

AIMS CDT SIGNAL PROCESSING

Introduction to graph neural networks

Dr. Dorina Thanou

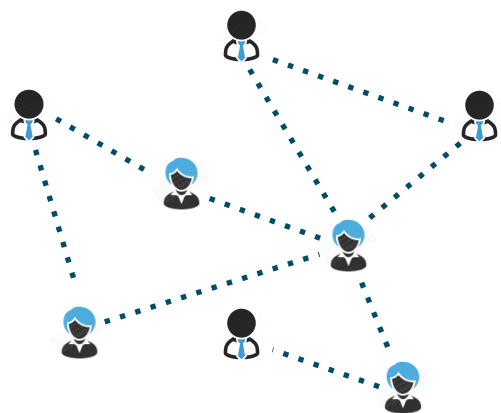
EPFL

21.10.2021

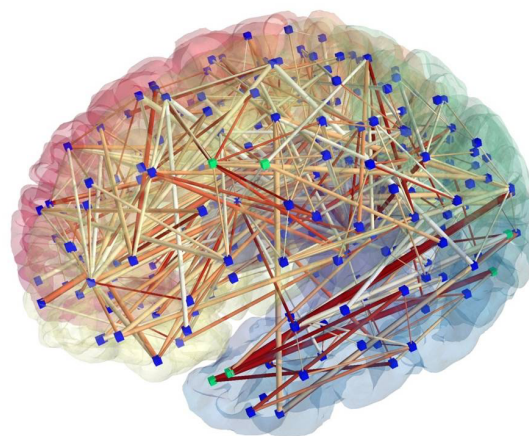
Outline

- Motivation: Why Graph Neural Networks?
- Basic definitions on graphs
- Partial historical overview
 - Graph CNN (main focus)
 - Graph autoencoders (briefly)
- Applications
 - Naturally graph-structured data
 - Images

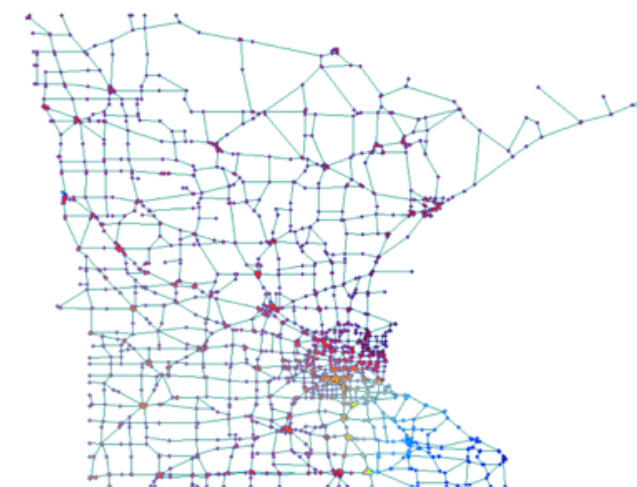
Network data are pervasive



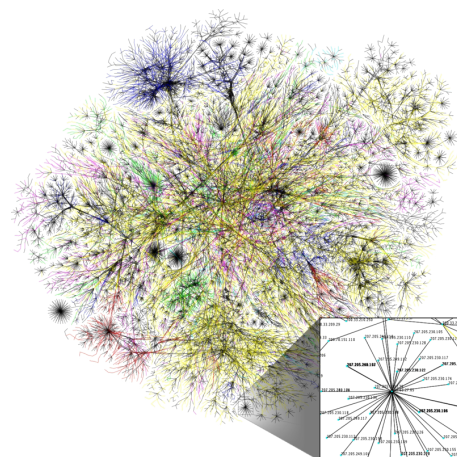
Social networks



Biological networks



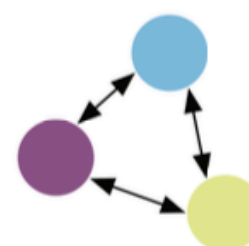
Transportation networks



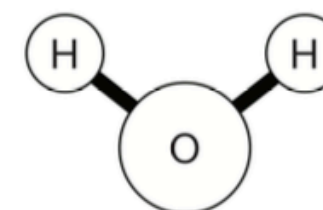
Communication networks



n-body system



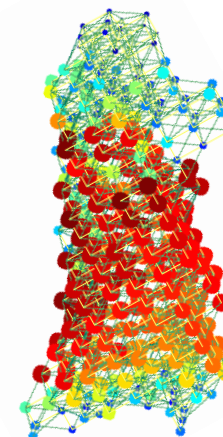
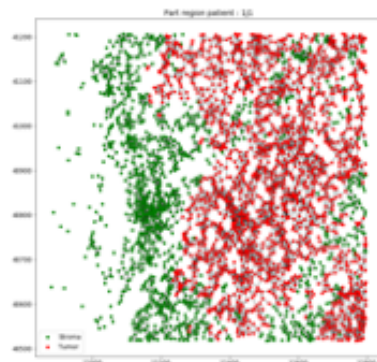
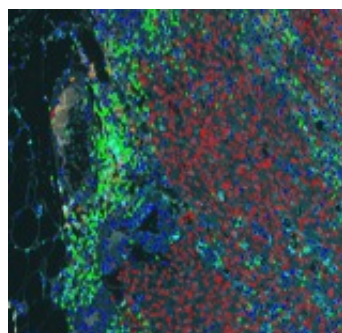
molecule



Graphs provide a mathematical representation for describing and modeling complex systems

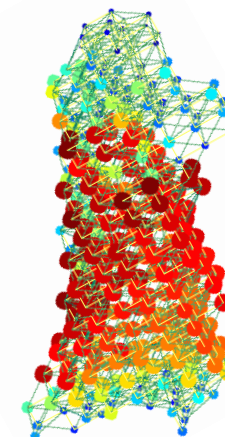
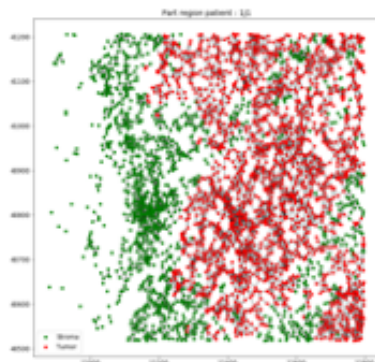
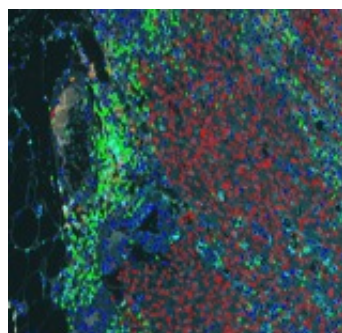
Why graph-based modelling?

