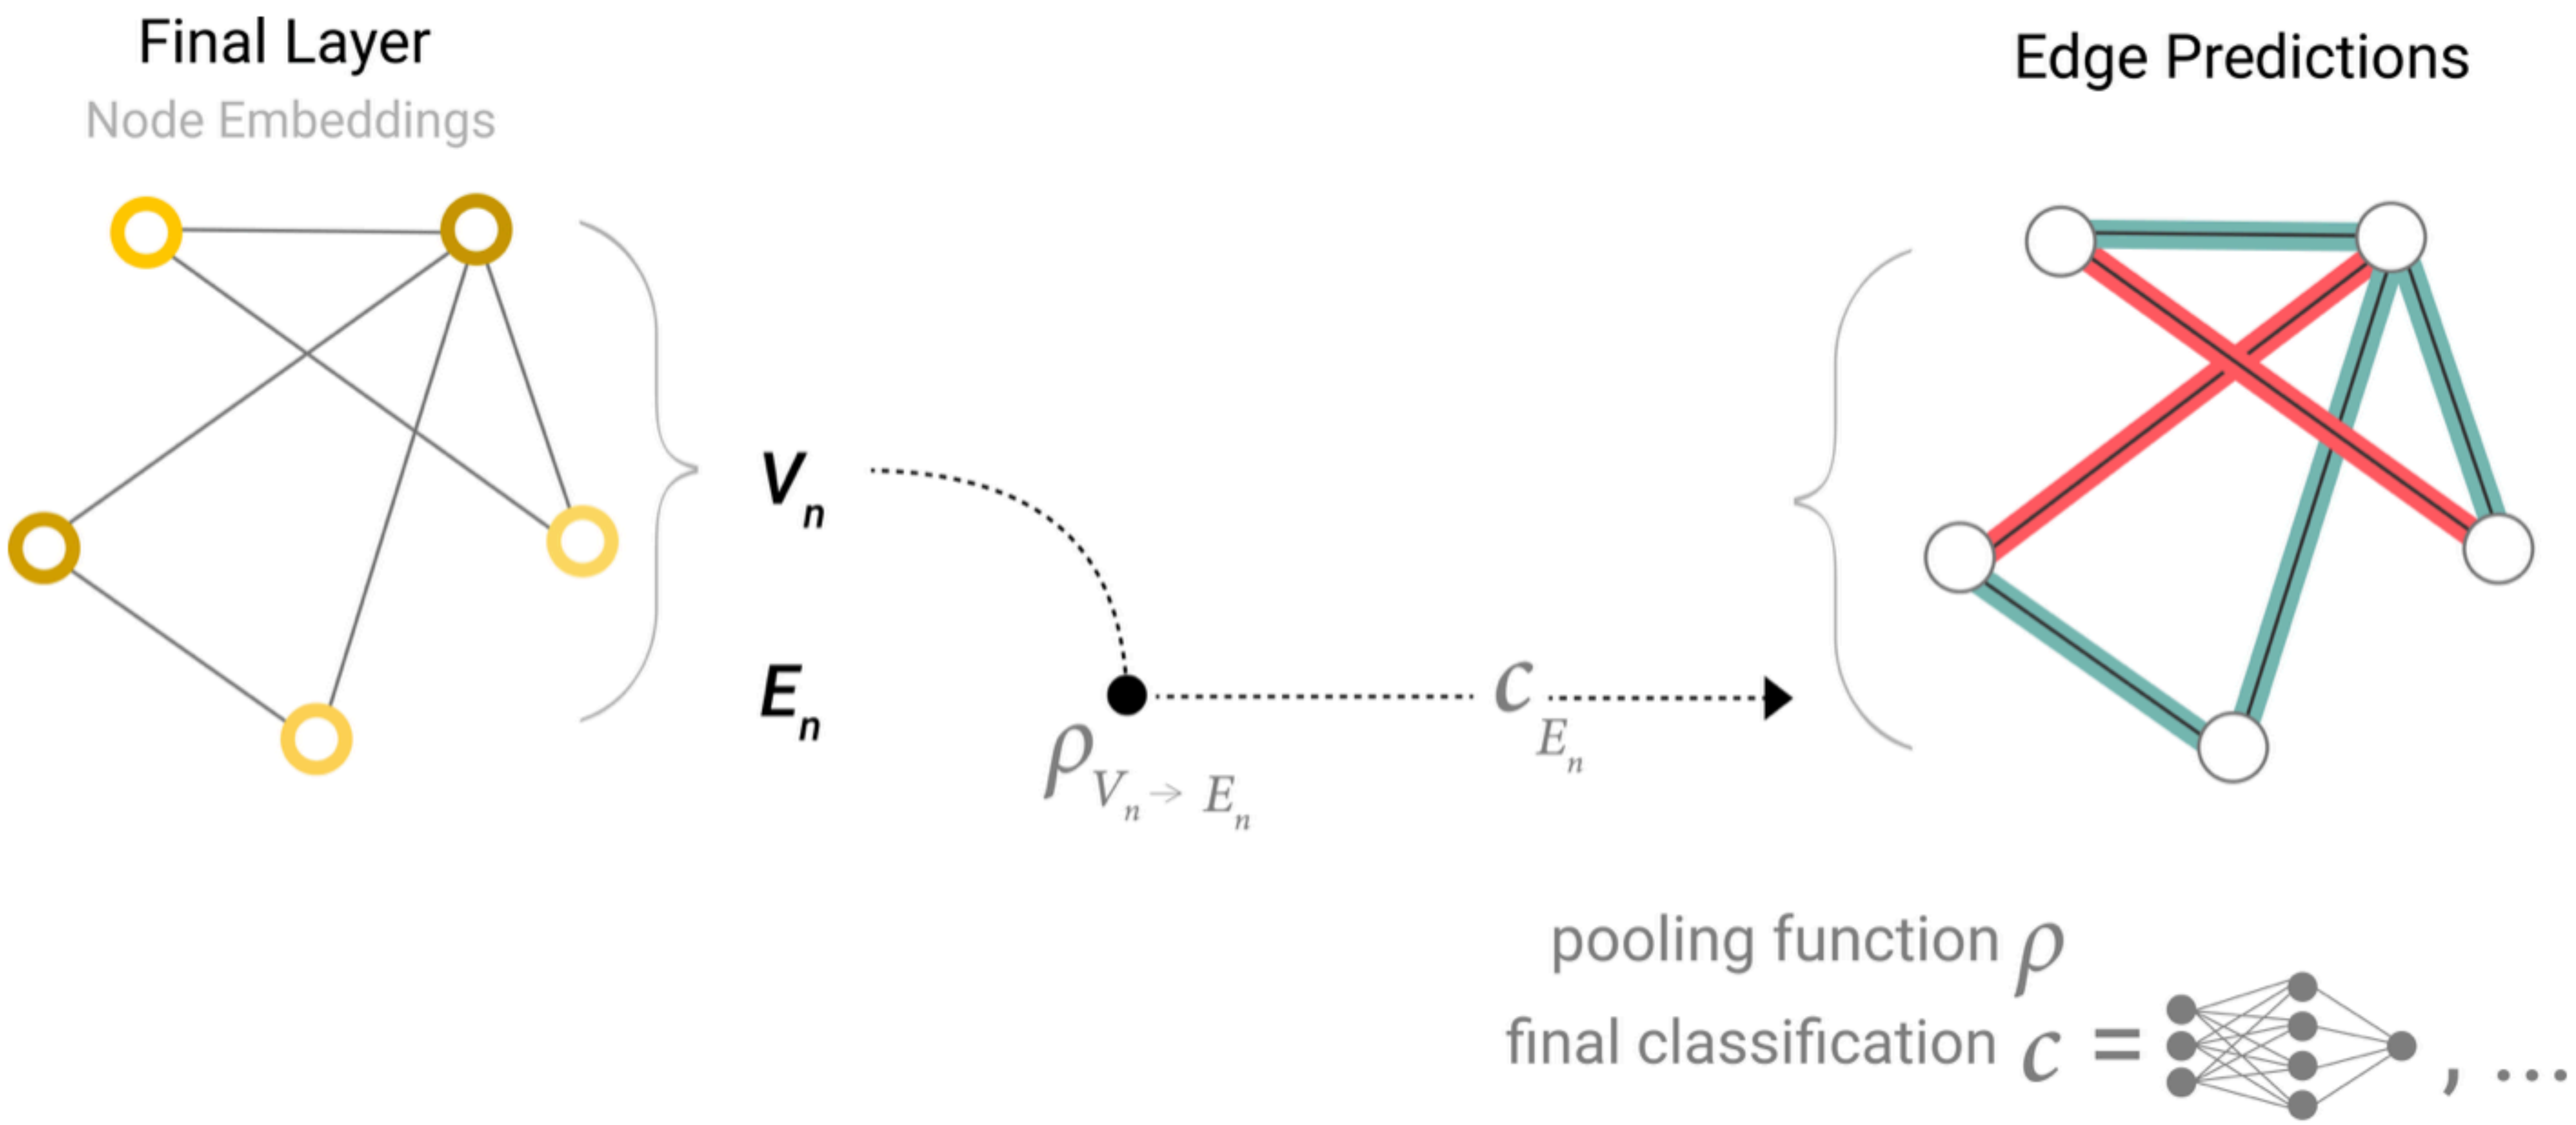

Predictions by Pooling Information

- If we only have edge/node-level features and need to predict a global property, need to gather all edge/node information and aggregate them