- Provide a different perspective of data

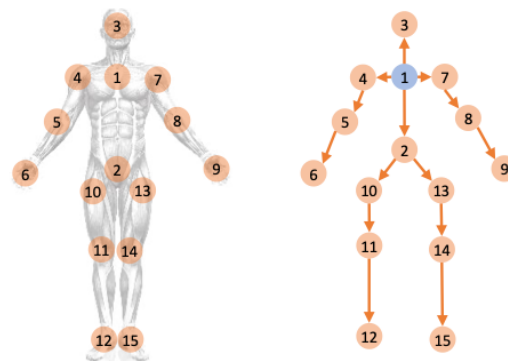


Why graph-based modelling?

- Provide a different perspective of data

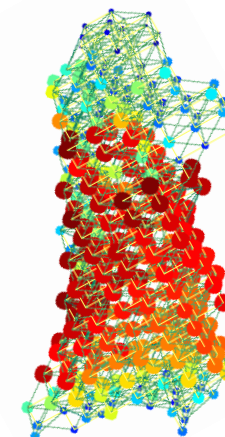
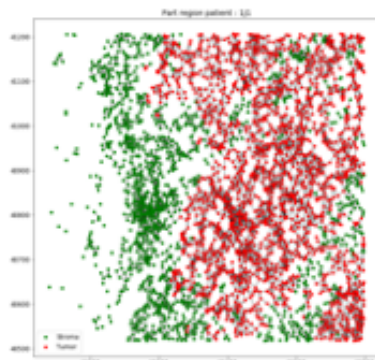
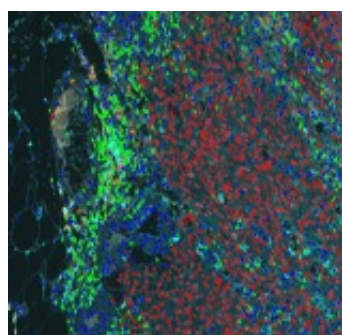


- Impose relational inductive bias in data

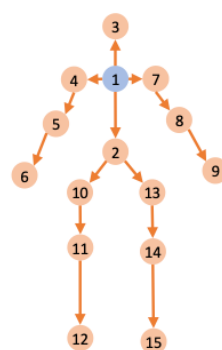
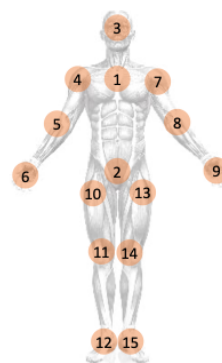


Why graph-based modelling?

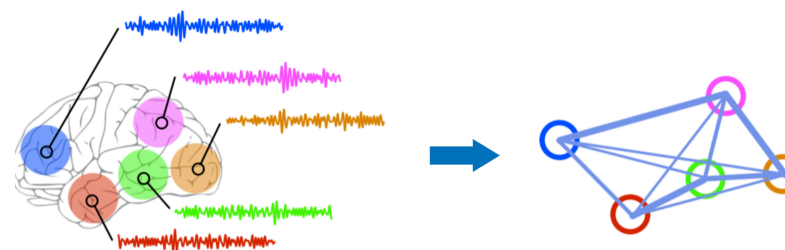
- Provide a different perspective of data



- Impose relational inductive bias in data

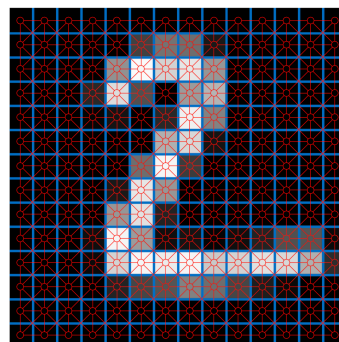


- Lead to knowledge discovery

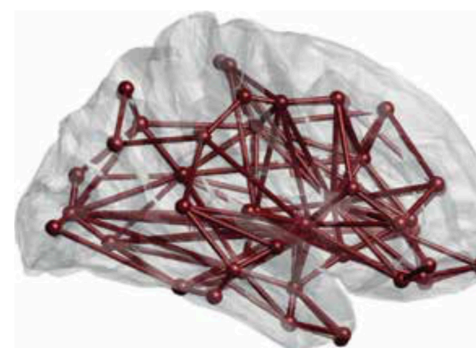


Typical learning tasks on graph-structured data

- Graph-wise classification

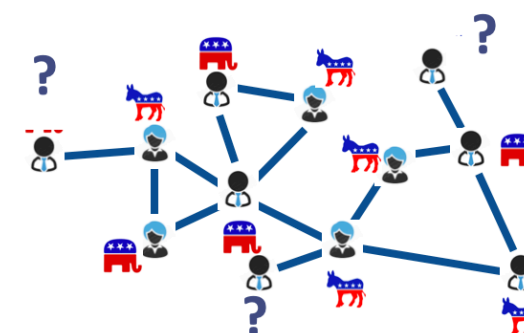


is it a 2?
is it a 4?

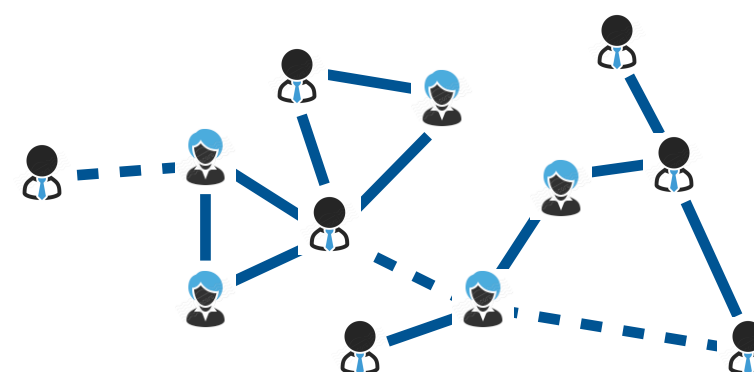


condition?
no condition?

- Vertex-wise classification/inference of missing values

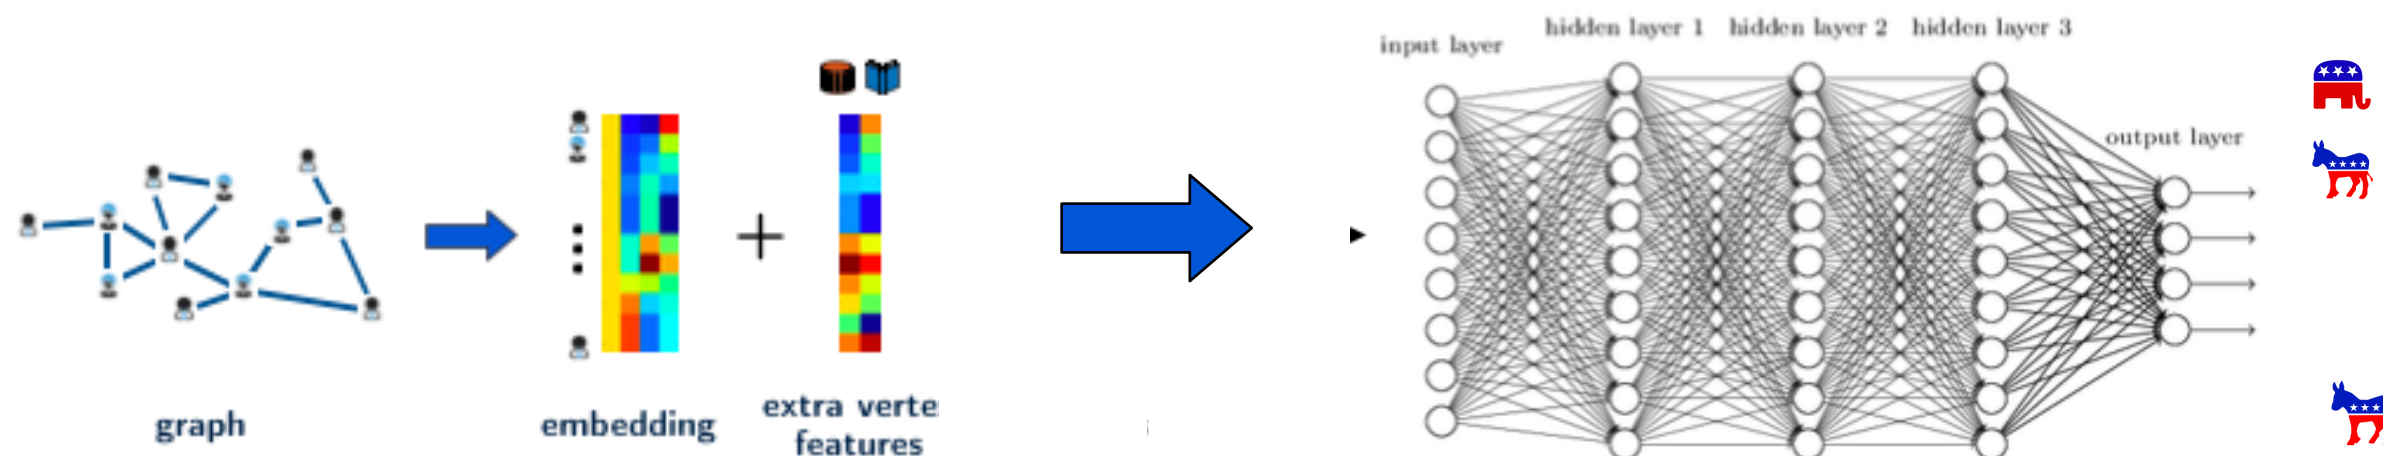


- Link prediction



A naive approach

- Embed graph and features into a Euclidean space
- Feed them into a deep neural net

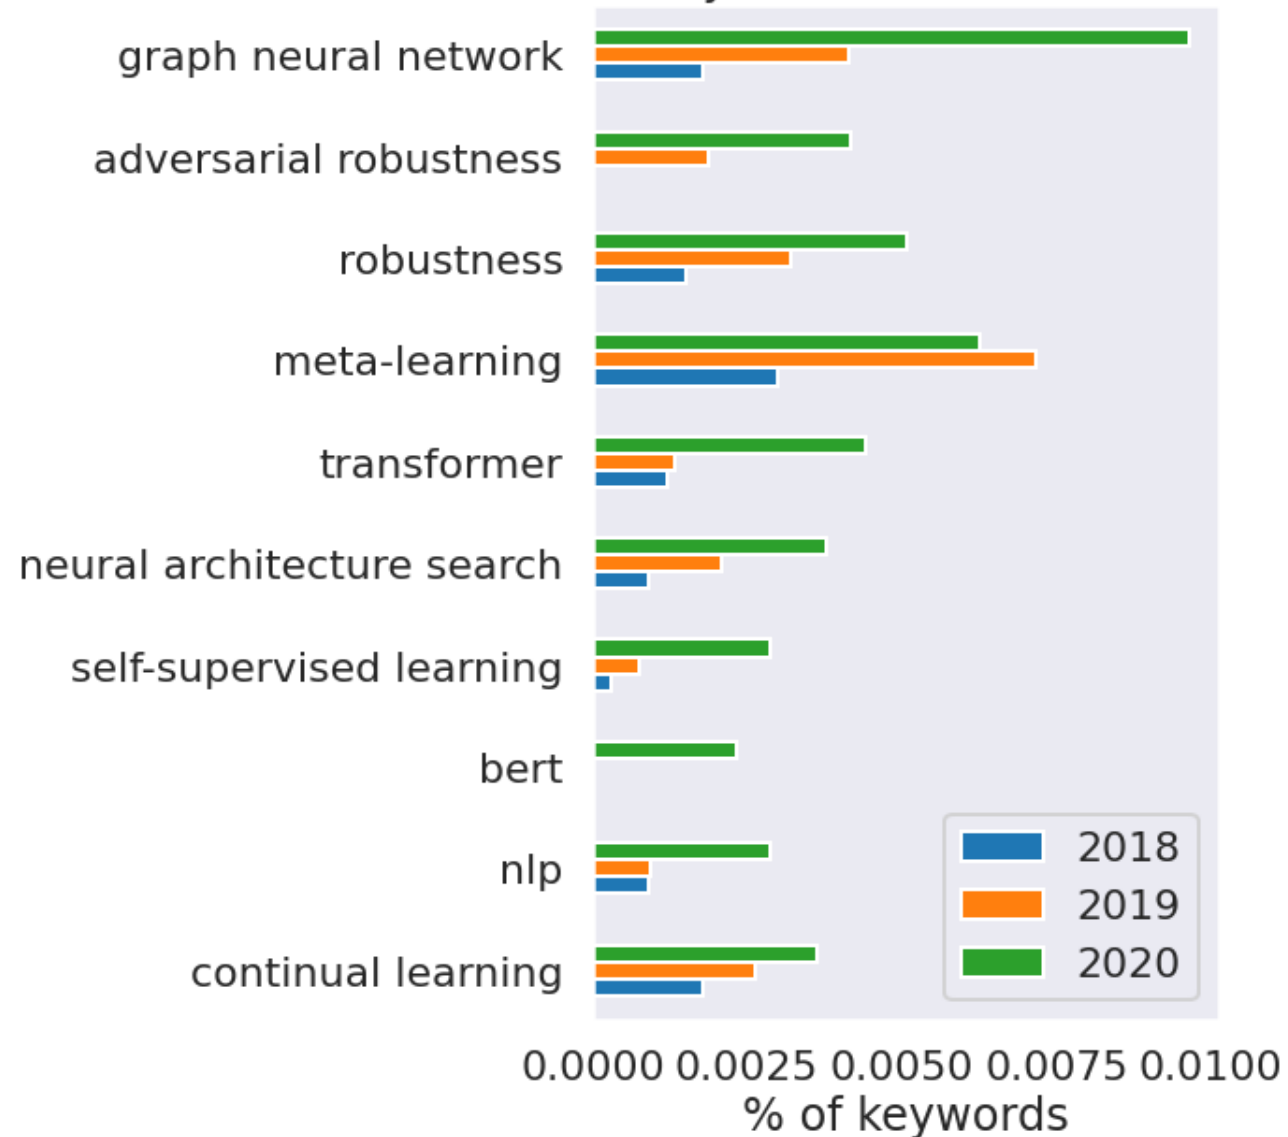


- Issues with that:
 - $O(N)$ parameters
 - Not applicable to graphs of different sizes
 - Not invariant to node ordering