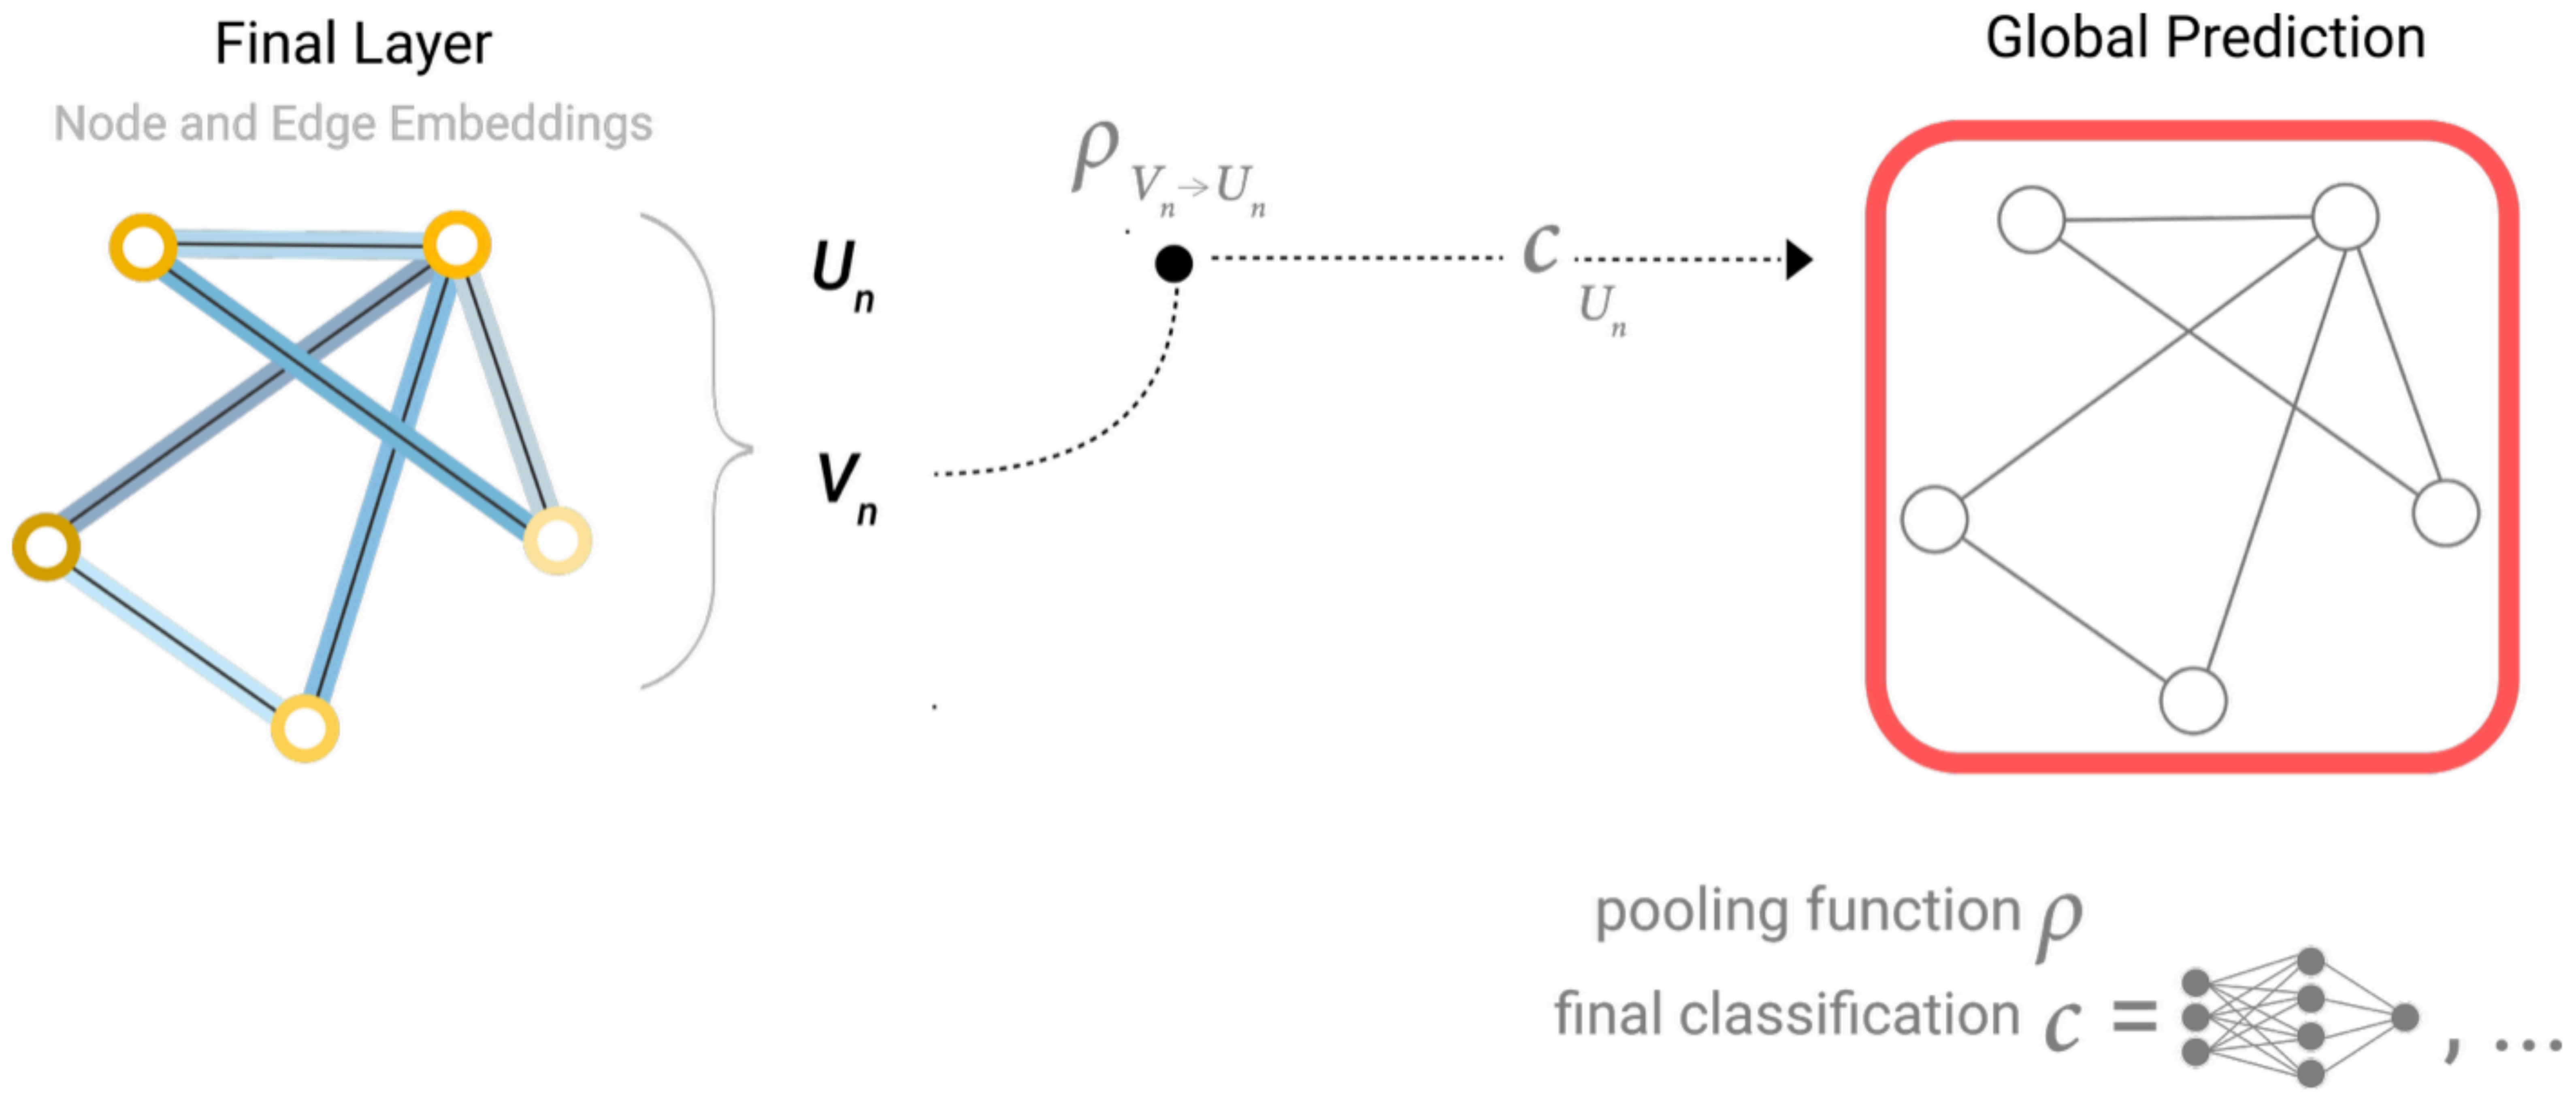

Prediction Review

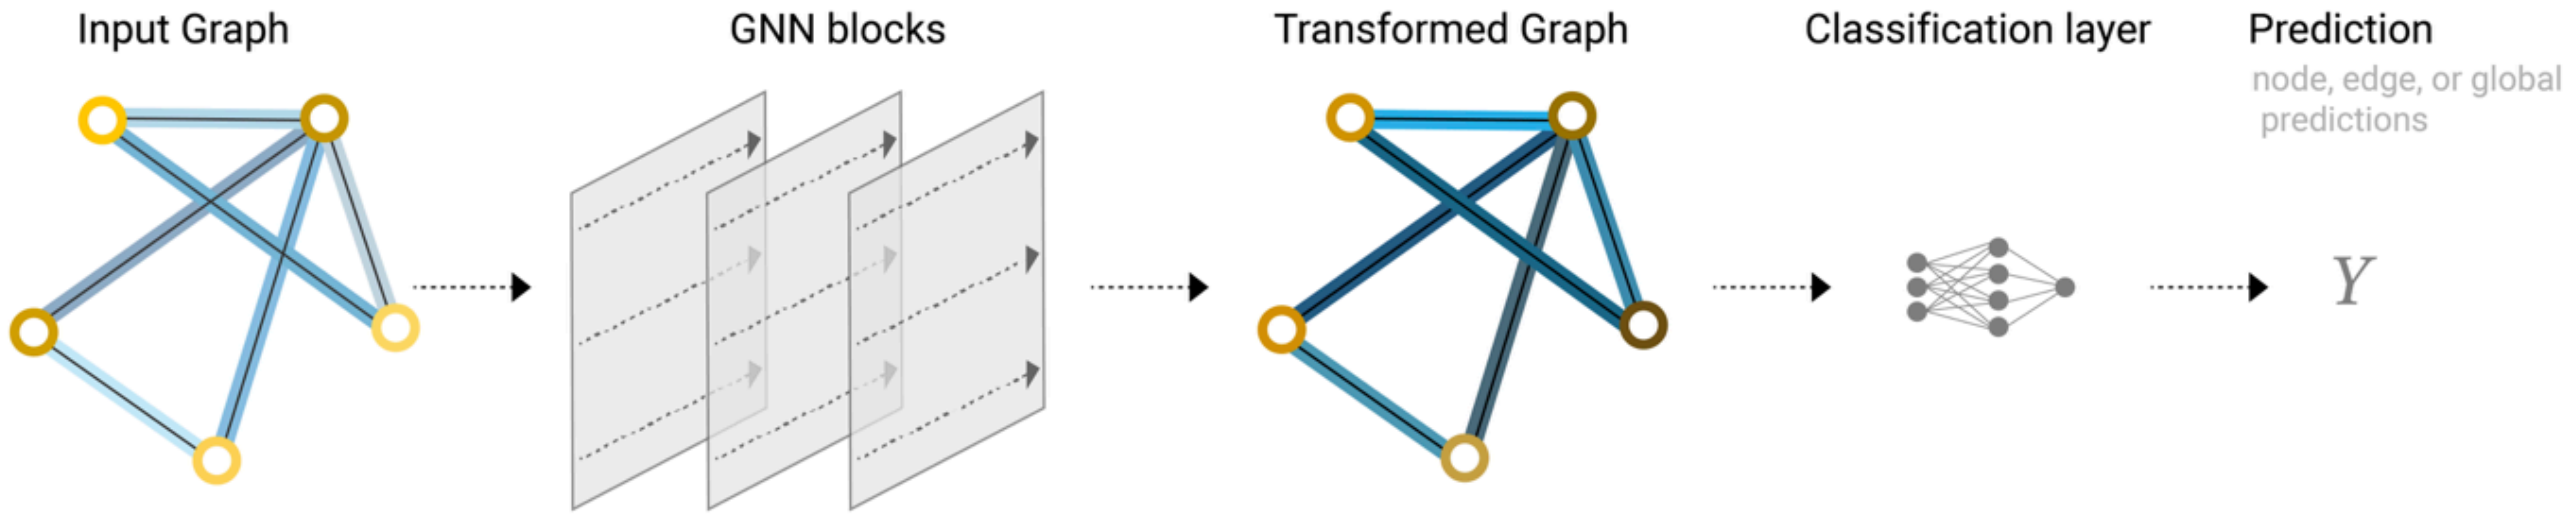

An end-to-end prediction task with a GNN model

Conceptual Application

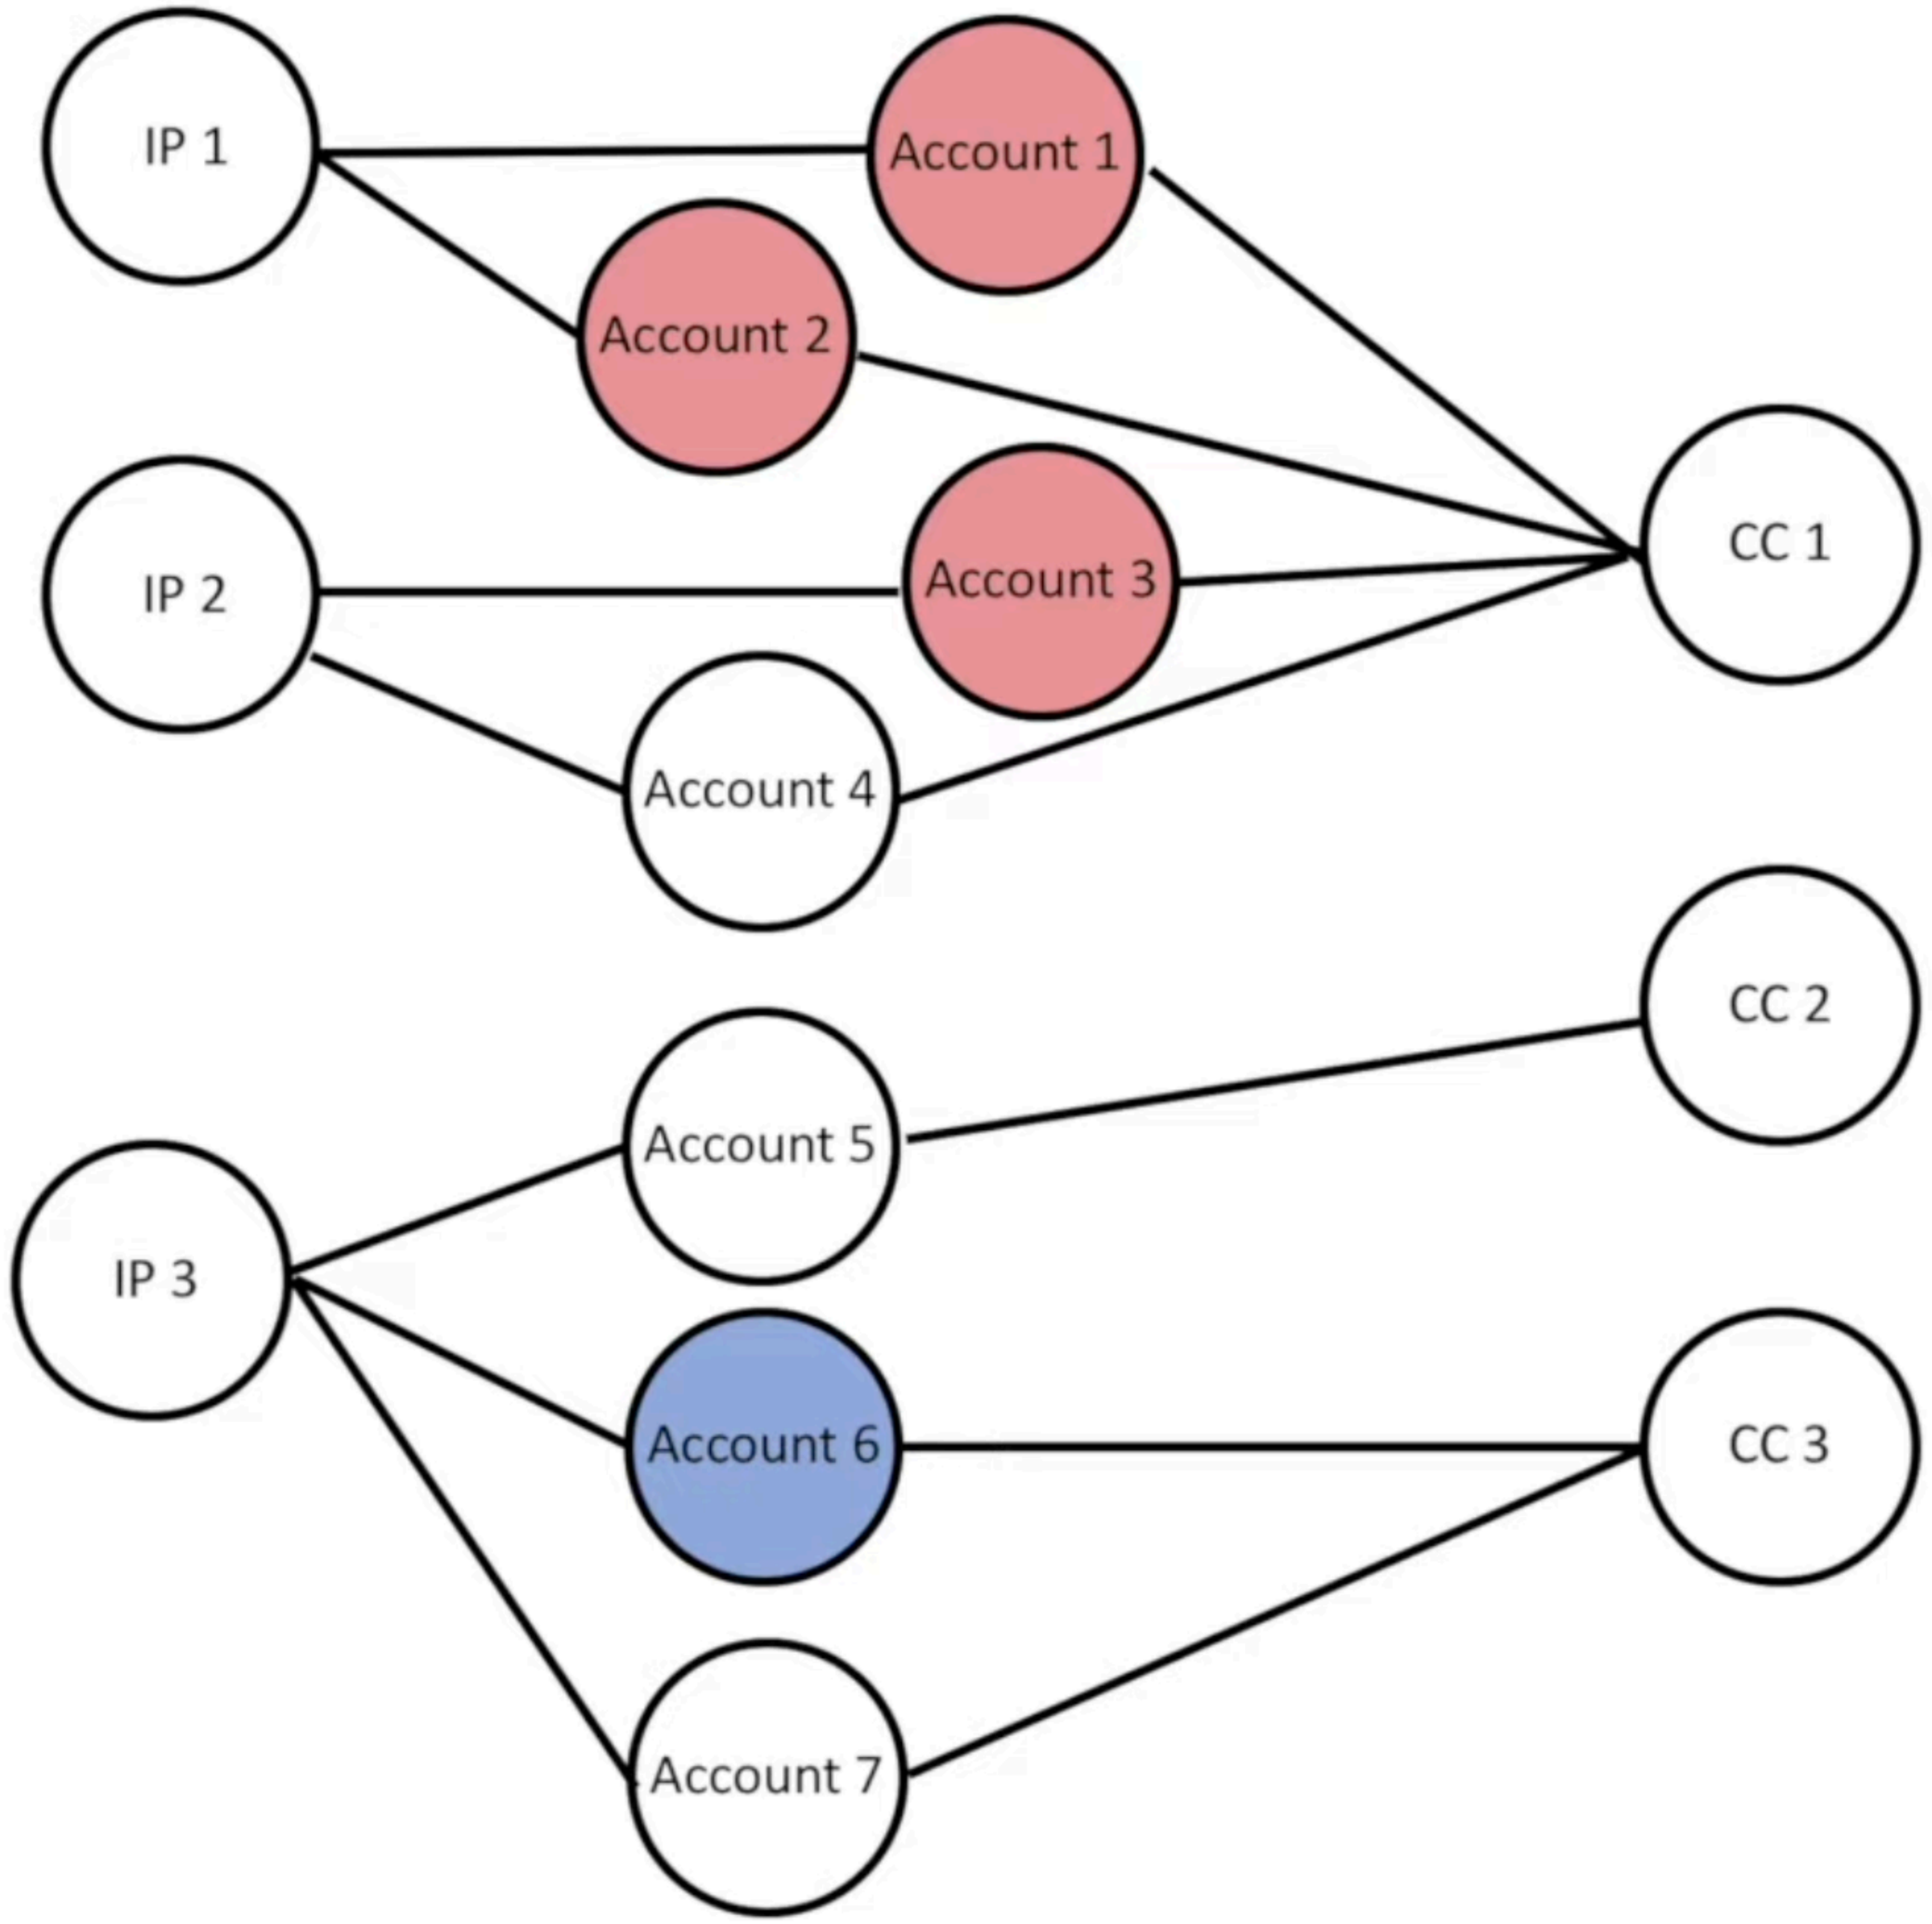

Conceptual Application

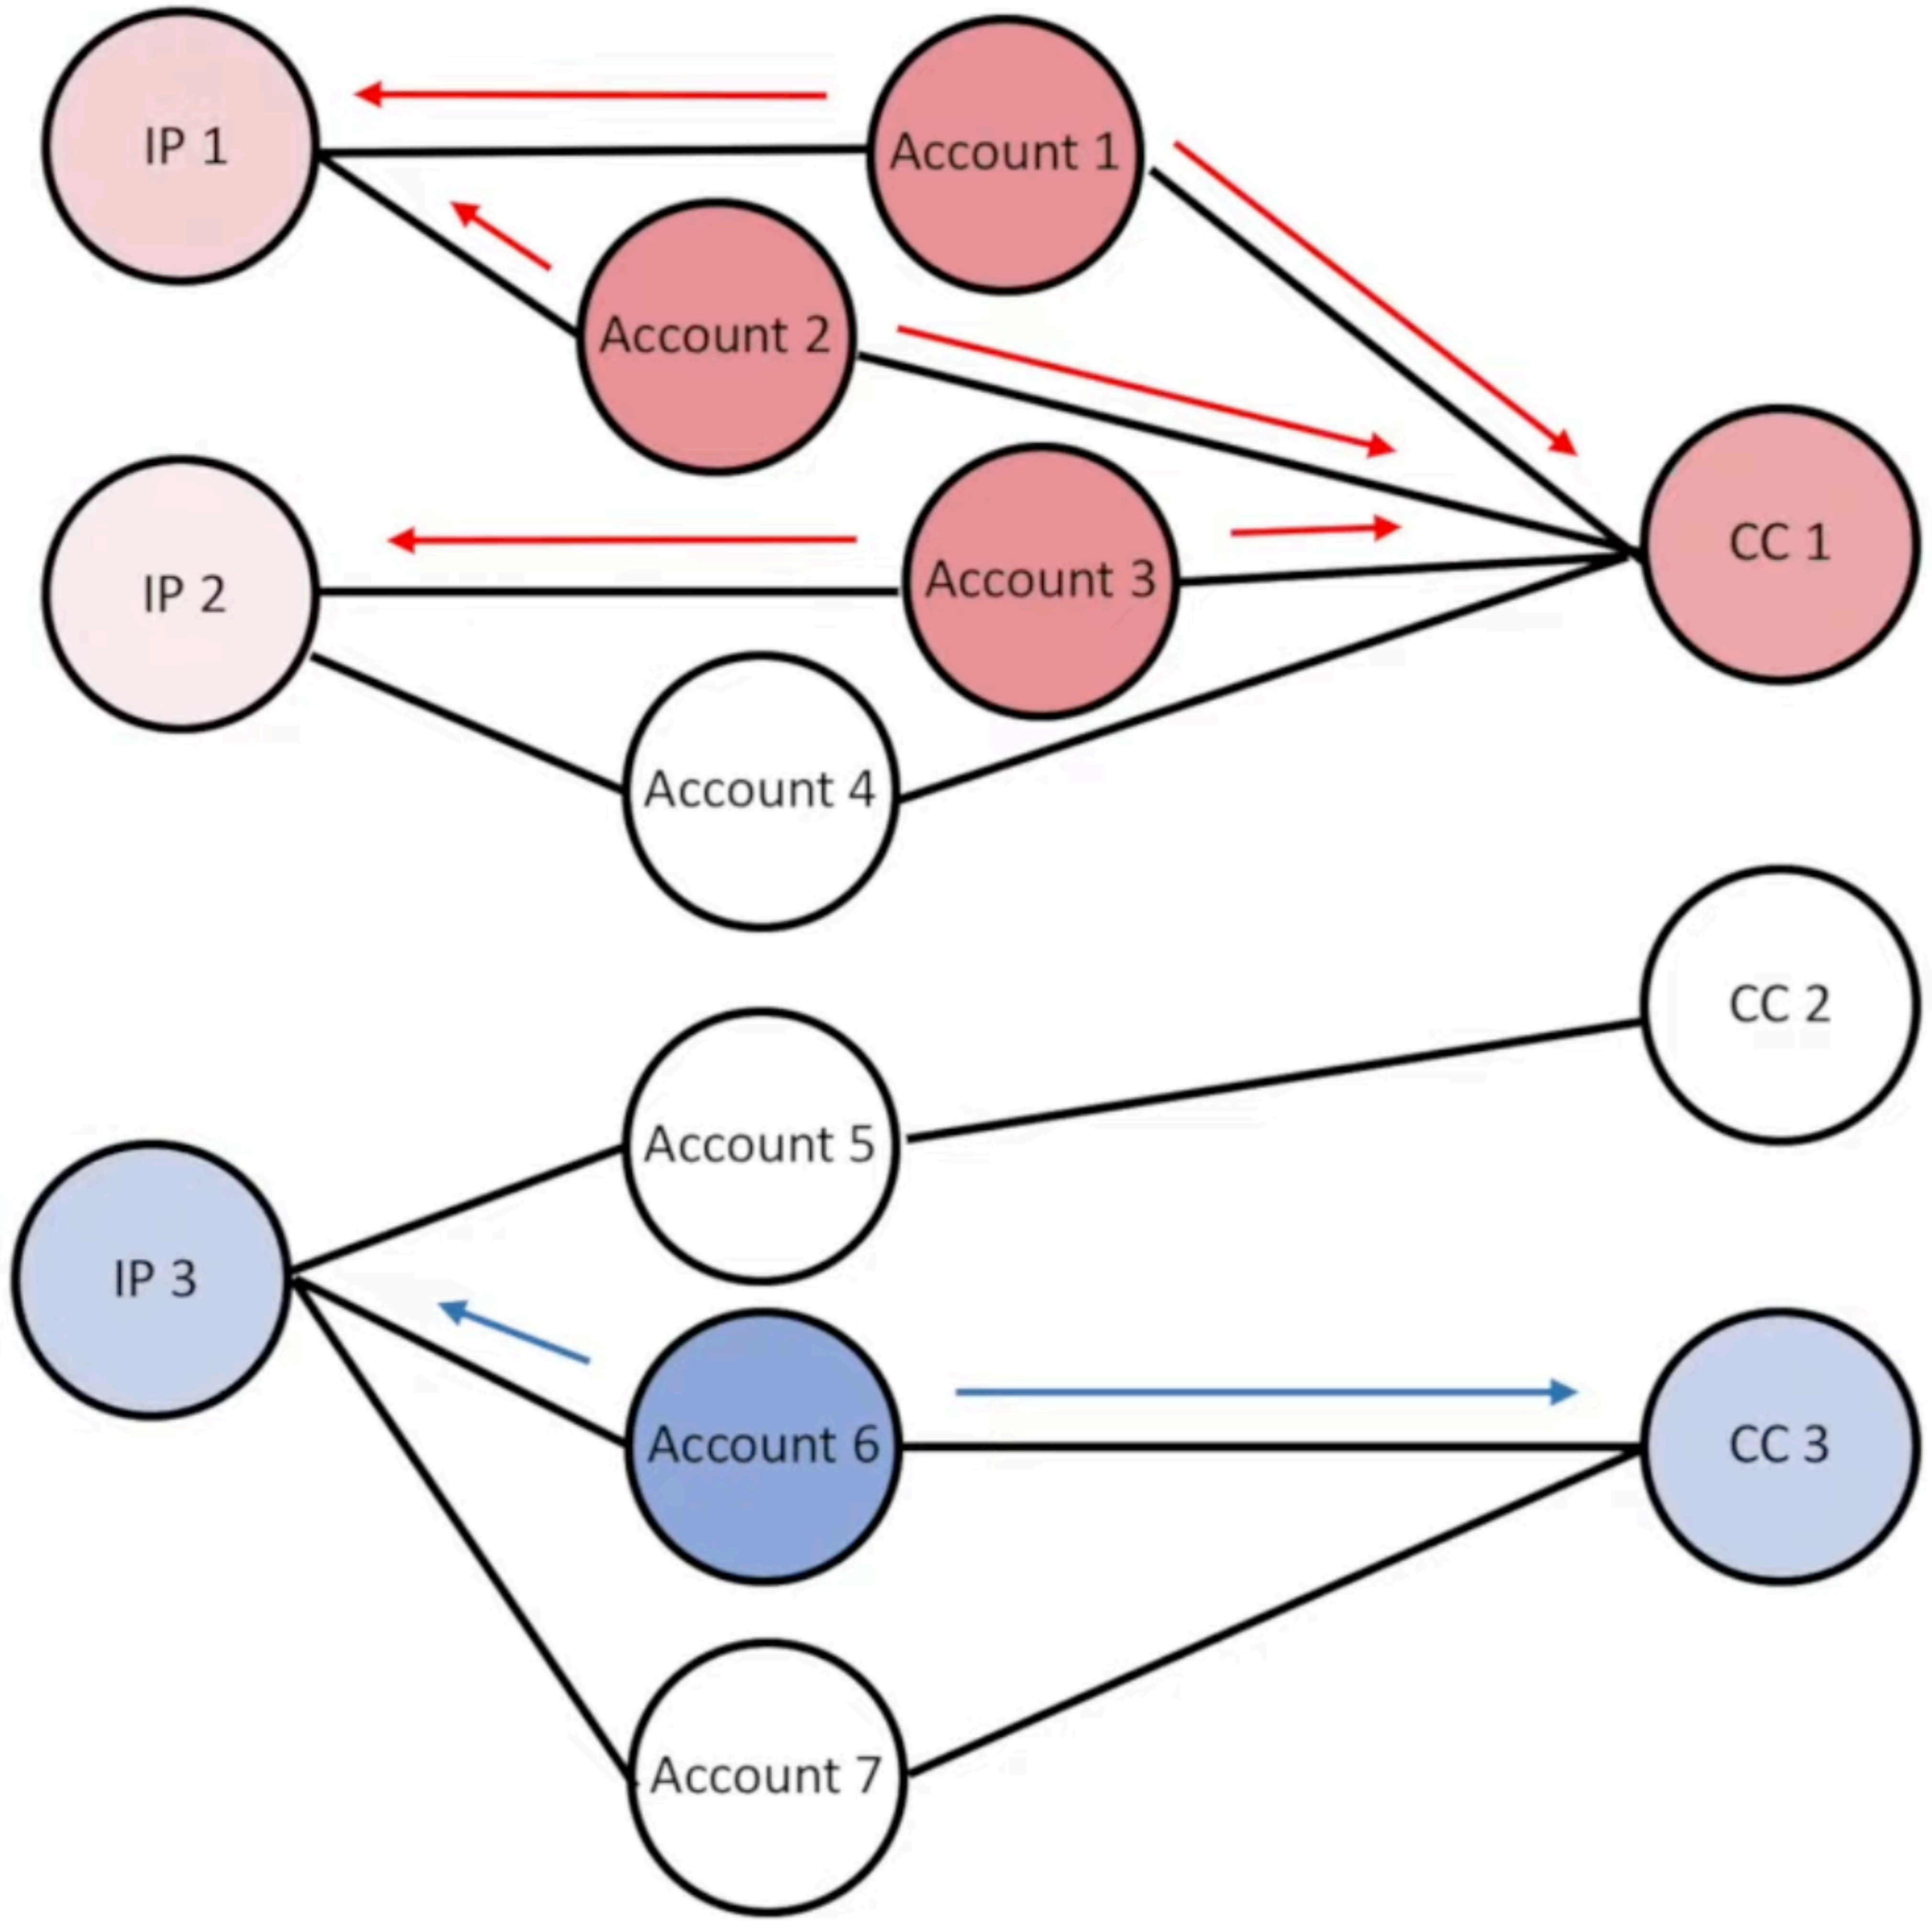

Conceptual Application

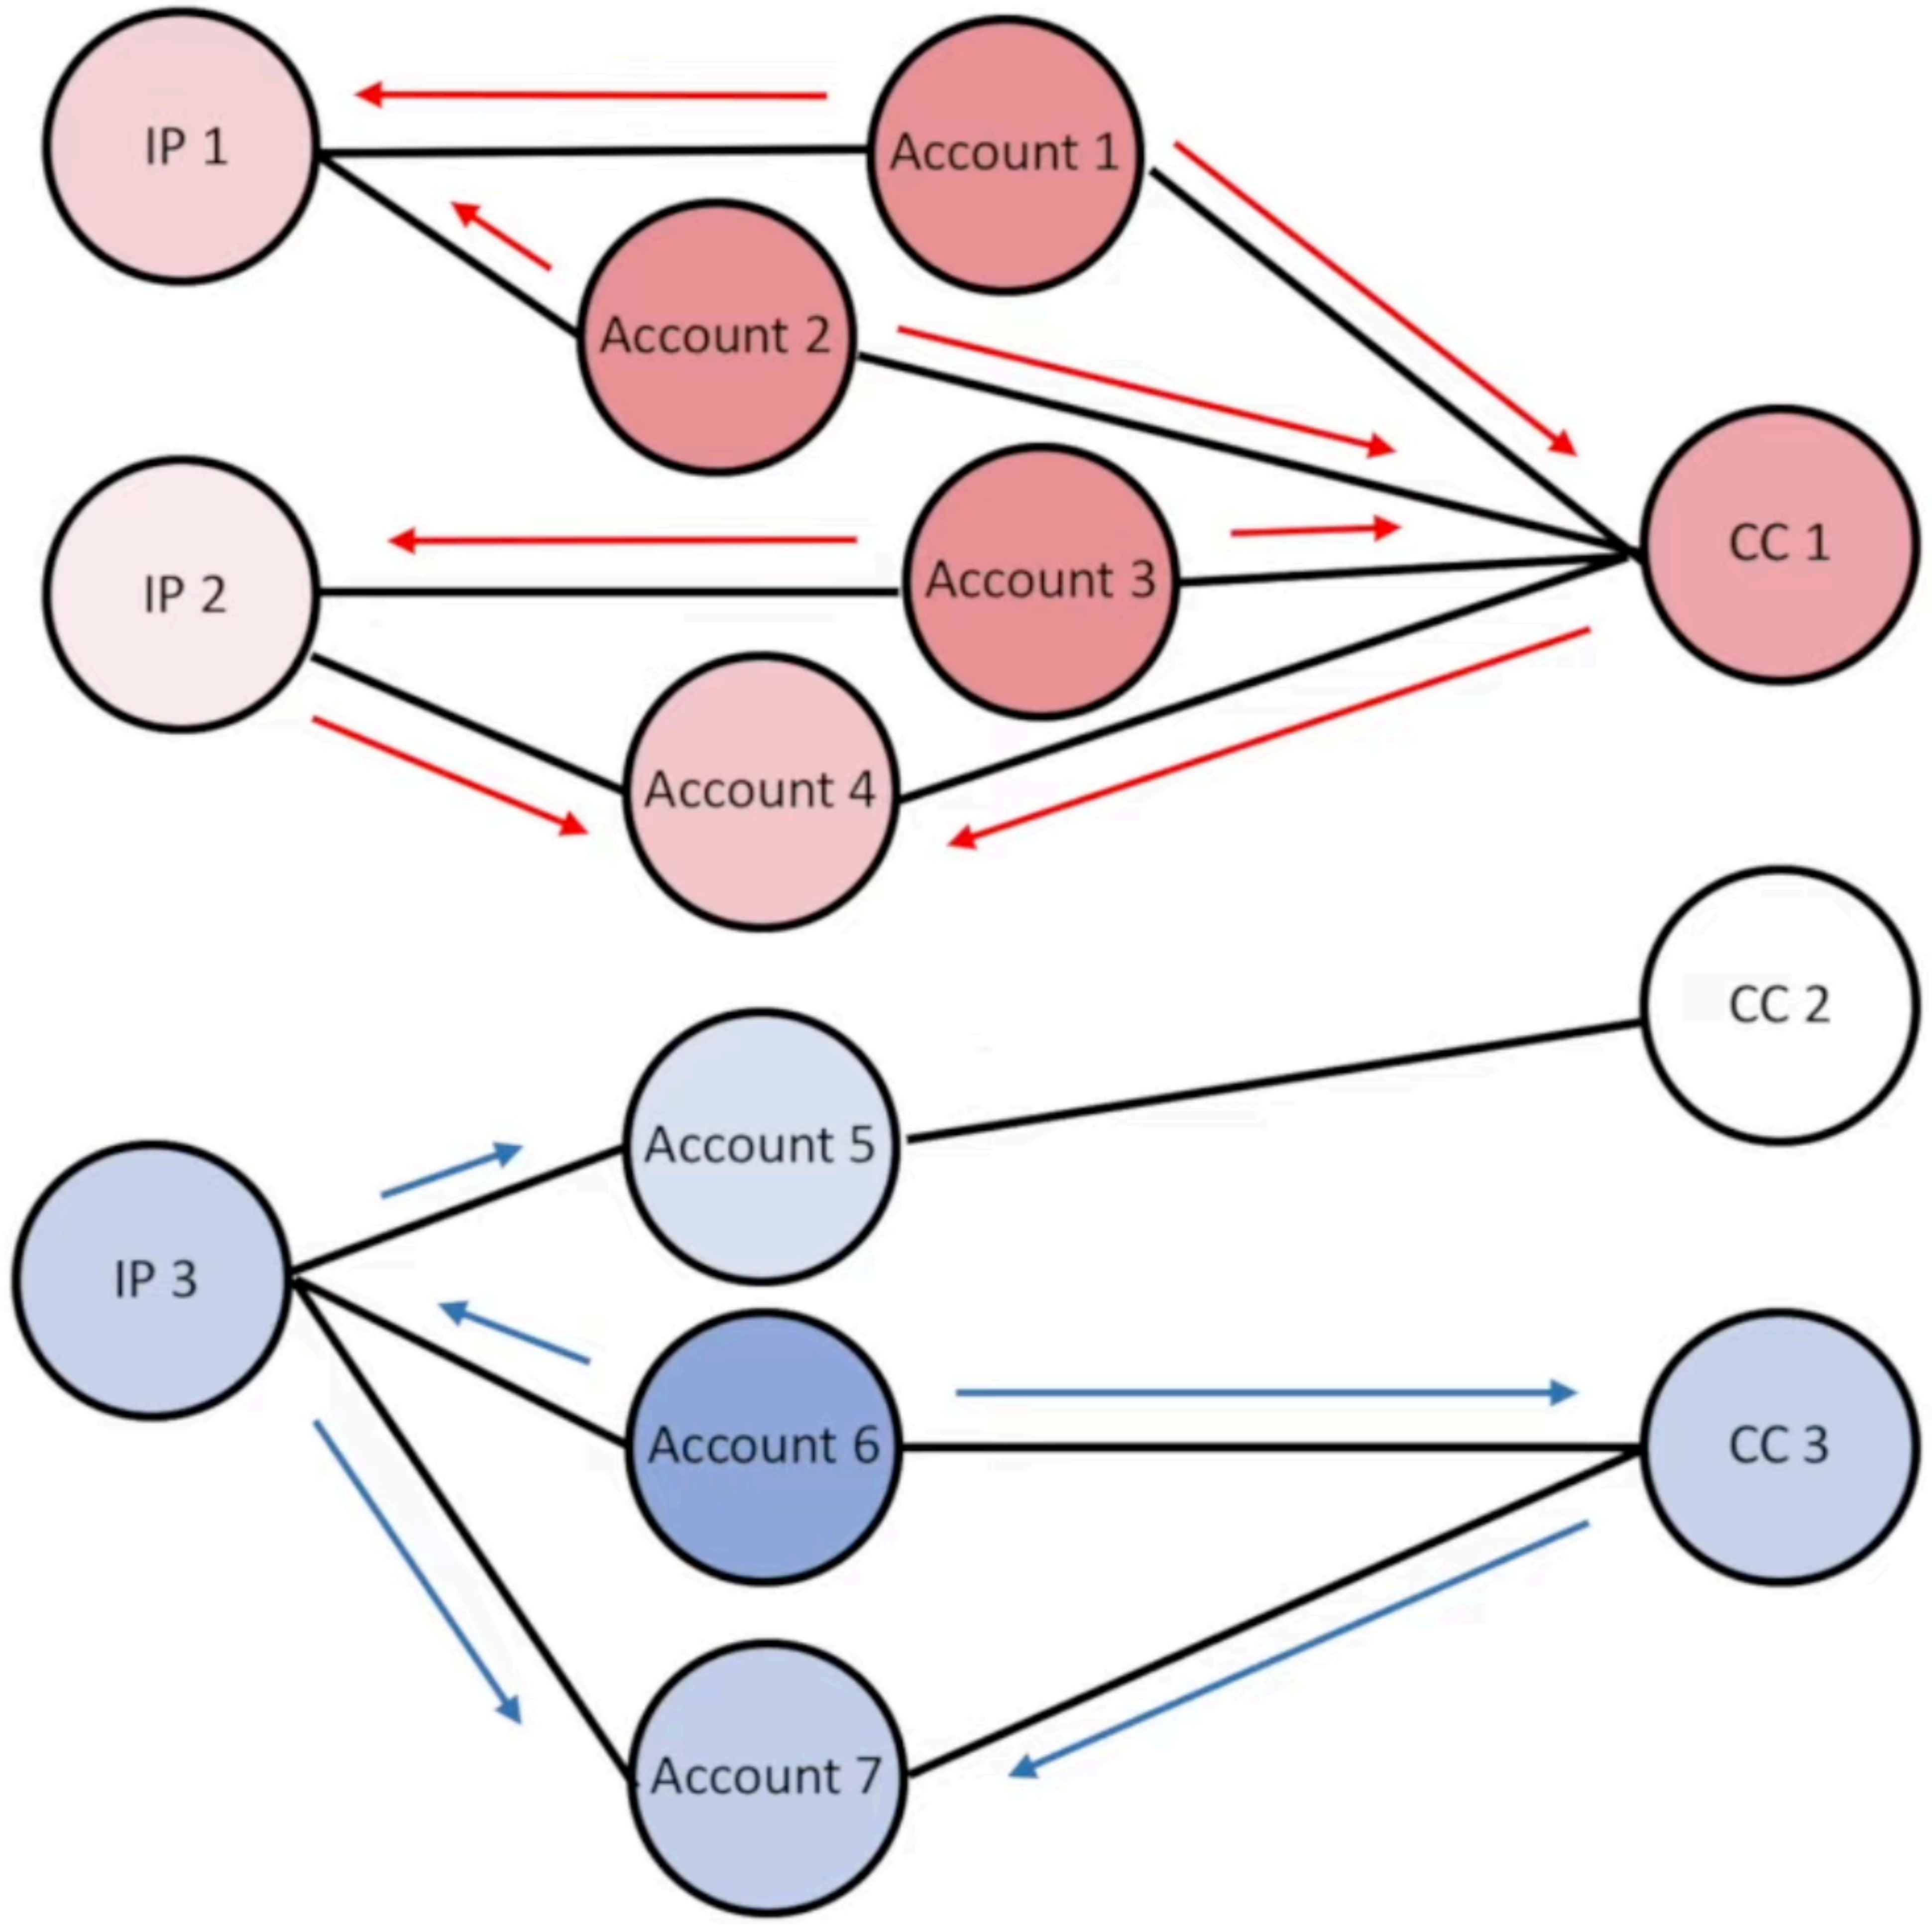

Conceptual Application

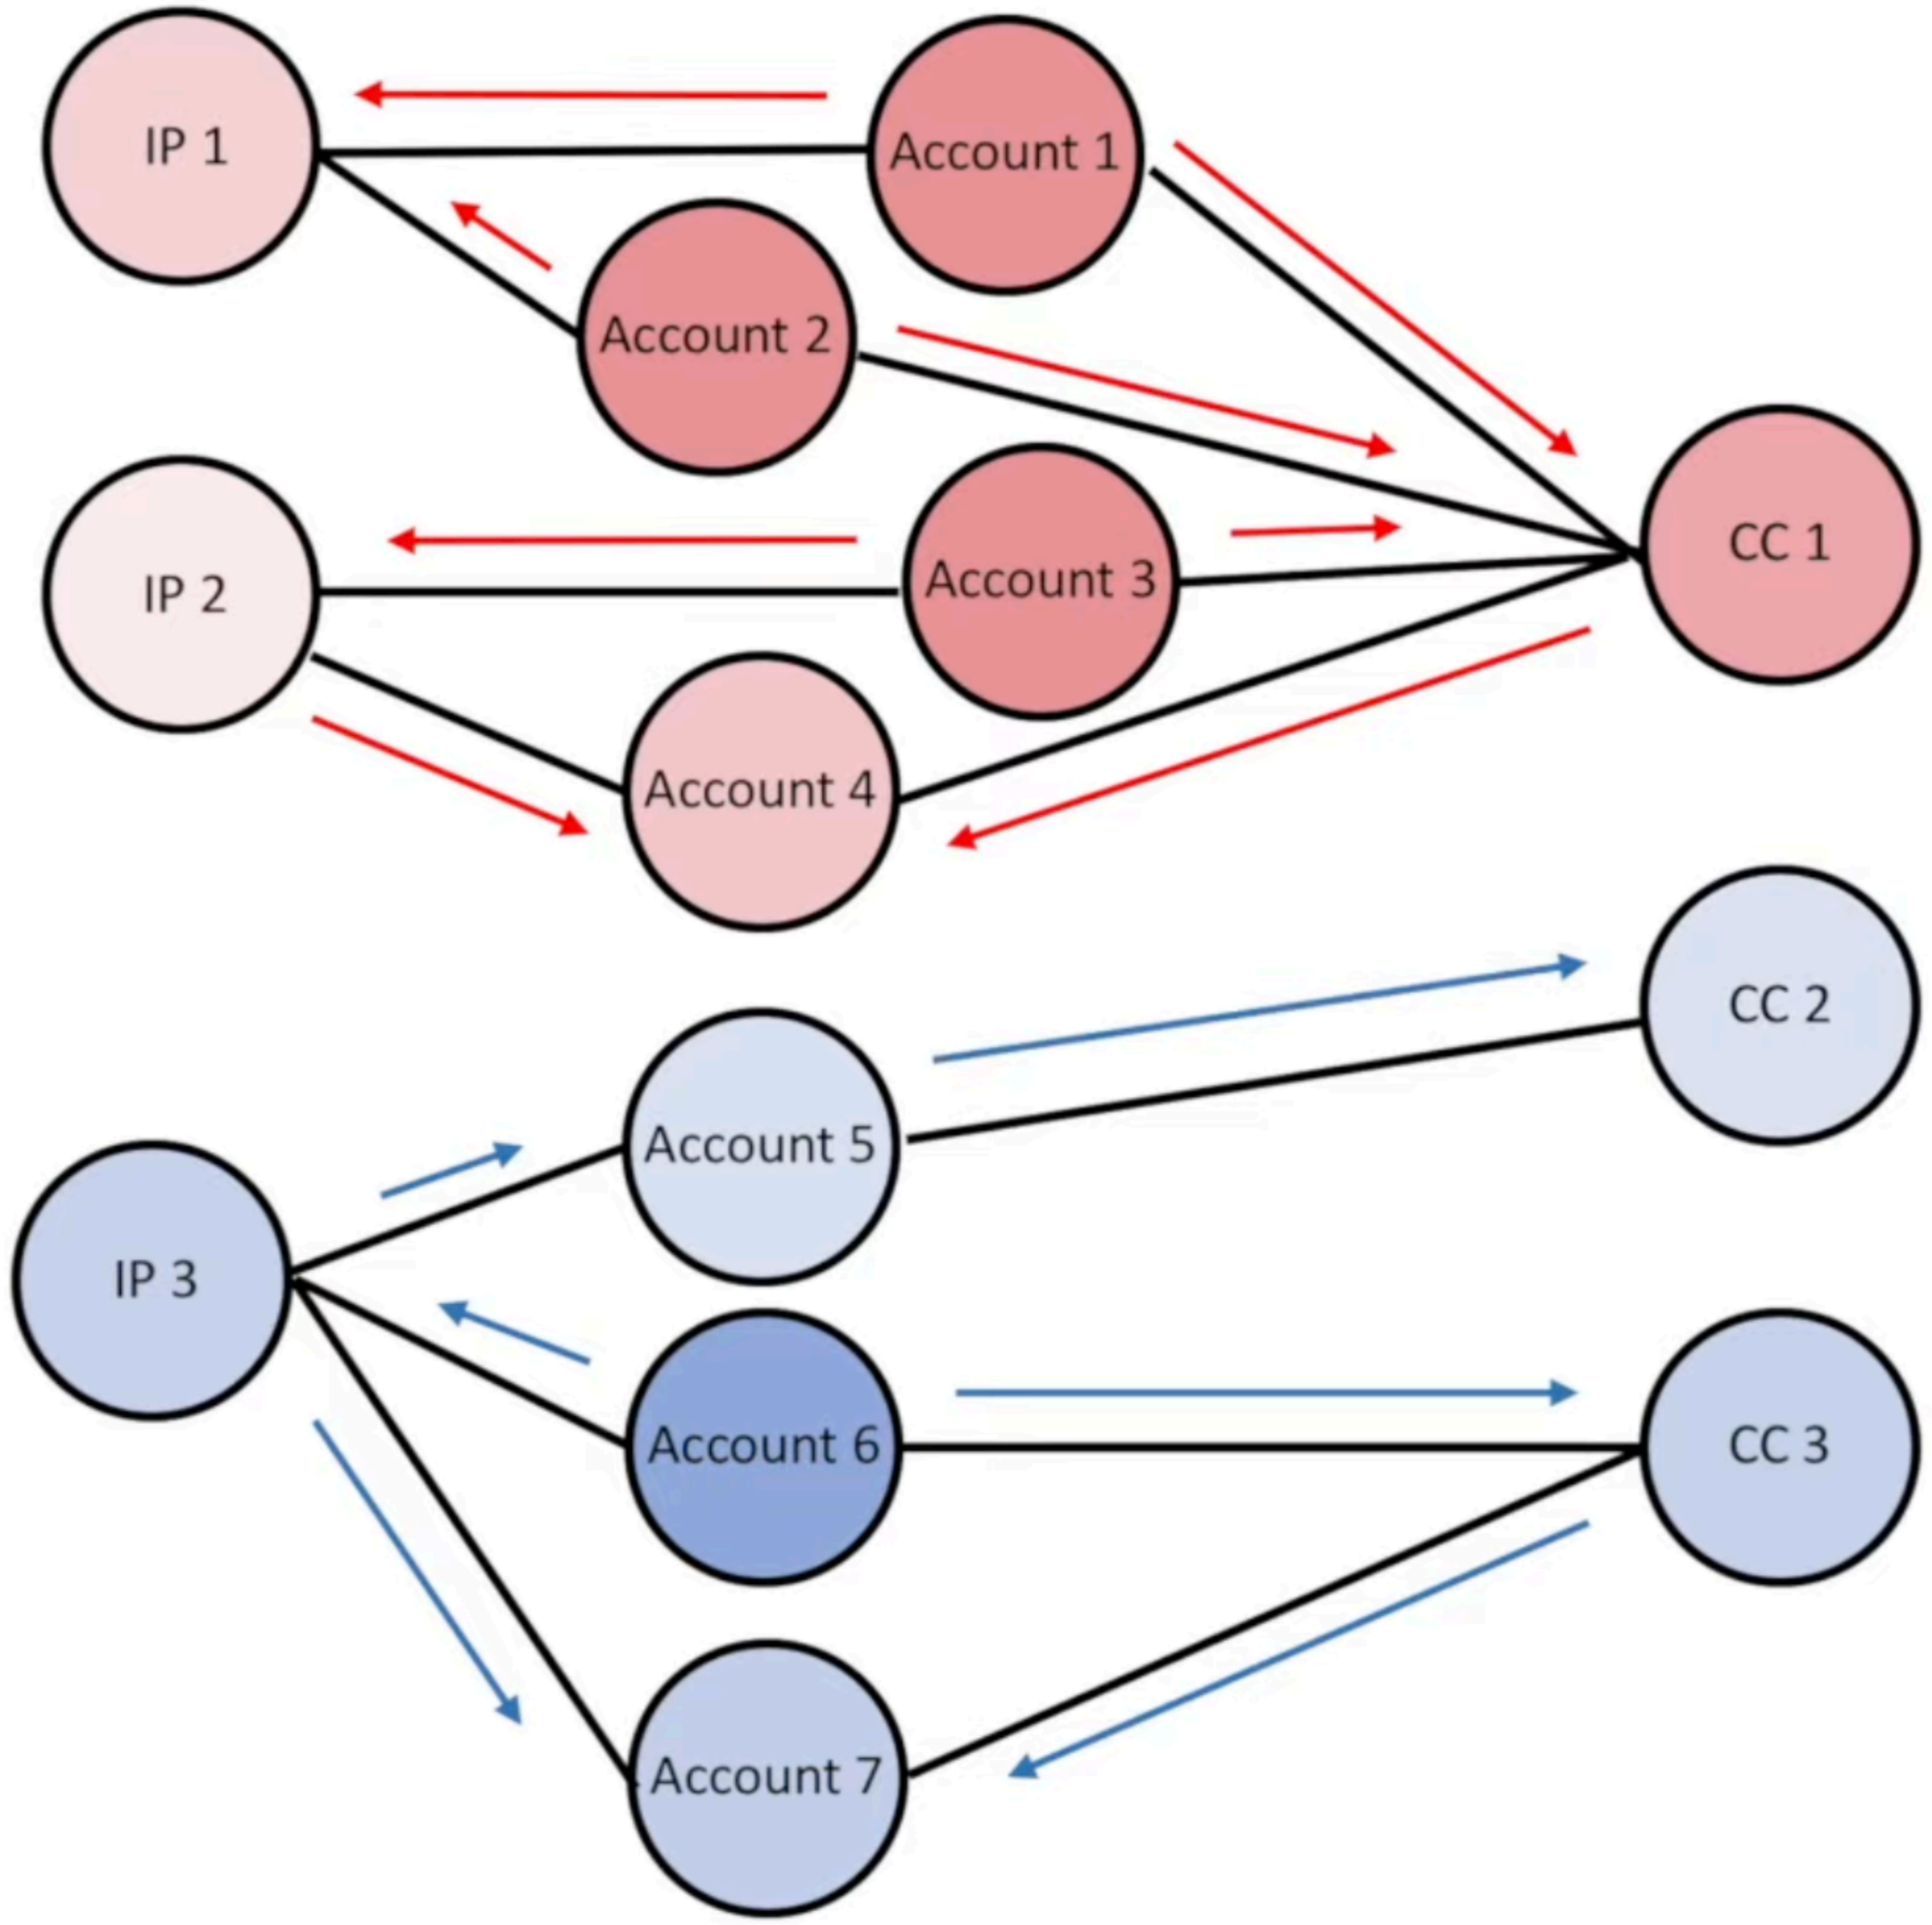

Simple Calculation in GNN Layer

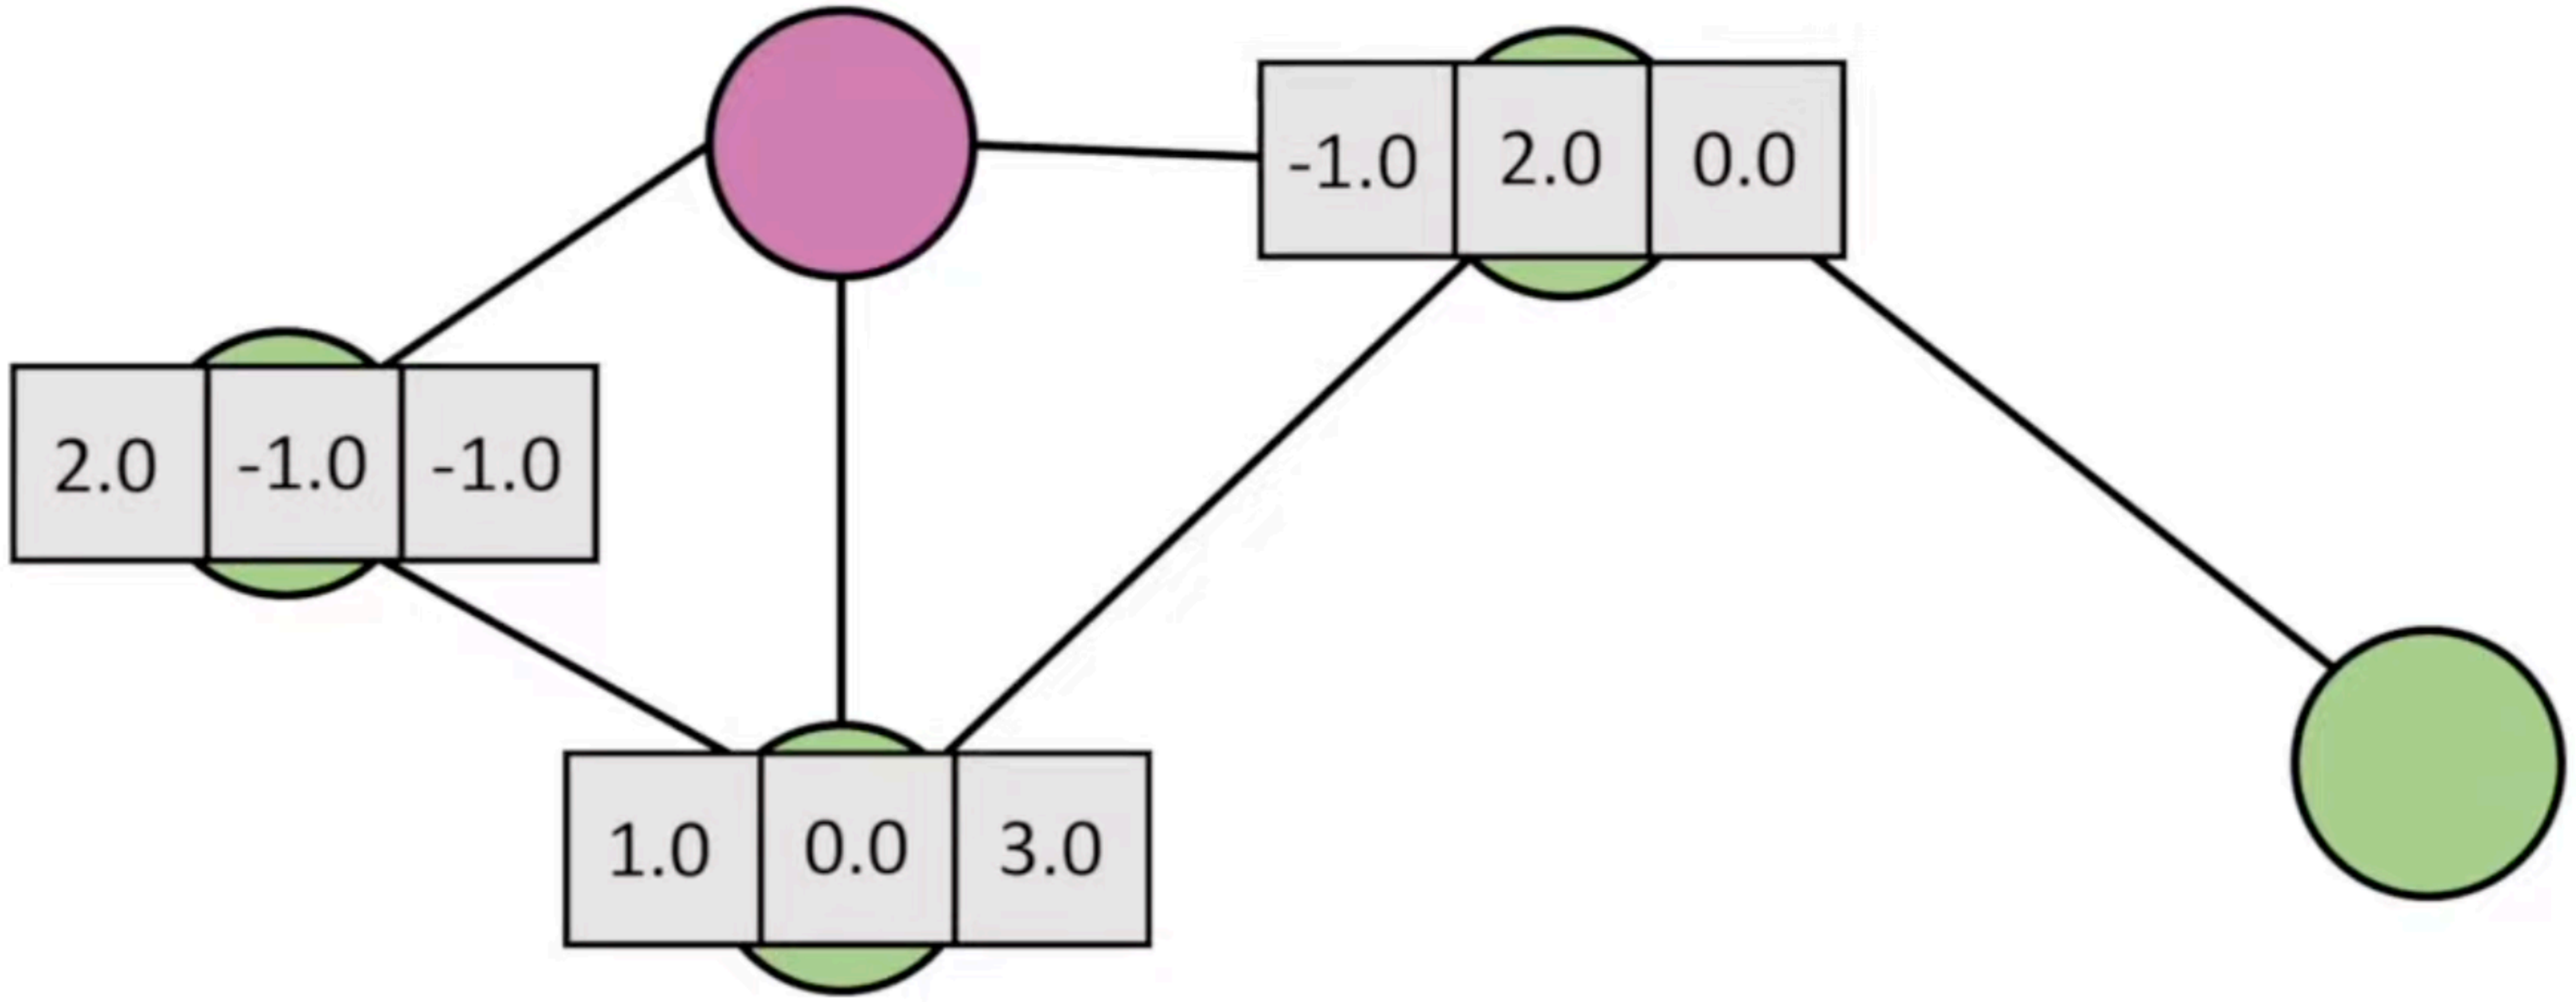

Simple Calculation in GNN Layer

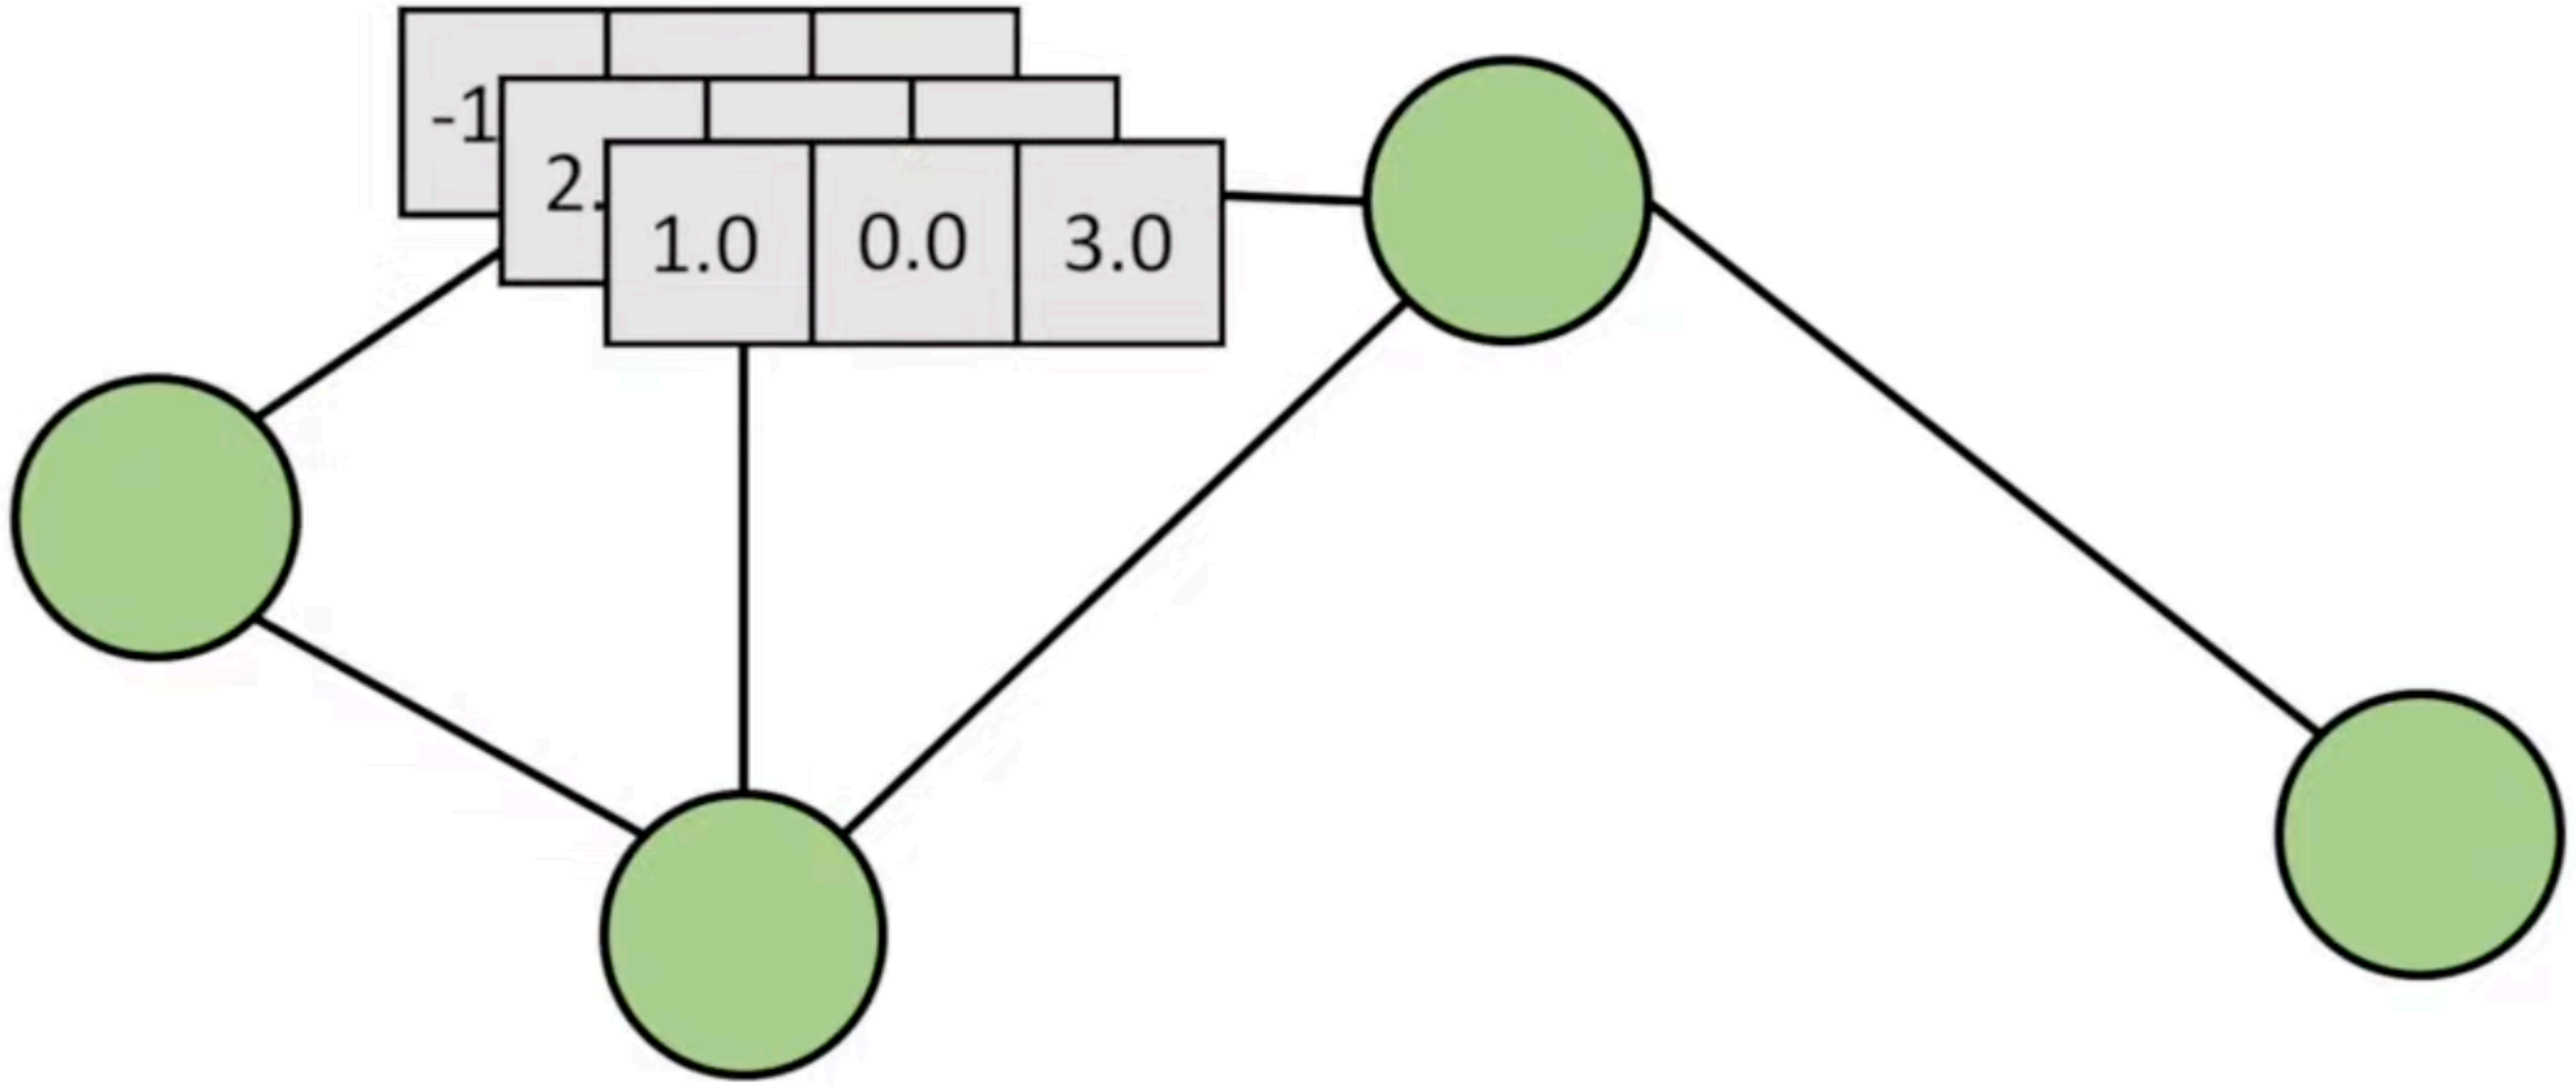

Simple Calculation in GNN Layer

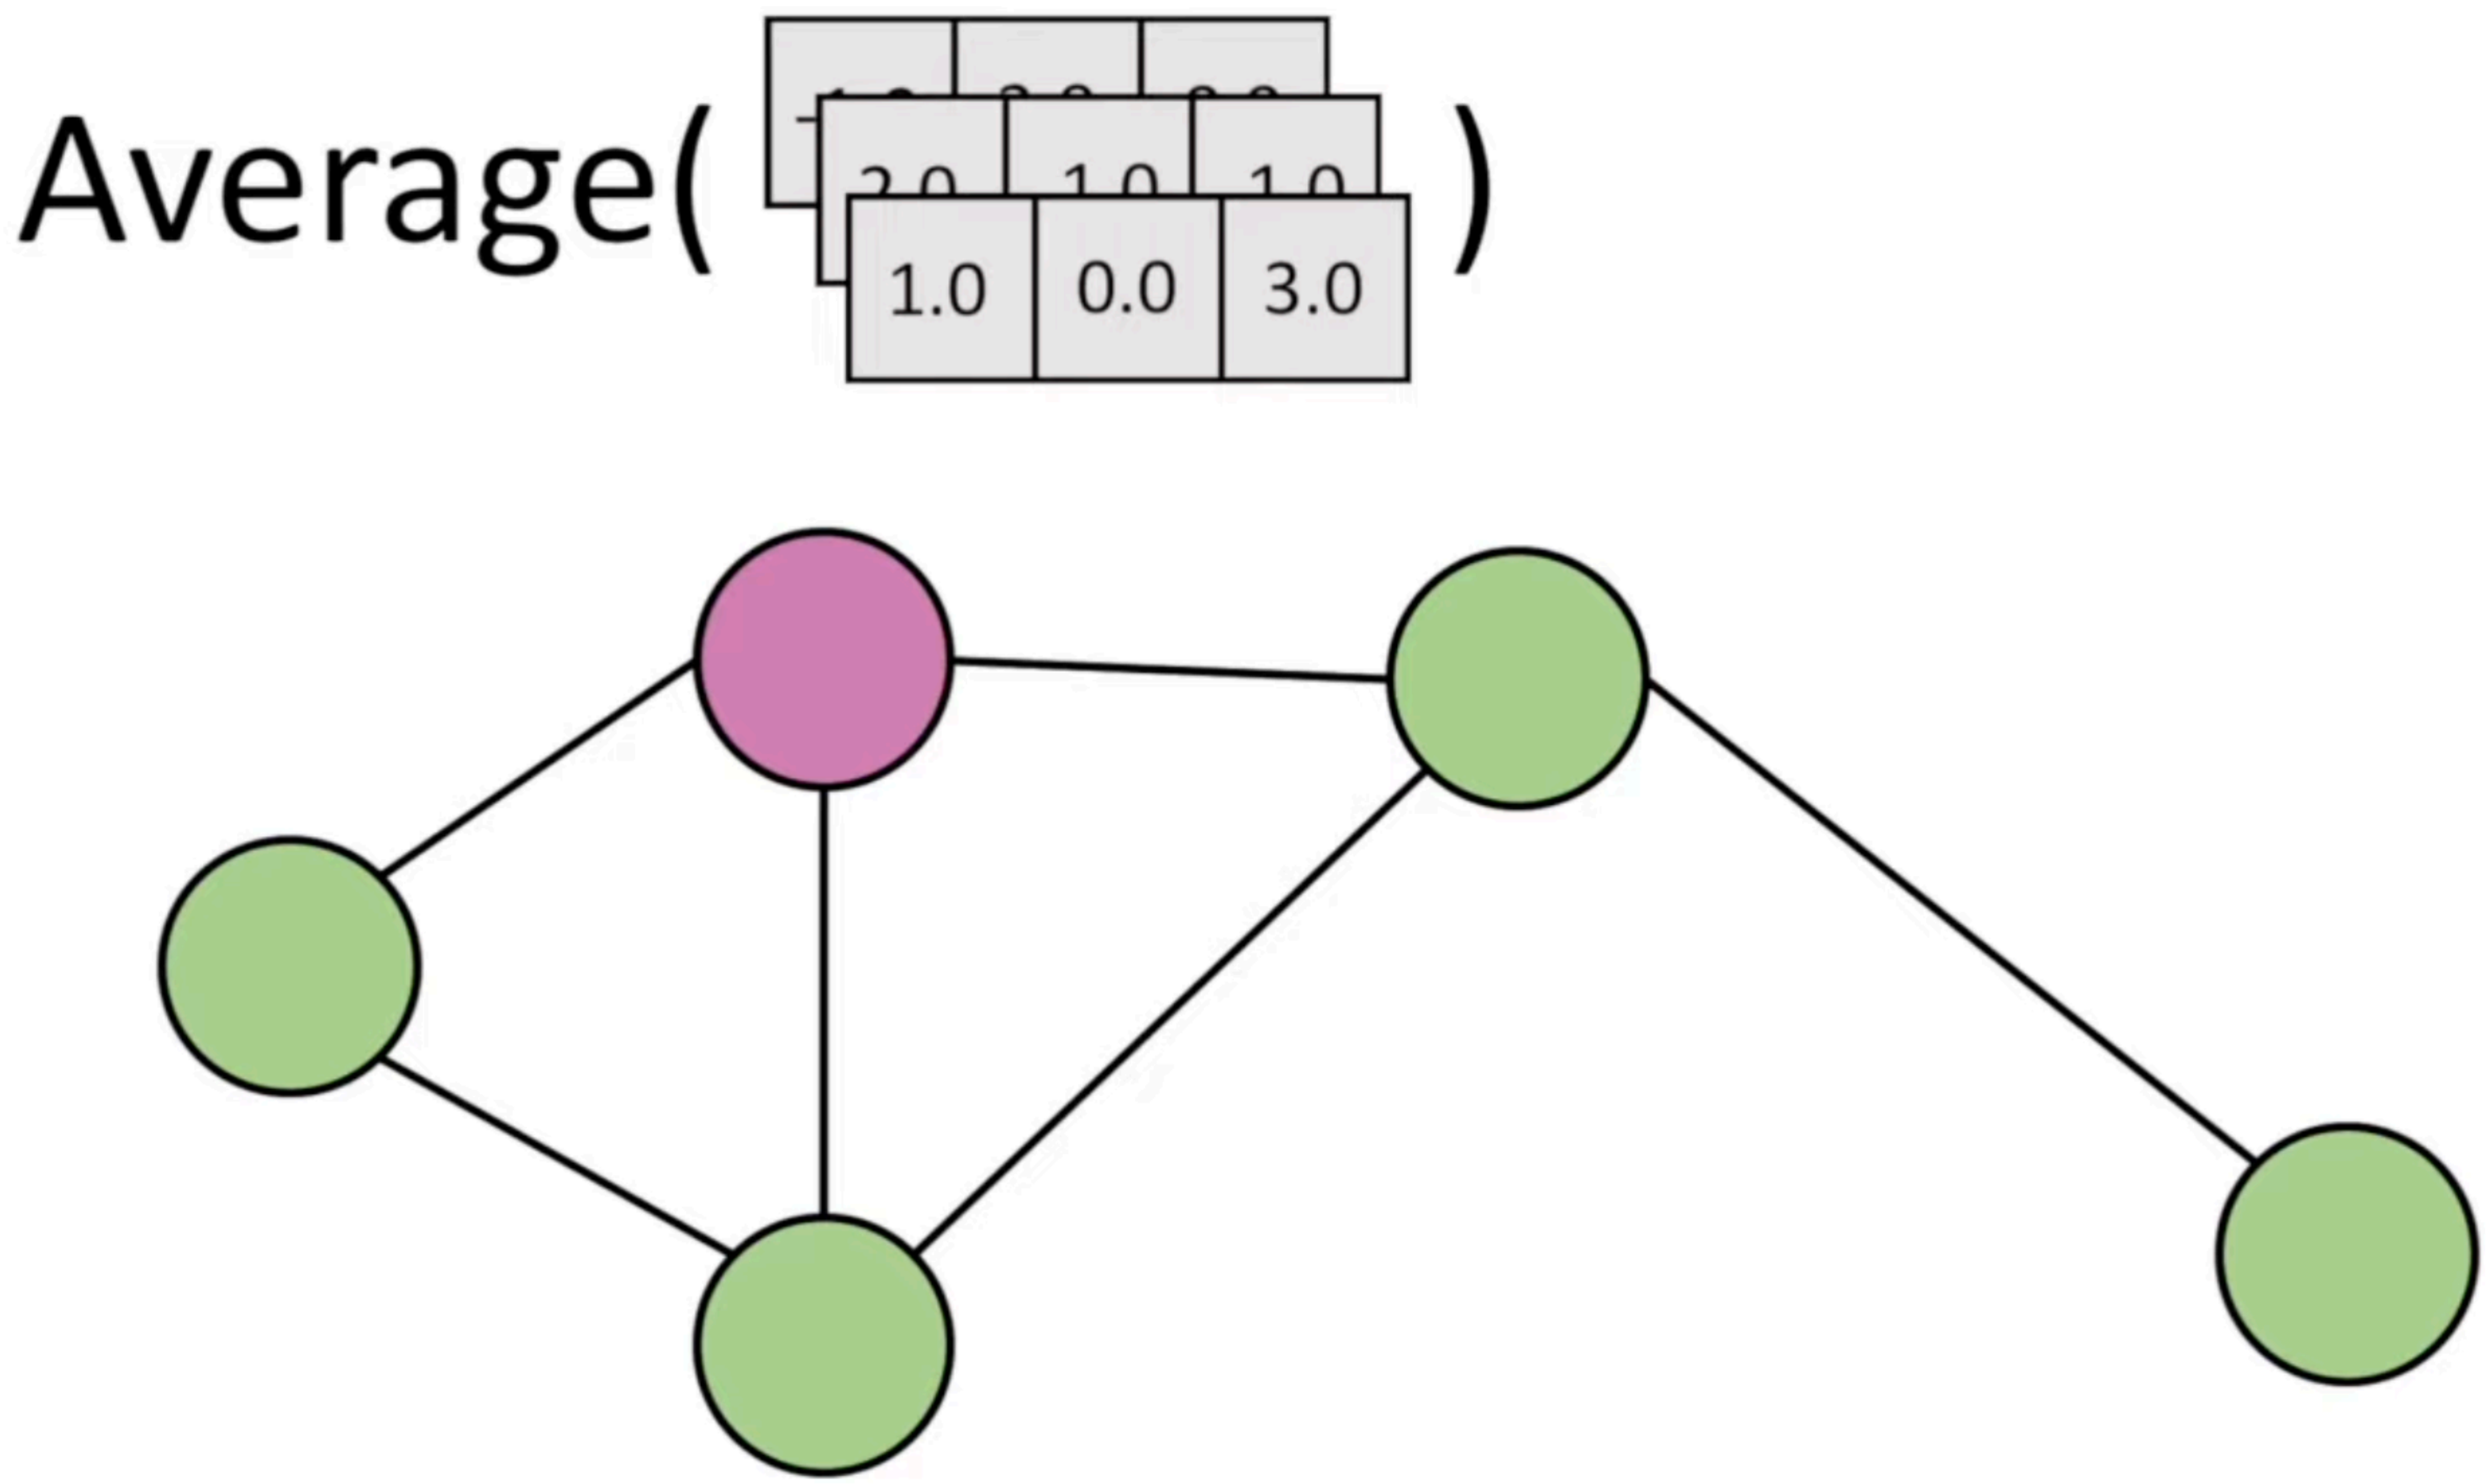

Simple Calculation in GNN Layer

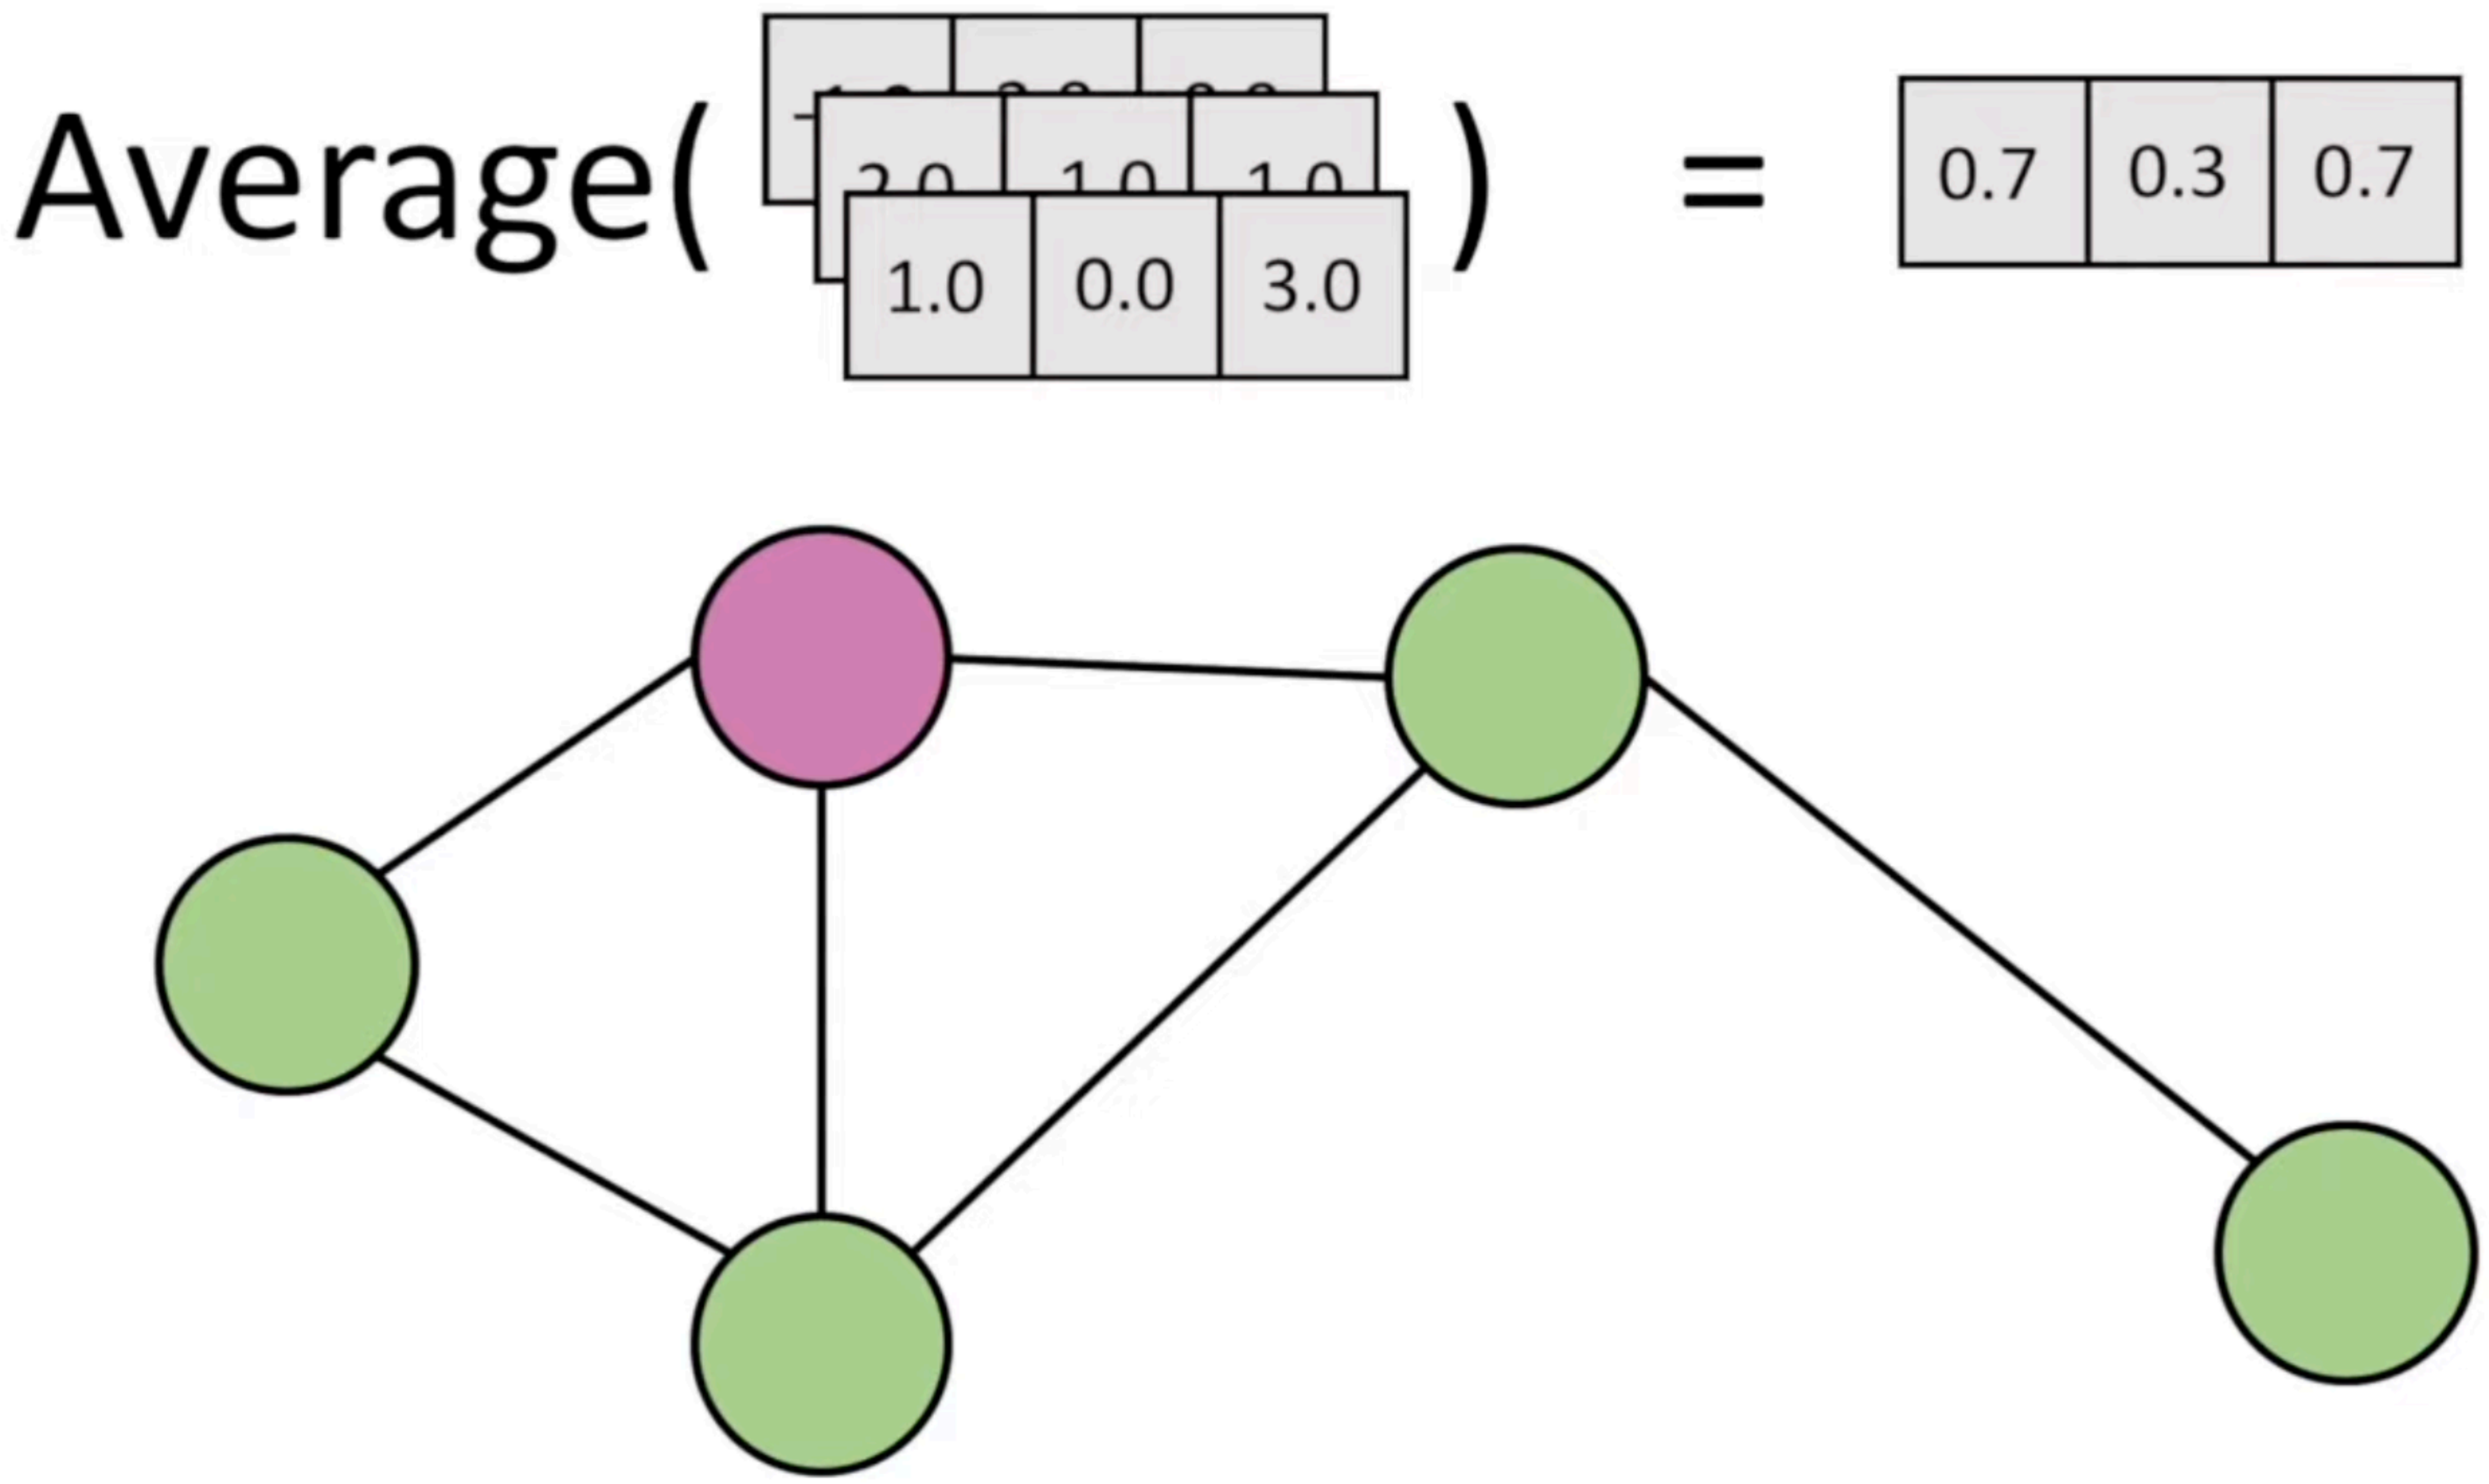

Simple Calculation in GNN Layer

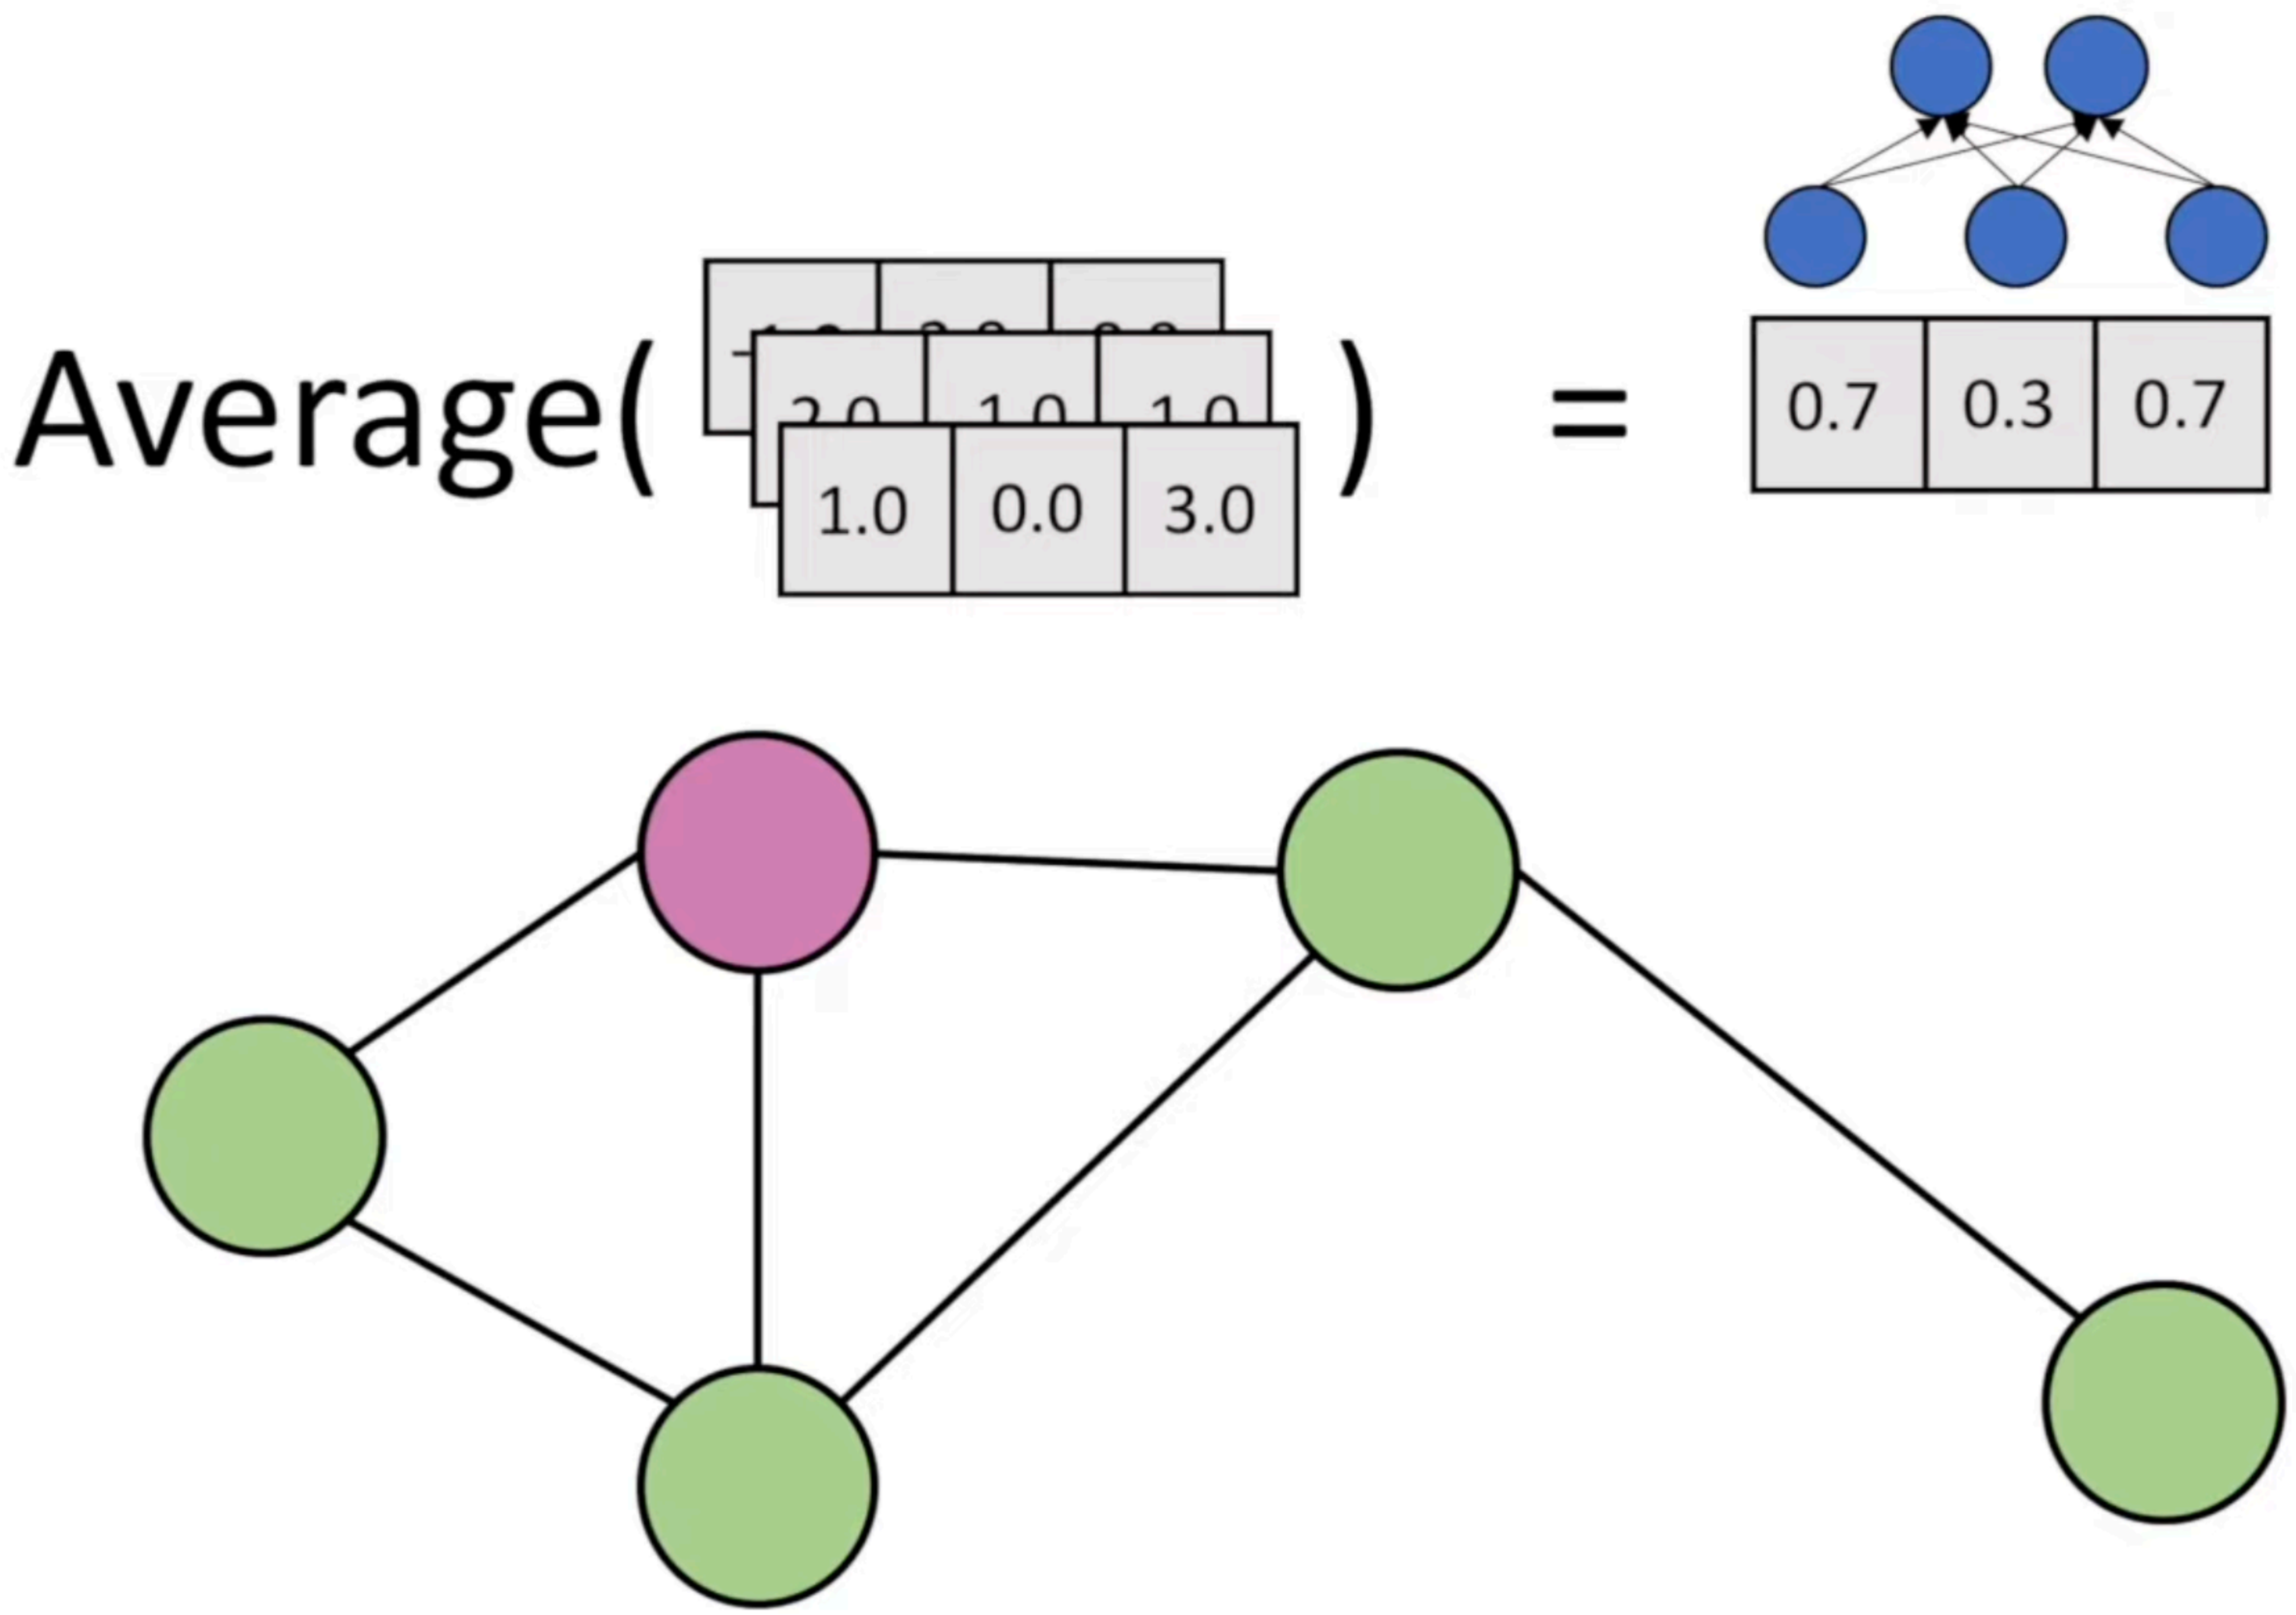