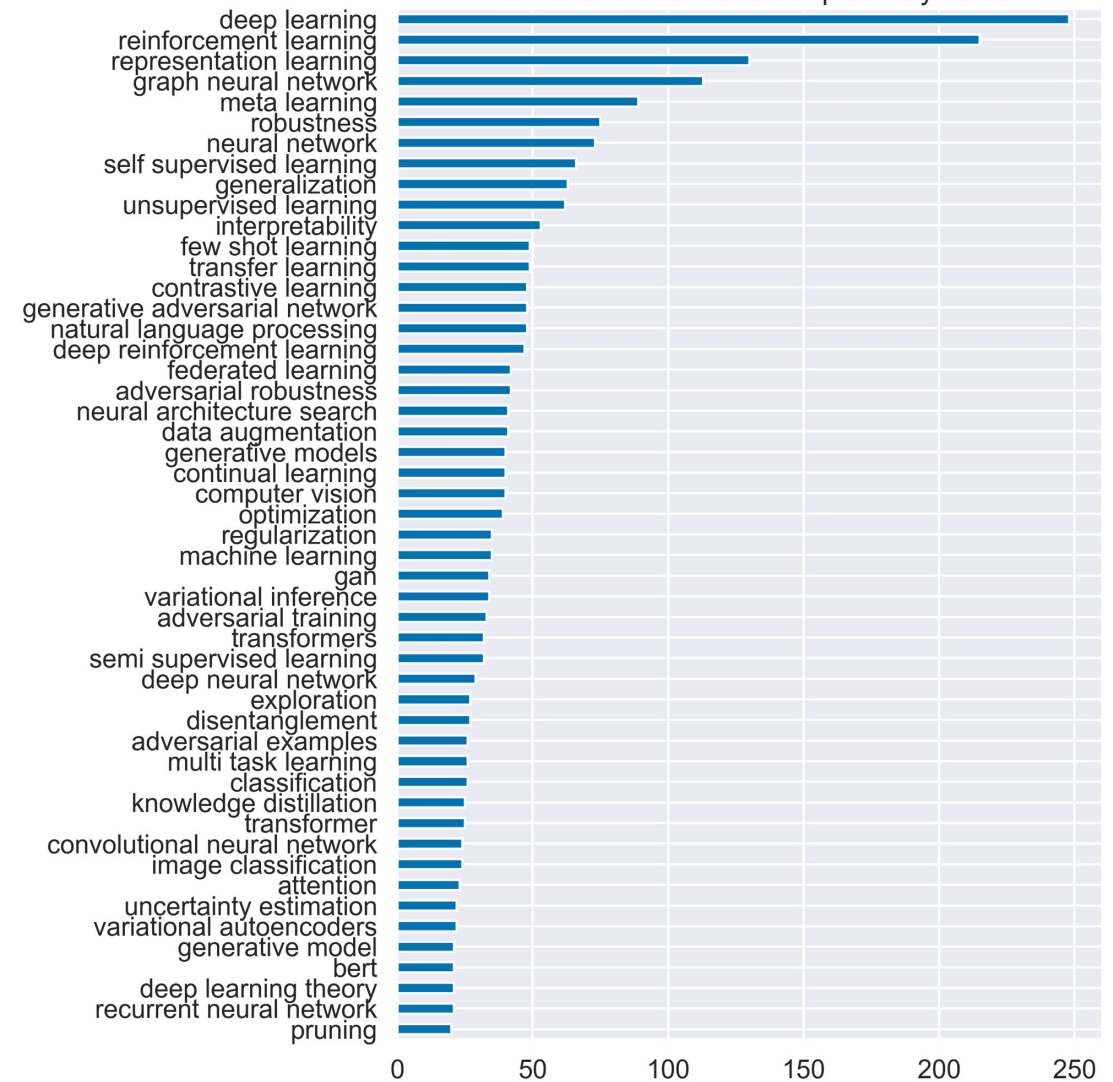
Can we do better?