Simple Calculation in GNN Layer

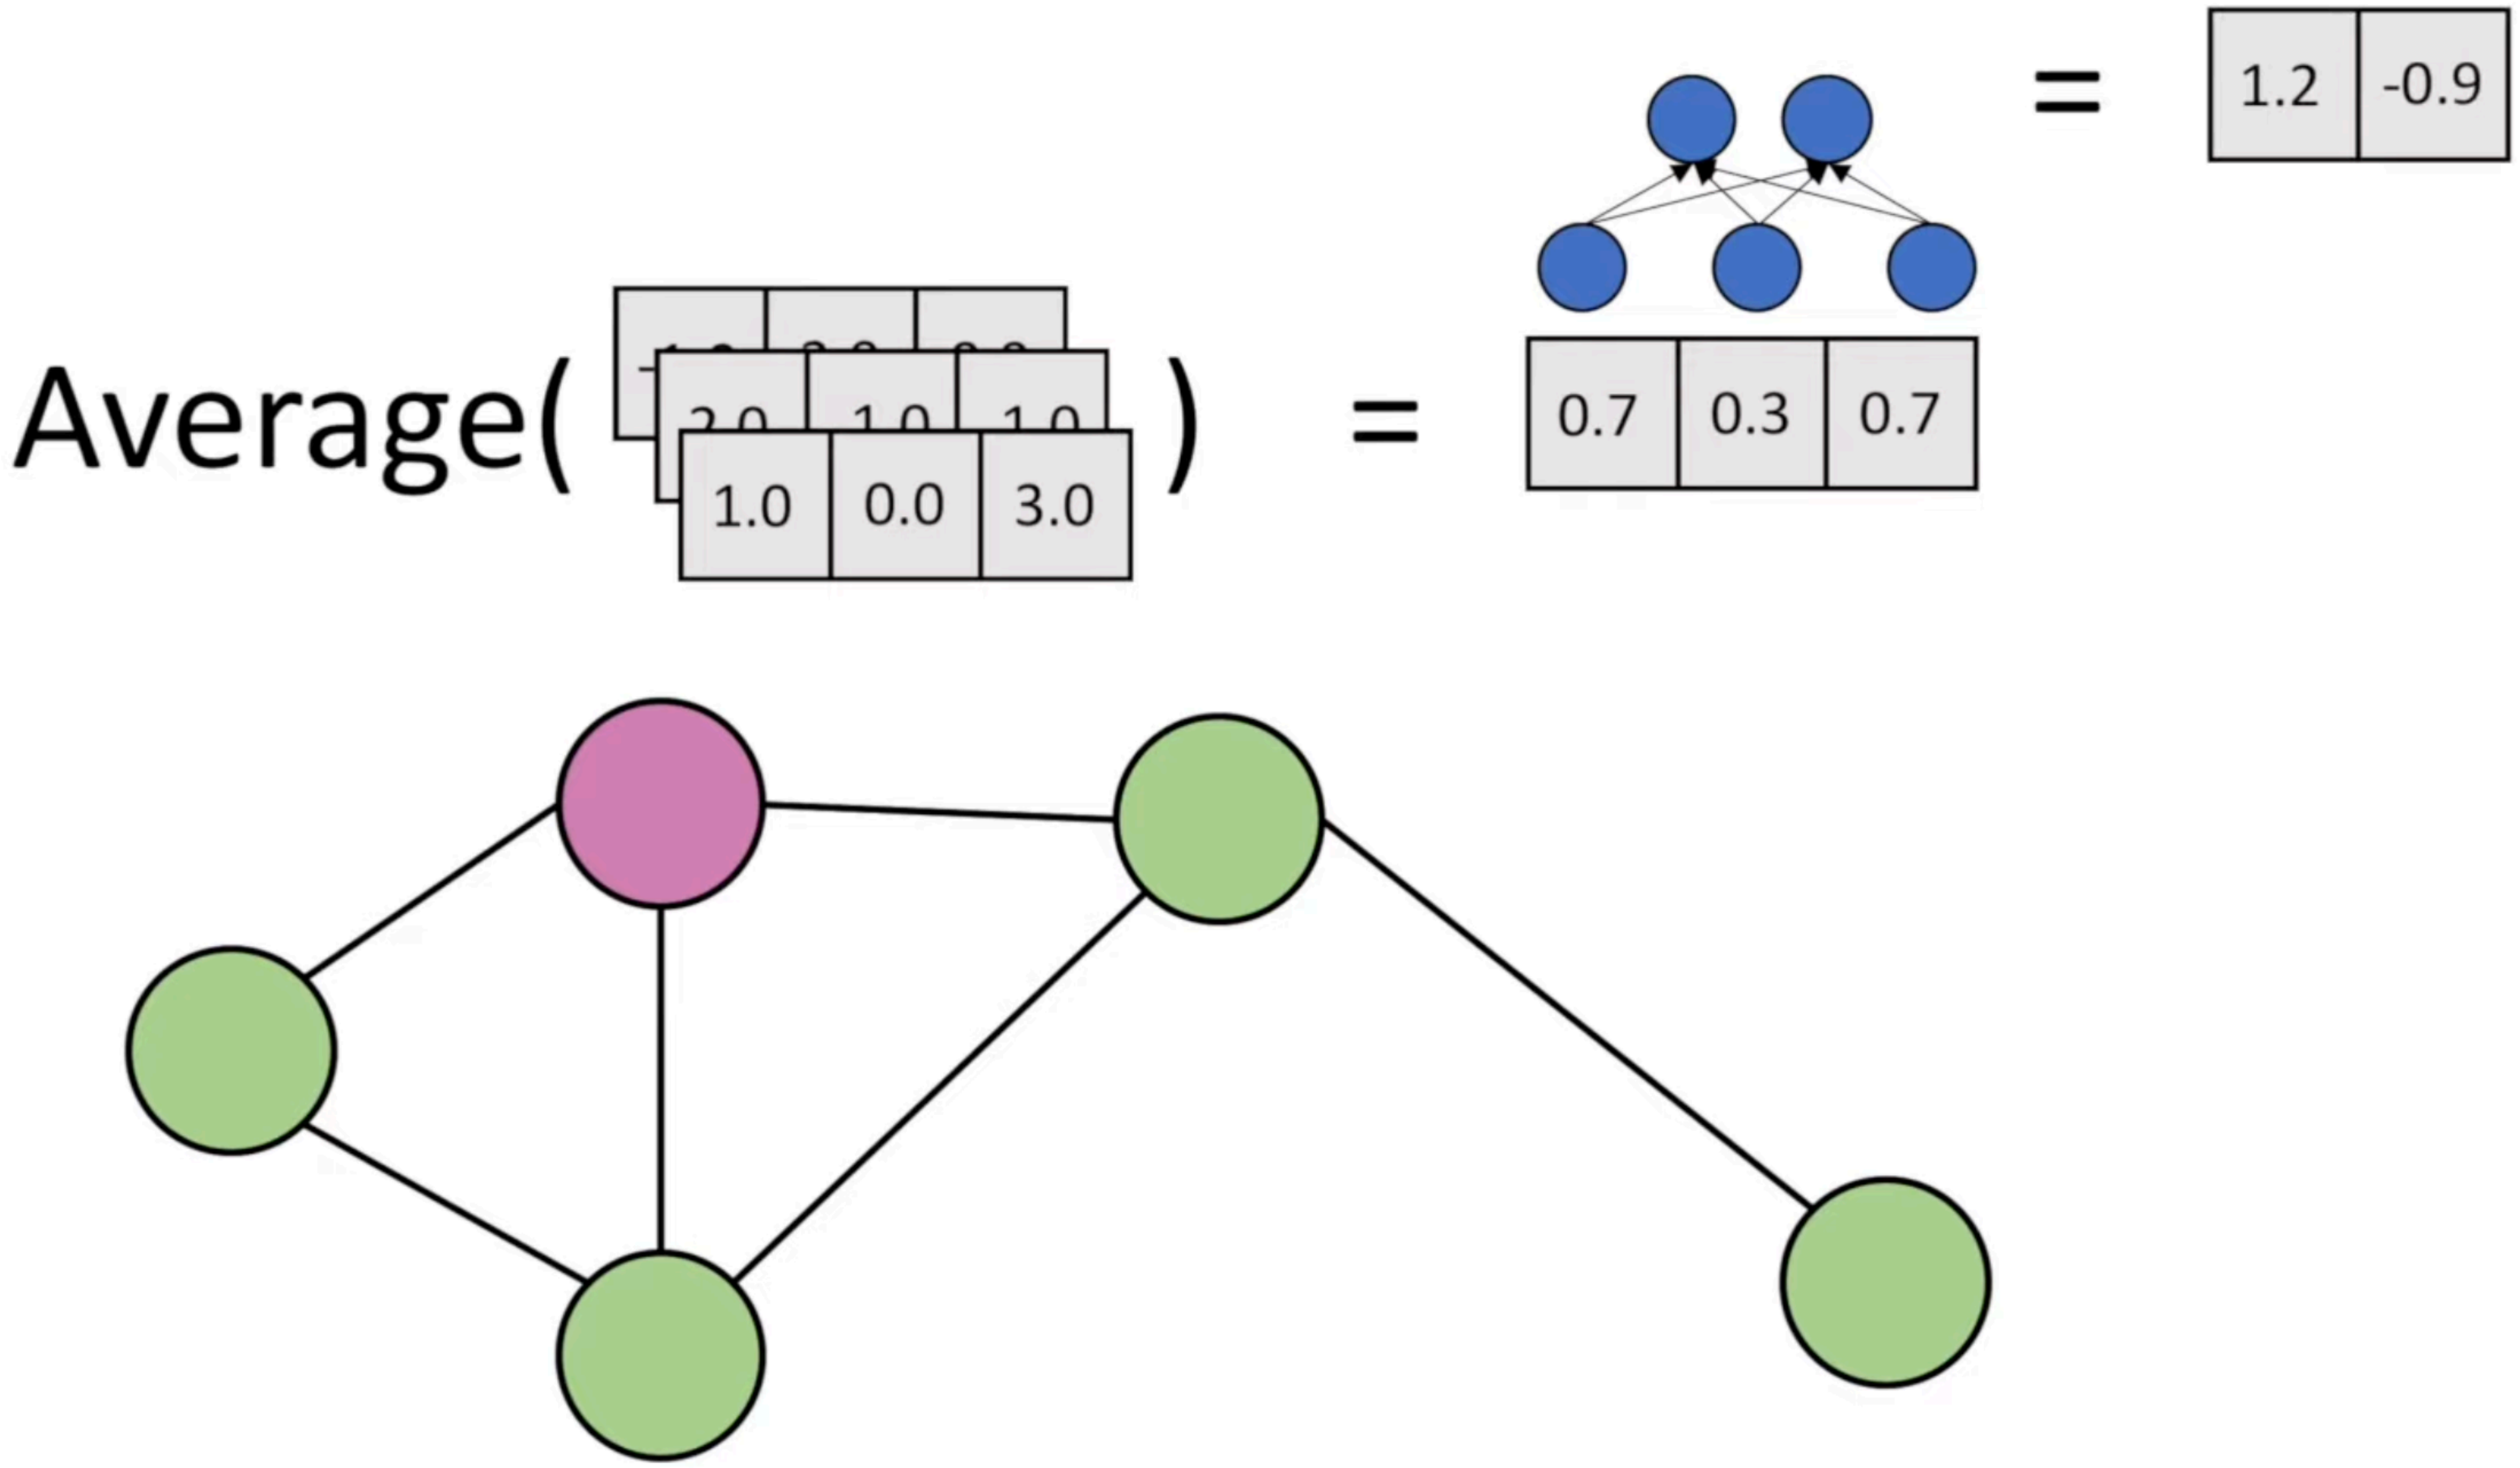

Simple Calculation in GNN Layer

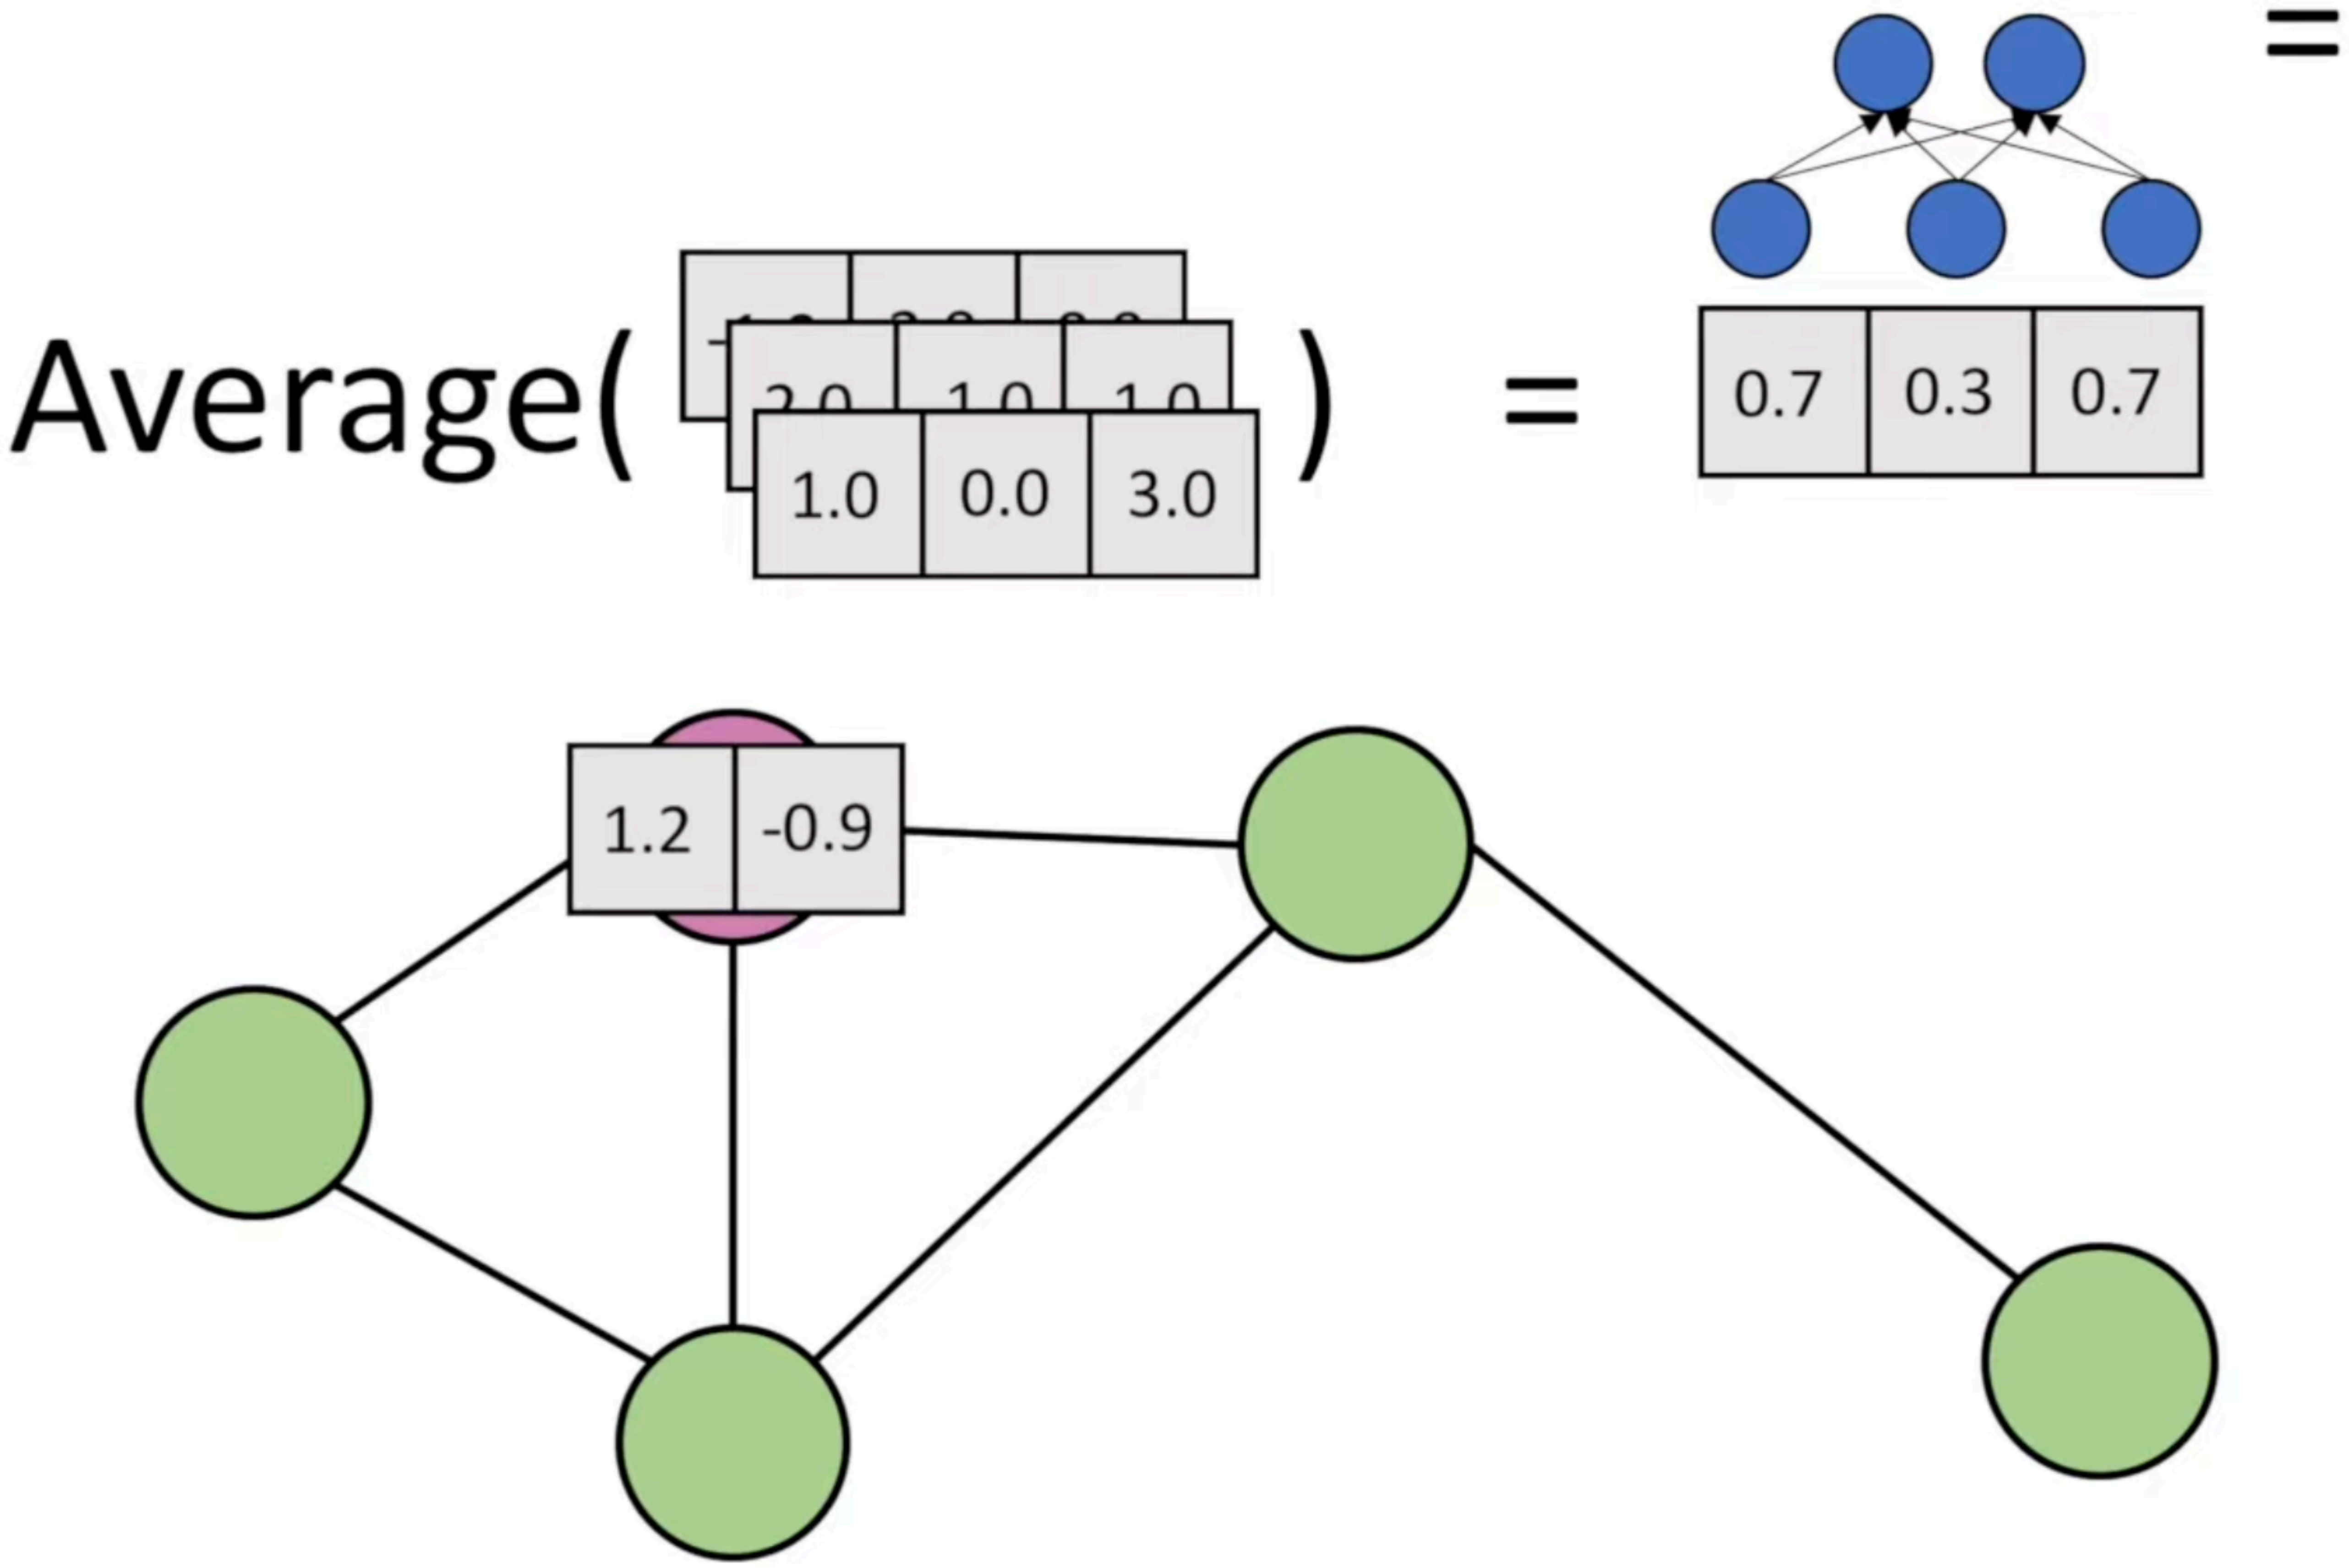

Example

Twitter Content Abuse

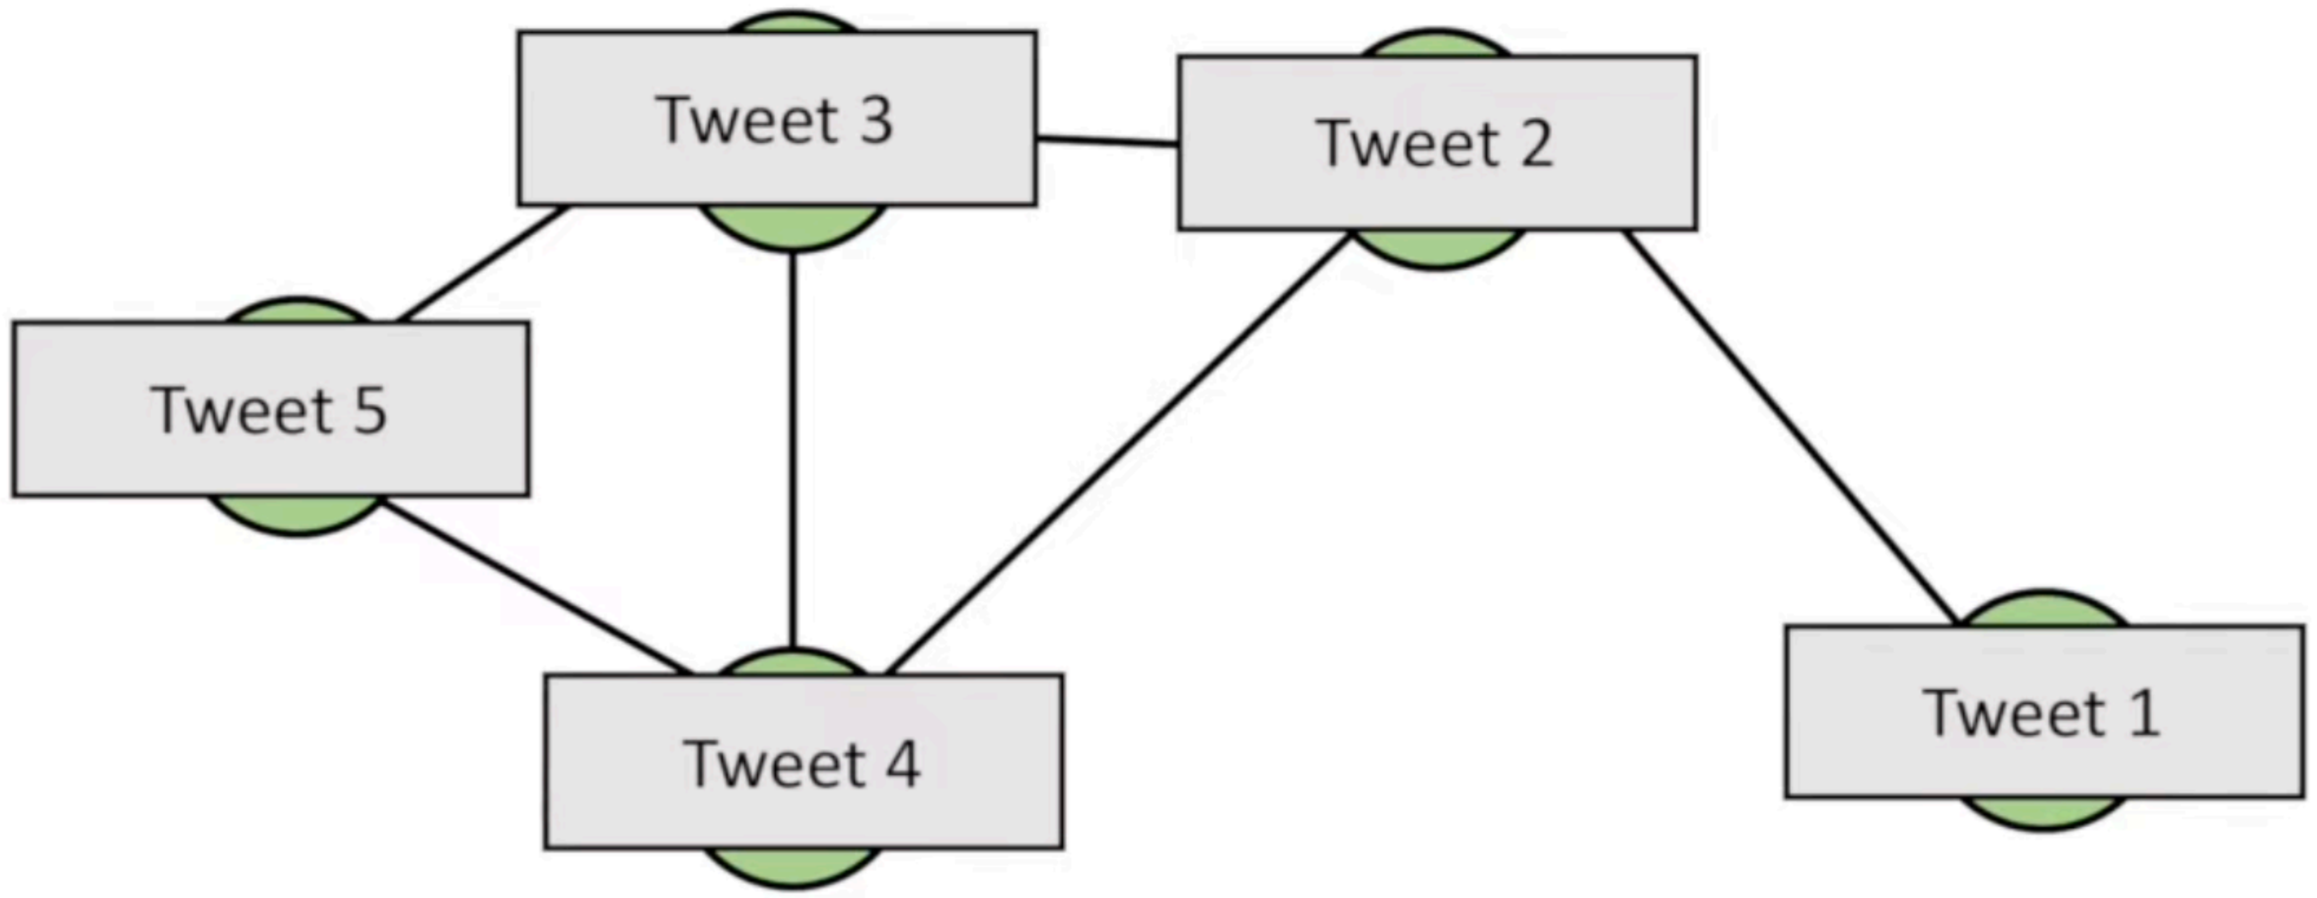

Example

Twitter Content Abuse

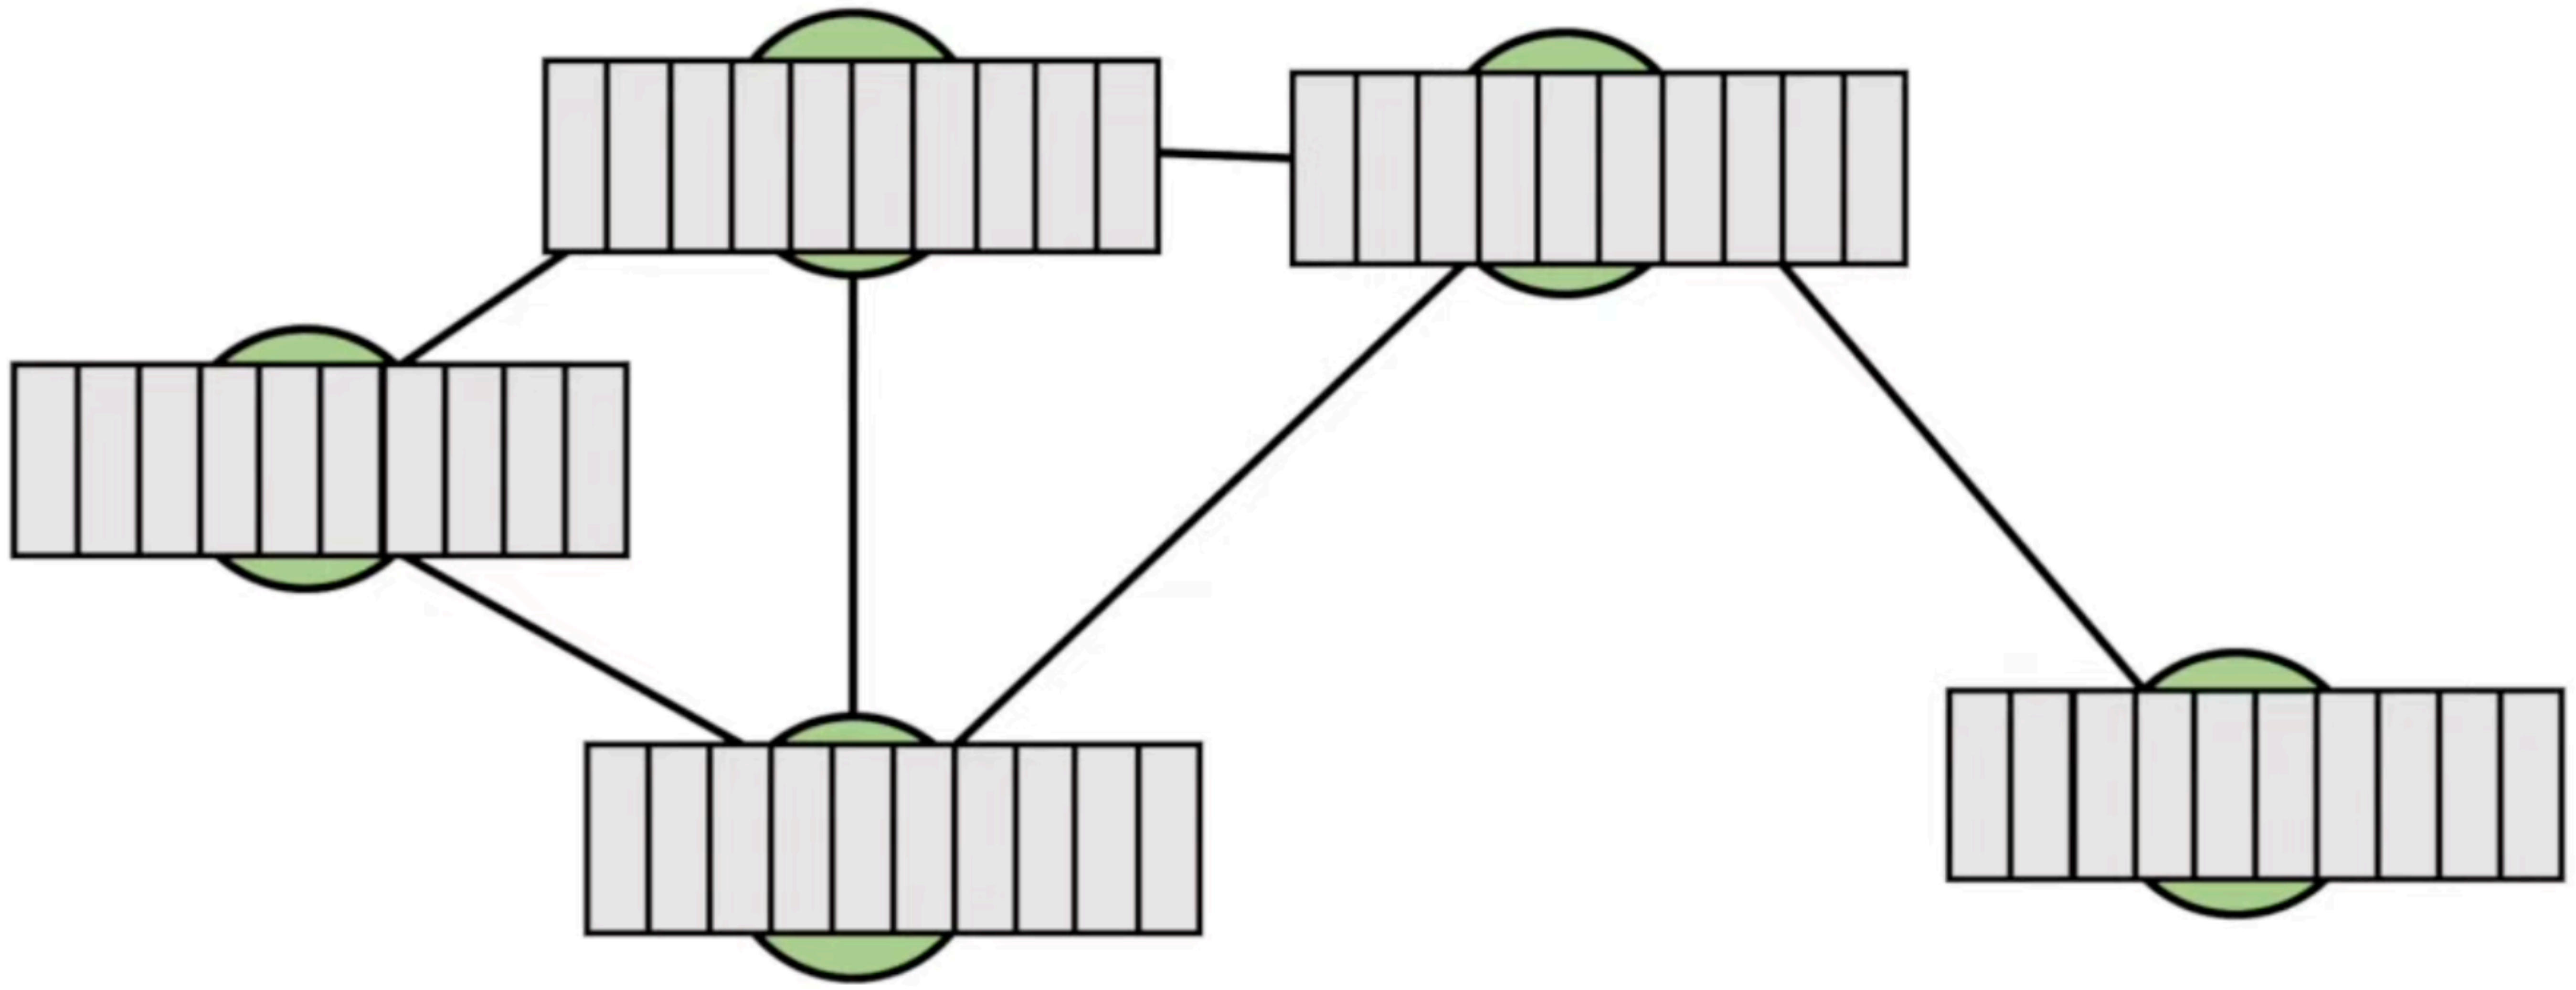

Example

Twitter Content Abuse

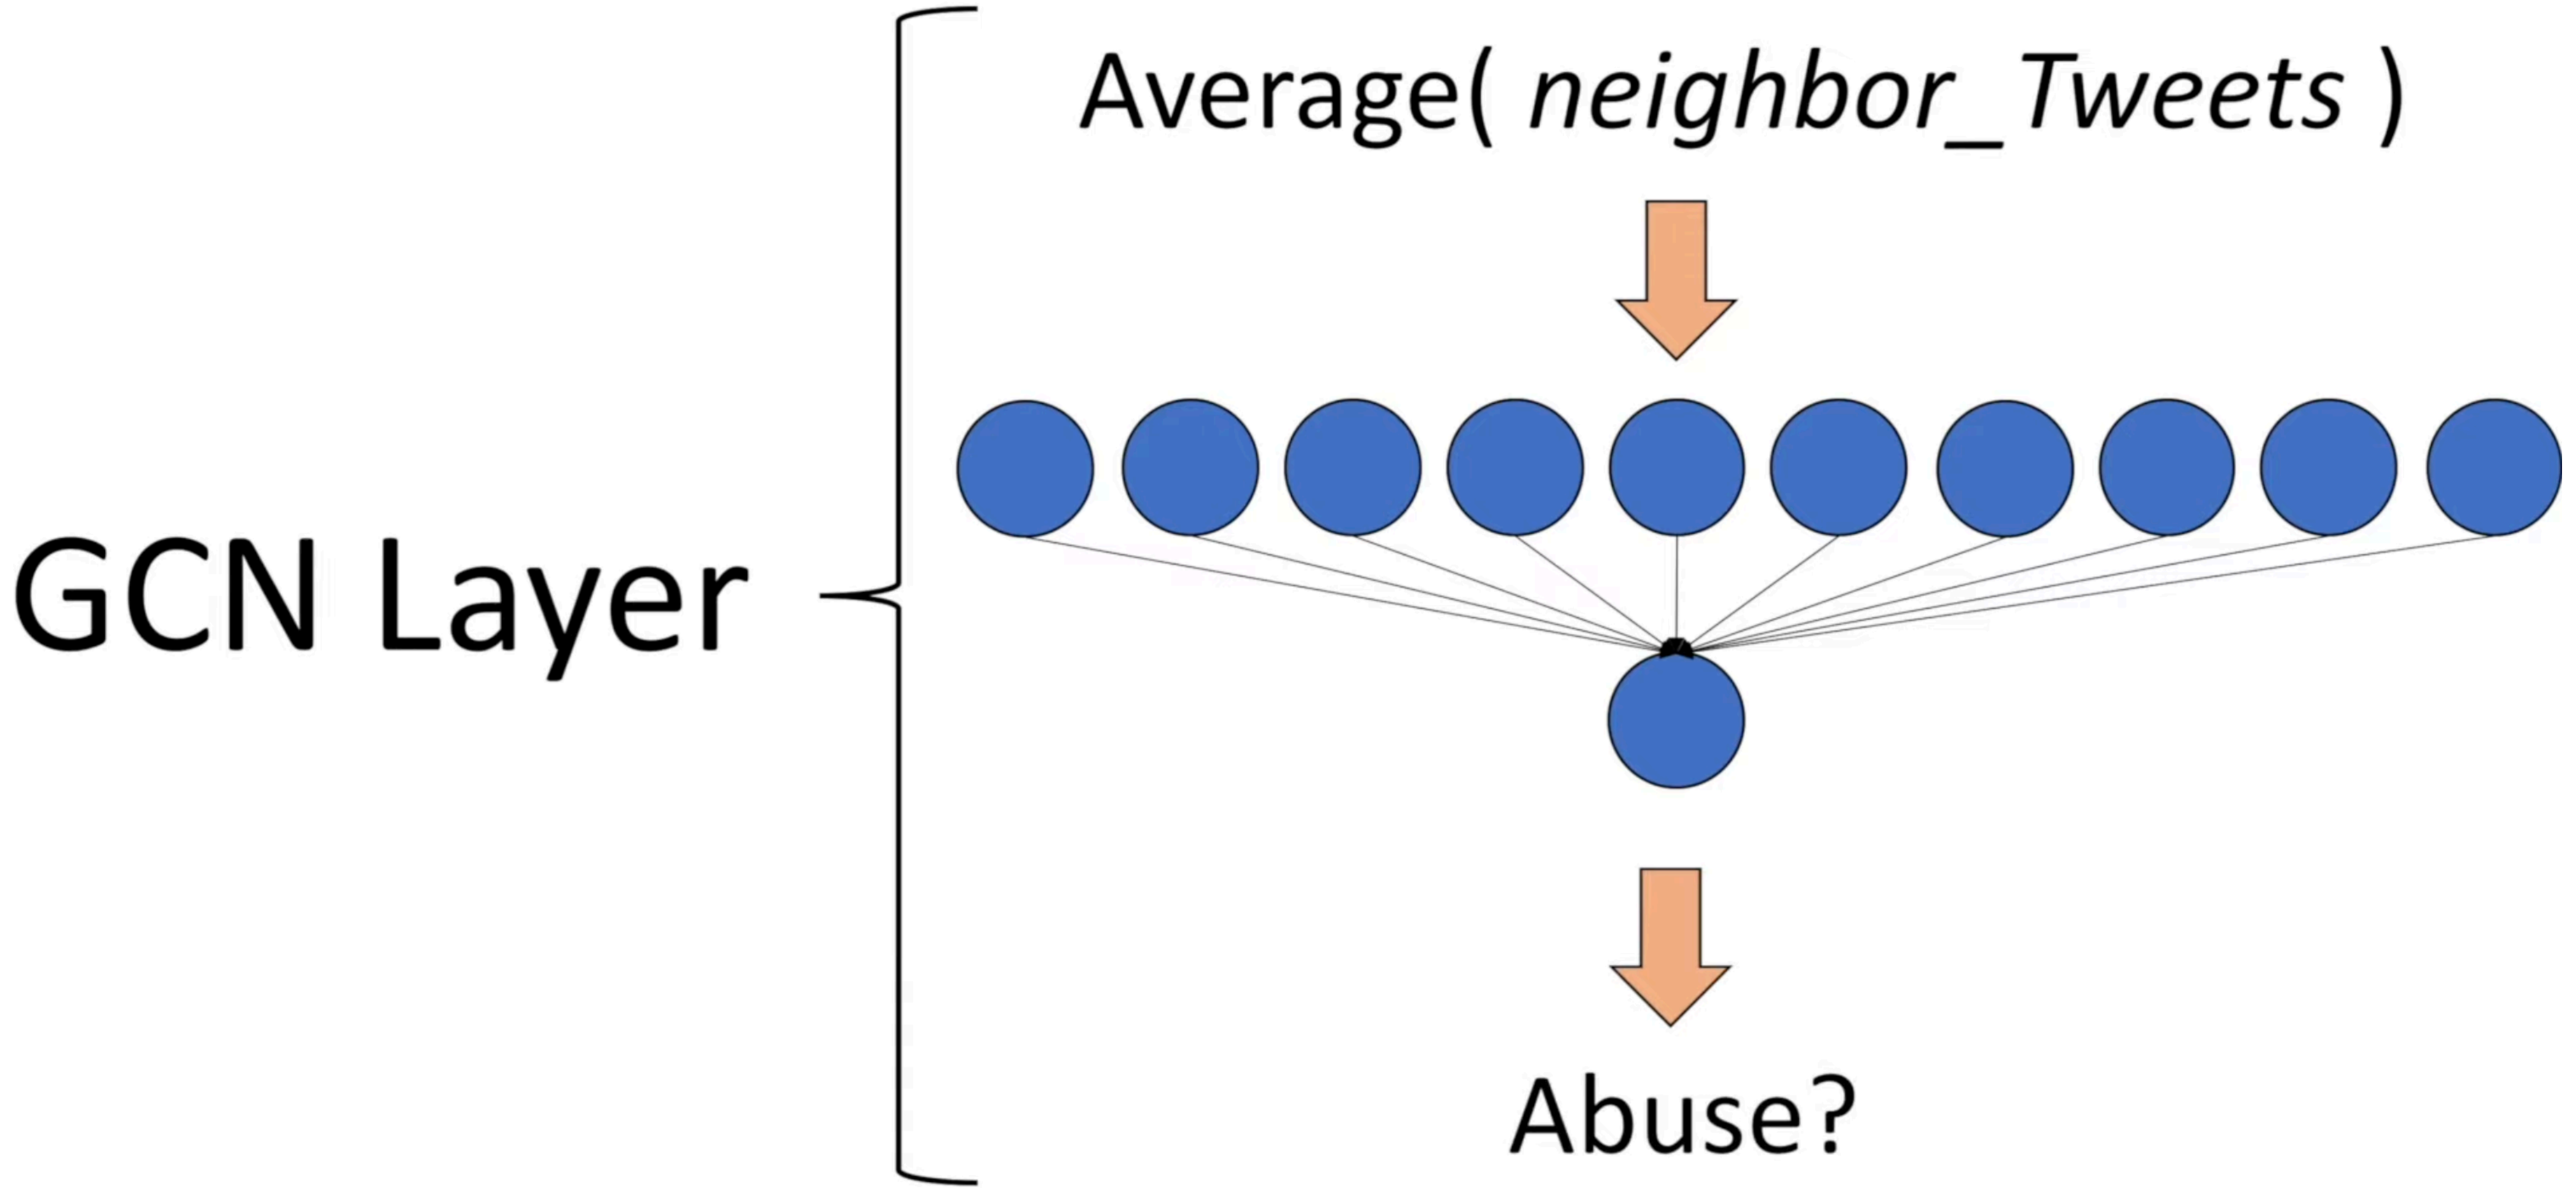

Two Layer Example

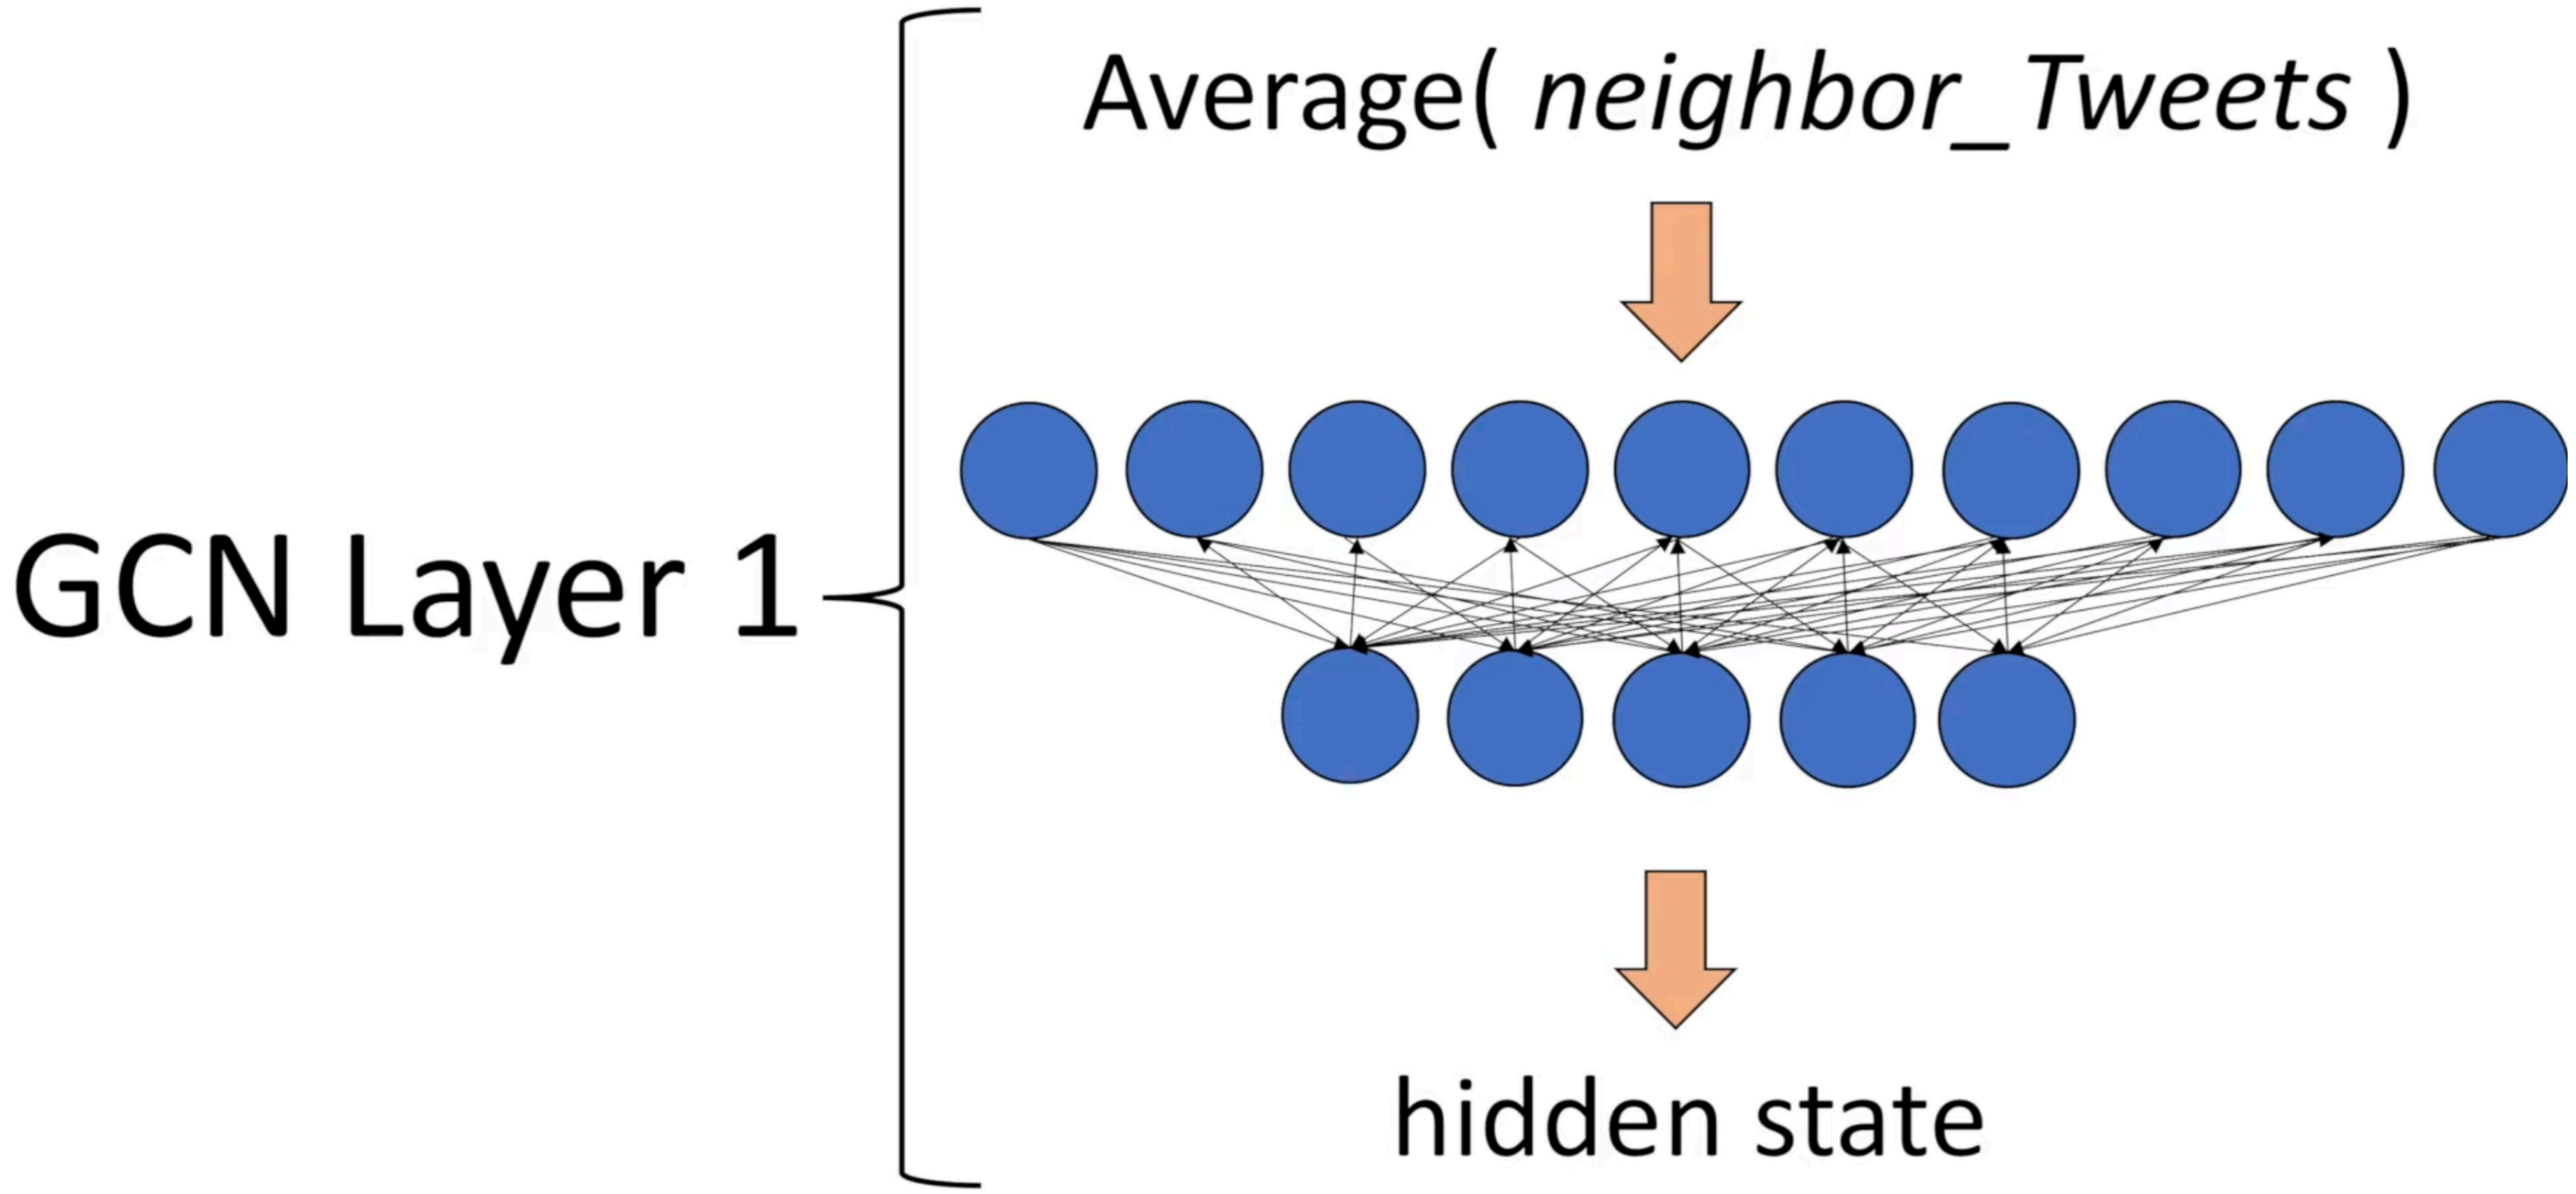

Two Layer Example

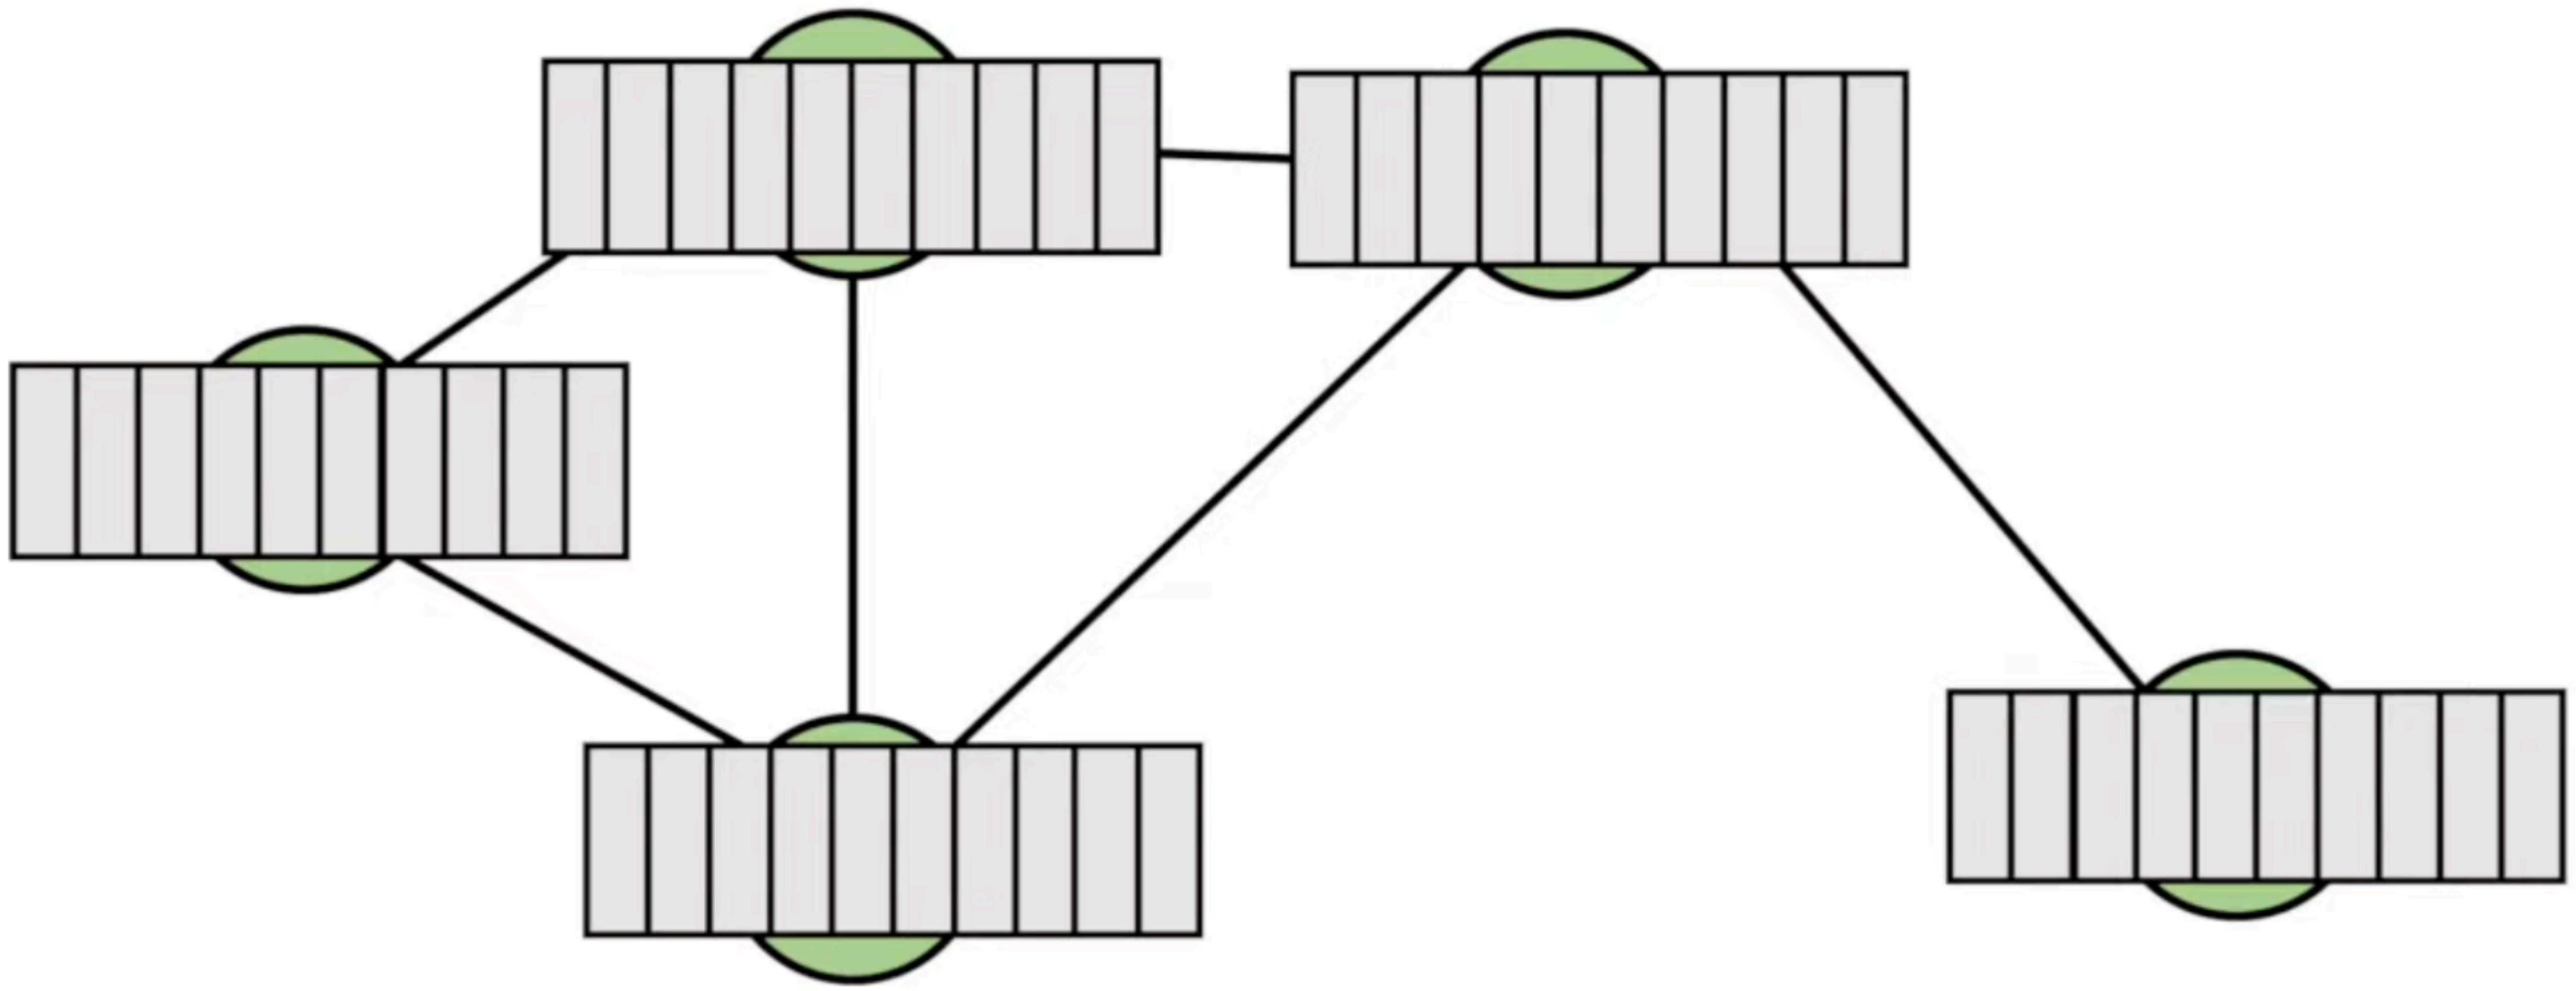

Two Layer Example

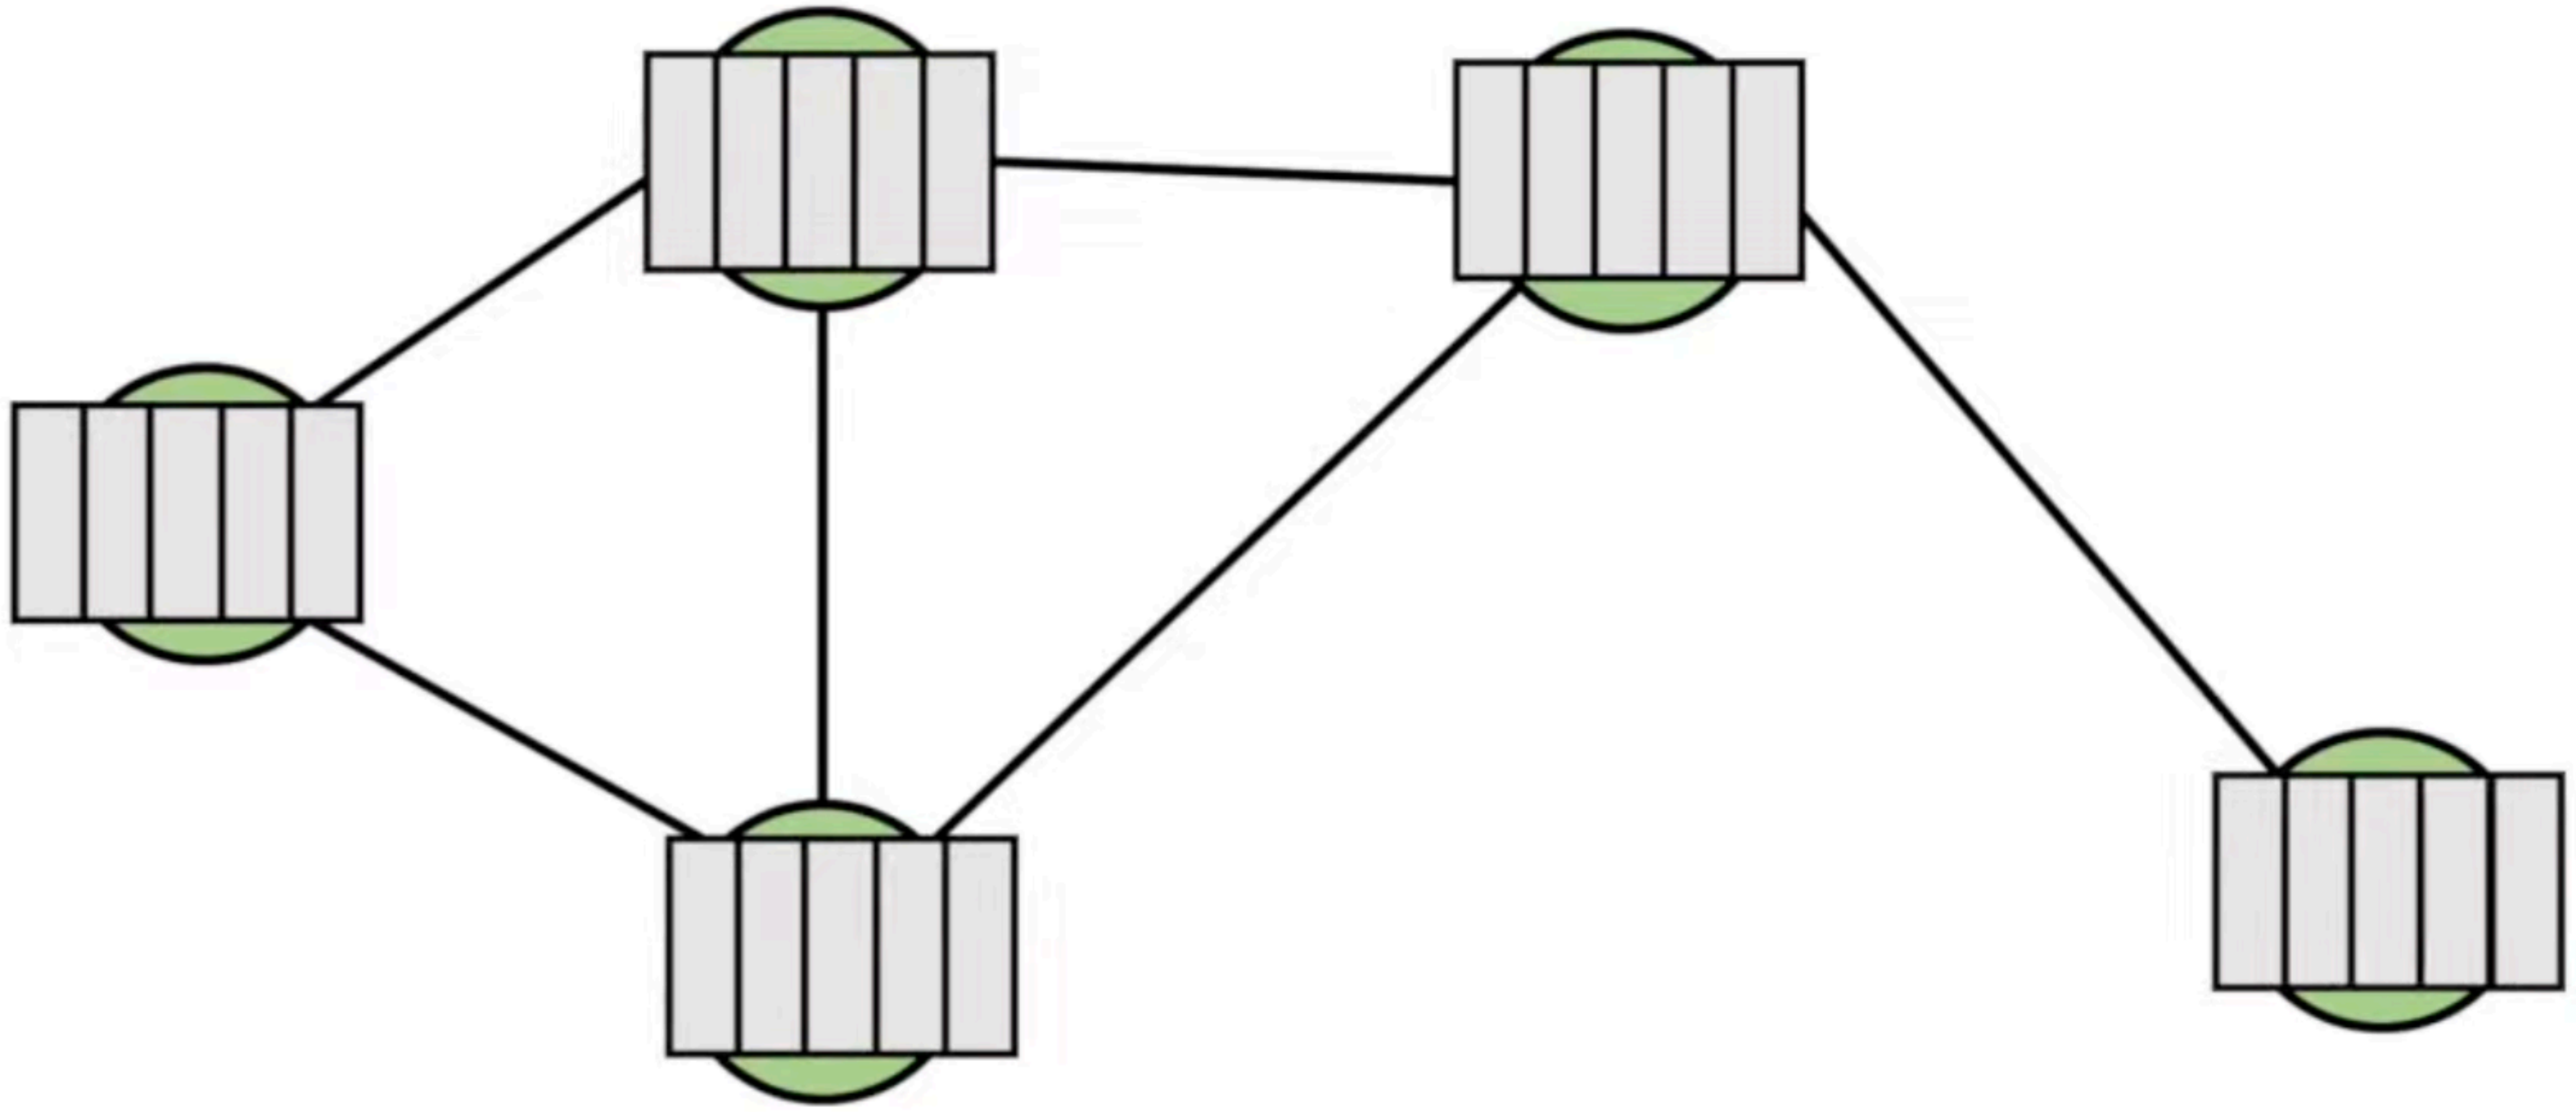

Two Layer Example

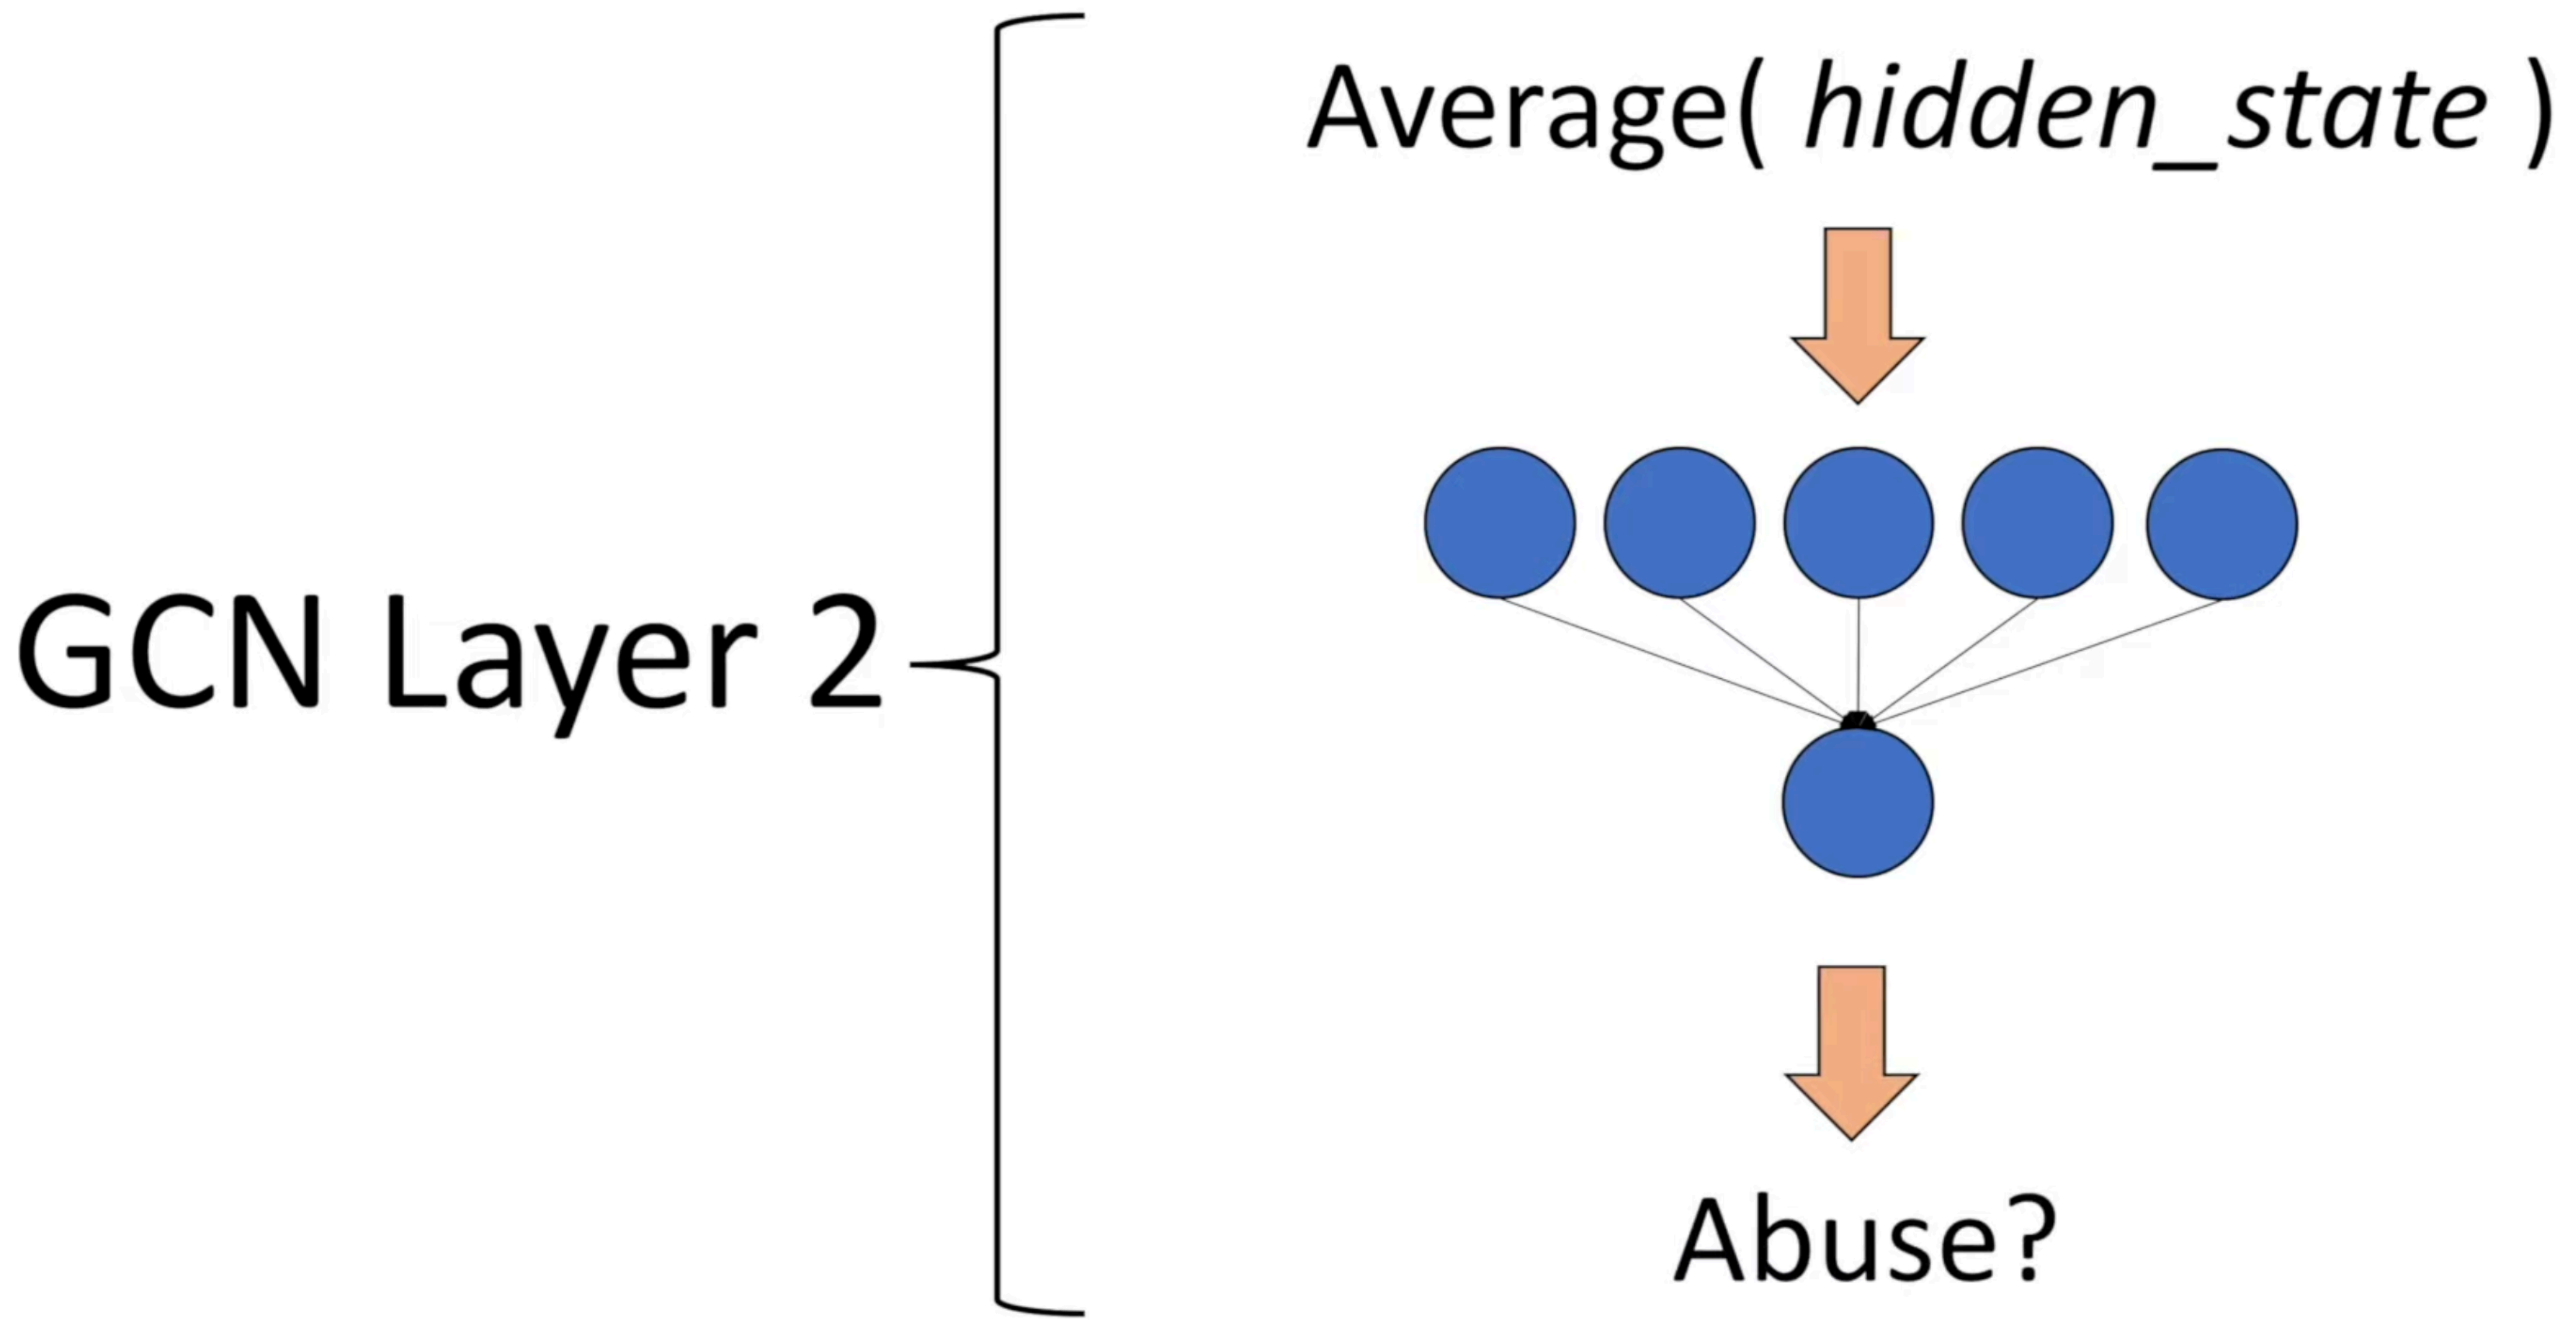

GNN Playground

- Graph-level prediction task with small molecular graphs
    - Each molecule is a graph
    - Atoms are nodes containing encoding for atomic identity
    - Bonds are edges containing encoding for bond type