GNN: A growing trend

ICLR Keyword Growth 2018-2020

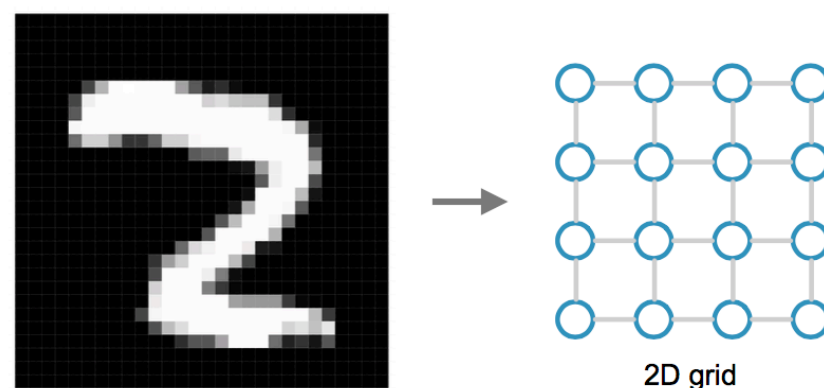


ICLR 2021 Submission Top 50 Keywords

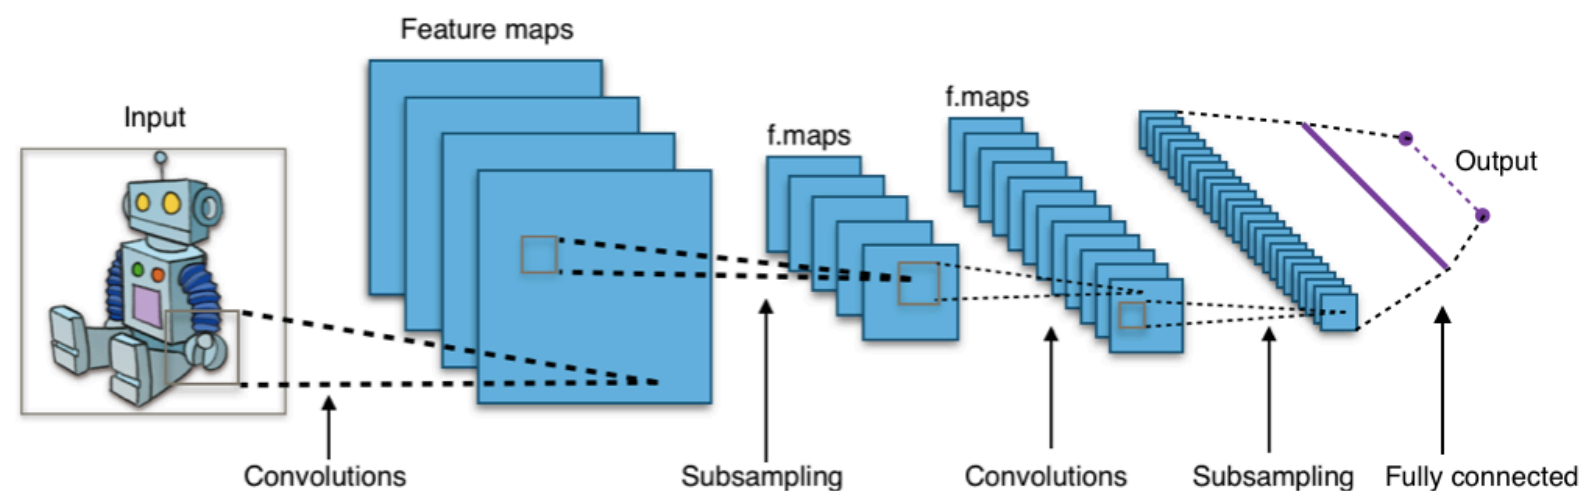


CNN on Euclidean data (I)

- **Main assumption:** data lives on a grid



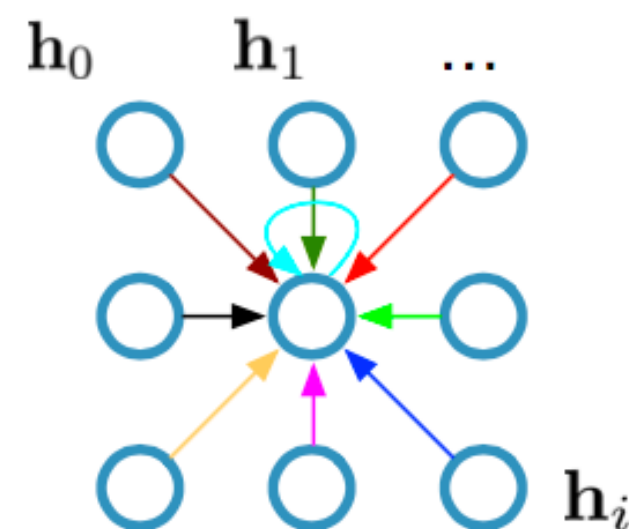
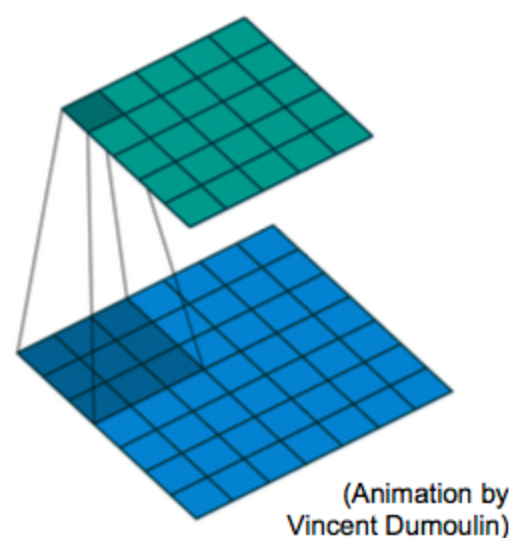
- Typical CNN architecture:



https://en.wikipedia.org/wiki/File:Typical_cnn.png

CNN on Euclidean data (II)

- Single CNN layer with 3x3 filter:



$$h_4^{(l+1)} = \sigma \left(\sum_{i=0}^8 W_i^{(l)} h_i^{(l)} \right)$$

Non-linearity

Filter weights

How can we extend CNN on graphs?

- Desirable properties
 - **Convolution:** how to achieve translation invariance
 - **Localization:** what is the notion of locality
 - **Graph pooling:** how to downsample on graphs
 - **Efficiency:** how to keep the computational complexity low
 - **Generalization:** how to build models that generalize to unseen graphs

Outline

- Motivation: Why Graph Neural Networks?
- **Basic definitions on graphs**
 - Convolution
 - Pooling
- Partial historical overview
 - Graph CNN (main focus)
 - Graph autoencoders (briefly)
- Applications
 - Naturally graph-structured data
 - Images

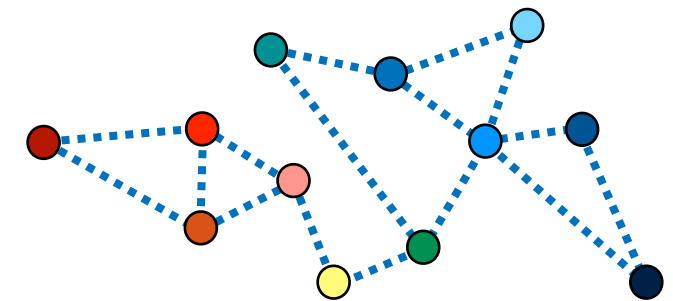
Graphs and signals on graphs

- **Irregular domain:** connected, undirected, weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$
- **Graph description:**
 - Degree matrix D : diagonal matrix with sum of weights of incident edges
 - Laplacian matrix L : $L = D - W$, $L = \chi \Lambda \chi^T$
 - Complete set of orthonormal eigenvectors $\chi = [\chi_0, \chi_1, \dots, \chi_{N-1}]$
 - Real, non-negative eigenvalue $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{N-1}$
- **Graph signal:** a function $y : \mathcal{V} \rightarrow \mathbb{R}$ that assigns real values to each vertex of the graph

Graphs and signals on graphs

- **Irregular domain:** connected, undirected, weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, W)$

- **Graph description:**



- Degree matrix D : diagonal matrix with sum of weights of incident edges
- Laplacian matrix L : $L = D - W$, $L = \chi \Lambda \chi^T$
 - Complete set of orthonormal eigenvectors $\chi = [\chi_0, \chi_1, \dots, \chi_{N-1}]$
 - Real, non-negative eigenvalue $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_{N-1}$
- **Graph signal:** a function $y : \mathcal{V} \rightarrow \mathbb{R}$ that assigns real values to each vertex of the graph

Frequency on the graph

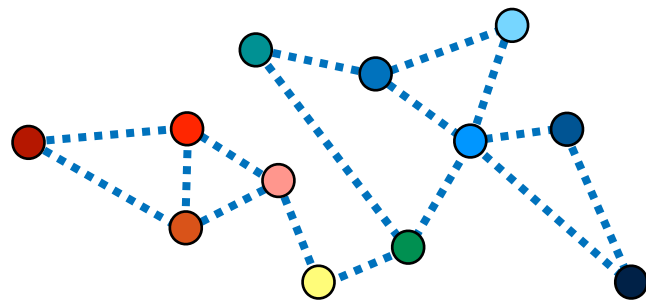
- Global smoothness on the graph

$$y^T L y = \sum_{m \in \mathcal{V}} \sum_{n \in \mathcal{N}_m} W_{m,n} [y(m) - y(n)]^2$$

Frequency on the graph

- Global smoothness on the graph

$$y^T L y = \sum_{m \in \mathcal{V}} \sum_{n \in \mathcal{N}_m} W_{m,n} [y(m) - y(n)]^2$$

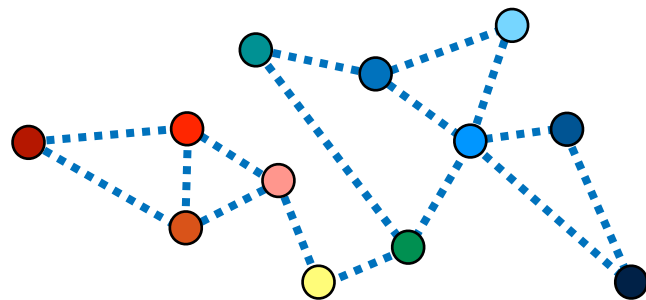


$$y^T \mathcal{L}_1 y = 0.15$$

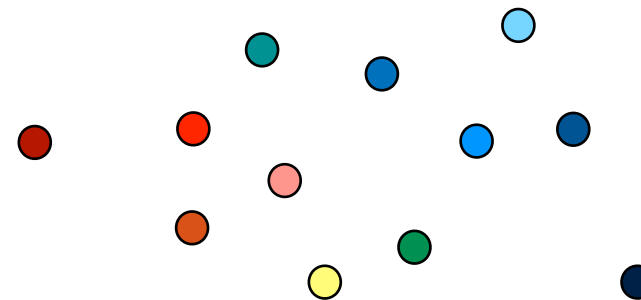
Frequency on the graph

- Global smoothness on the graph

$$y^T L y = \sum_{m \in \mathcal{V}} \sum_{n \in \mathcal{N}_m} W_{m,n} [y(m) - y(n)]^2$$



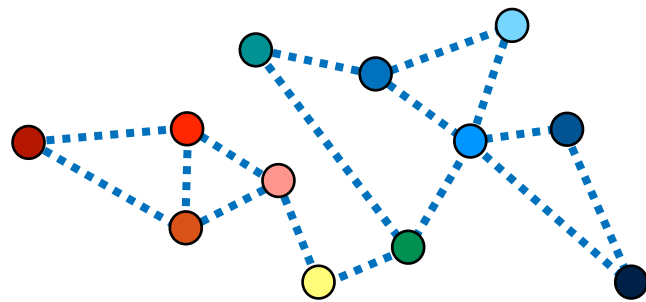
$$y^T \mathcal{L}_1 y = 0.15$$



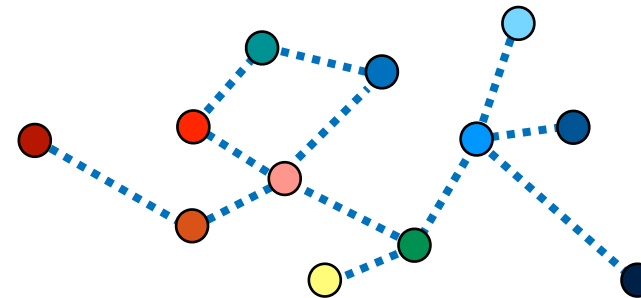
Frequency on the graph

- Global smoothness on the graph

$$y^T L y = \sum_{m \in \mathcal{V}} \sum_{n \in \mathcal{N}_m} W_{m,n} [y(m) - y(n)]^2$$