A Gentle Introduction to Graph Neural Networks
Benjamin Sanchez-Lengeling, Emily Reif, Adam Pearce, Alexander B. Wiltschko
Google Research, https://doi.org/10.23915/distill.00033

Resources and Sources

A Gentle Introduction to Graph Neural Networks
Benjamin Sanchez-Lengeling, Emily Reif, Adam Pearce, Alexander B. Wiltschko
Google Research, https://doi.org/10.23915/distill.00033

Graph Convolutional Neural Networks (GCNs) Made Simple
WelcomeAIOverloards
https://youtu.be/2KRAOZIULzw

An Example of Graph Convolutional Networks
Zak Jost
https://blog.zakjost.com/post/gcn_citeseer/

Neural Message Passing for Quantum Chemistry
J. Gilmer, S.S. Schoenholz, P.F. Riley, O. Vinyals, G.E. Dahl.
Proceedings of the 34th International Conference on Machine Learning, Vol 70, pp. 1263--1272. PMLR. 2017.

Relational inductive biases, deep learning, and graph networks
P.W. Battaglia, J.B. Hamrick, V. Bapst, A. Sanchez-Gonzalez, V. Zambaldi, M. Malinowski, A. Tacchetti, D. Raposo, A. Santoro, R. Faulkner, C. Gulcehre, F. Song, A. Ballard, J. Gilmer, G. Dahl, A. Vaswani, K. Allen, C. Nash, V. Langston, C. Dyer, N. Heess, D. Wierstra, P. Kohli, M. Botvinick, O. Vinyals, Y. Li, R. Pascanu.
2018.