$$y^T \mathcal{L}_1 y = 0.15$$



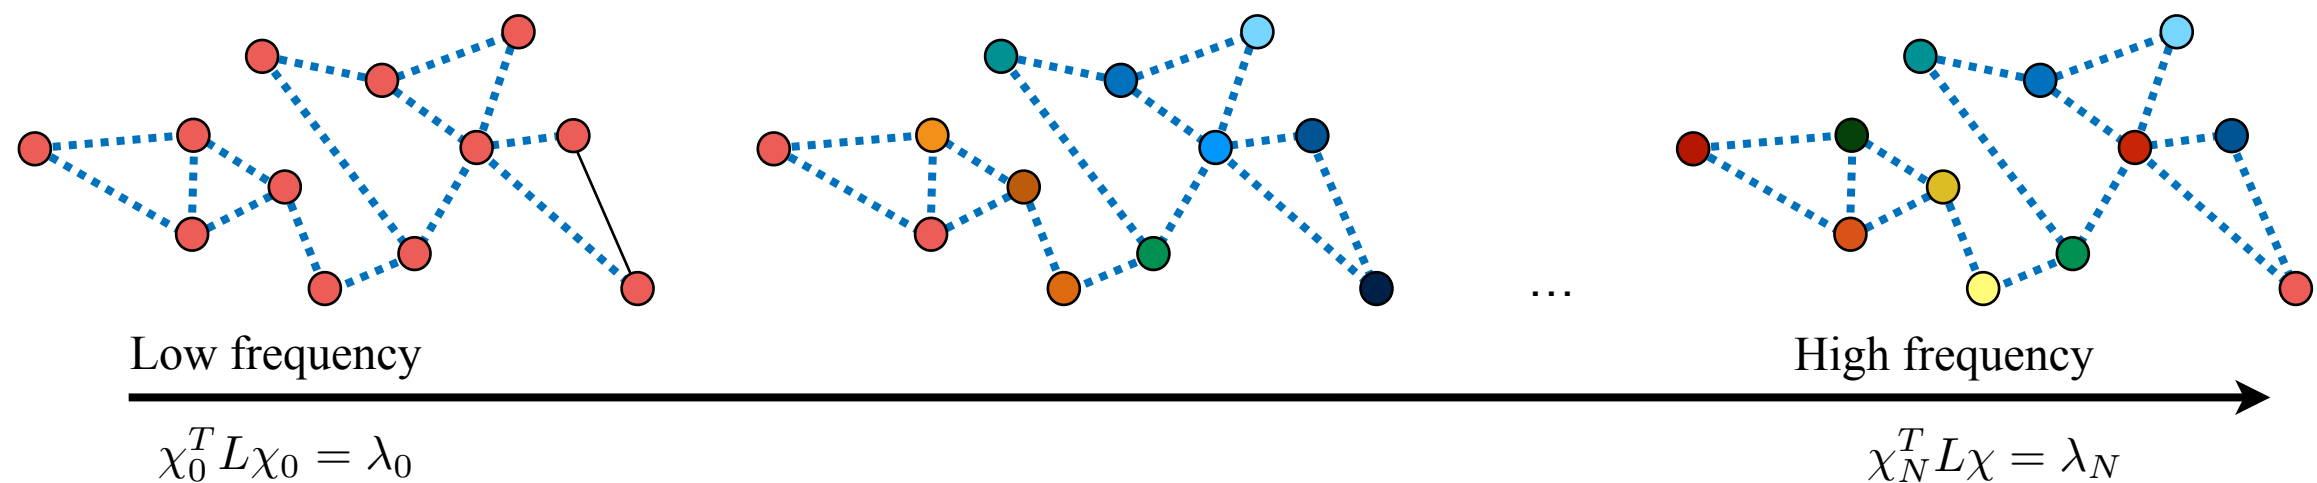
$$y^T \mathcal{L}_2 y = 1.8$$

Frequency on the graph

- Global smoothness on the graph

$$y^T L y = \sum_{m \in \mathcal{V}} \sum_{n \in \mathcal{N}_m} W_{m,n} [y(m) - y(n)]^2$$

- The eigenvectors of the Laplacian permit to define a harmonic (Fourier) analysis of graph signals

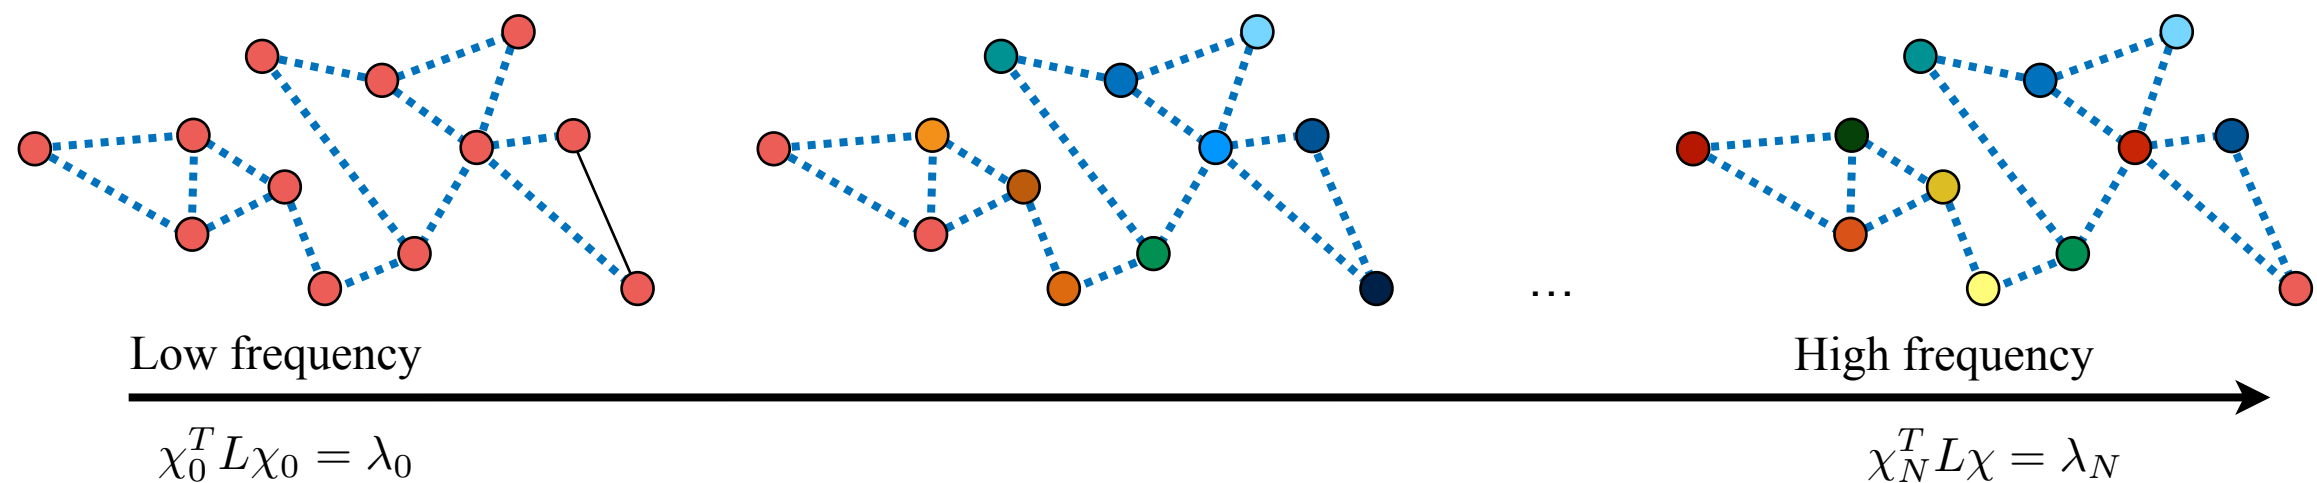


Frequency on the graph

- Global smoothness on the graph

$$y^T L y = \sum_{m \in \mathcal{V}} \sum_{n \in \mathcal{N}_m} W_{m,n} [y(m) - y(n)]^2$$

- The eigenvectors of the Laplacian permit to define a harmonic (Fourier) analysis of graph signals



GFT $\hat{y}(\lambda_\ell) = \langle y, \chi_\ell \rangle = \sum_{n=1}^N y(n) \chi_\ell^*(n), \quad \ell = 0, 1, \dots, N-1$

Filtering on graphs

- Filtering in the spectral domain with a transfer function $\hat{g}(\cdot)$

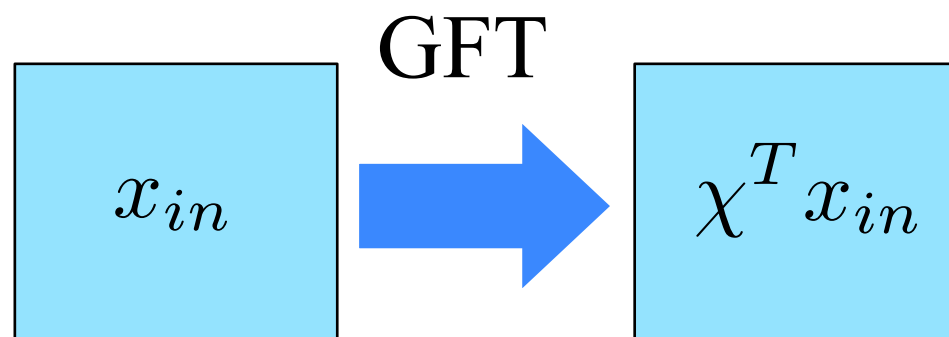
Filtering on graphs

- Filtering in the spectral domain with a transfer function $\hat{g}(\cdot)$



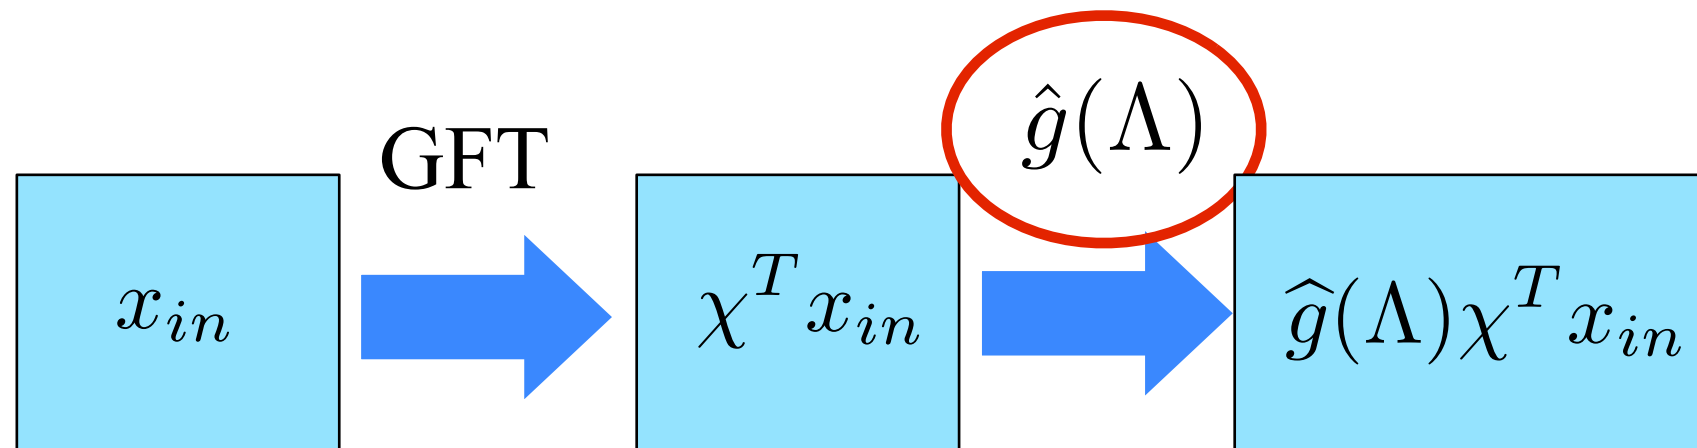
Filtering on graphs

- Filtering in the spectral domain with a transfer function $\hat{g}(\cdot)$



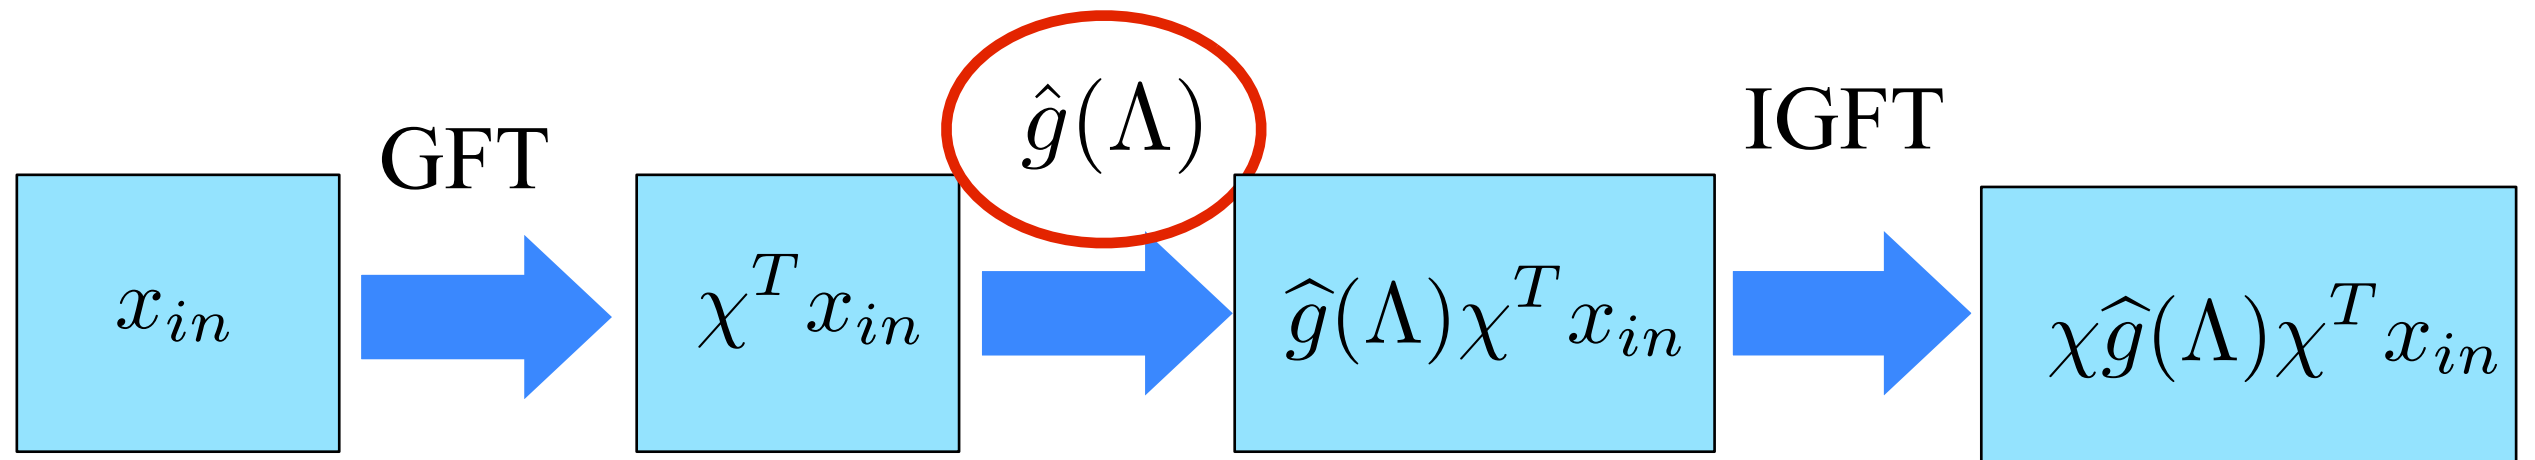
Filtering on graphs

- Filtering in the spectral domain with a transfer function $\hat{g}(\cdot)$



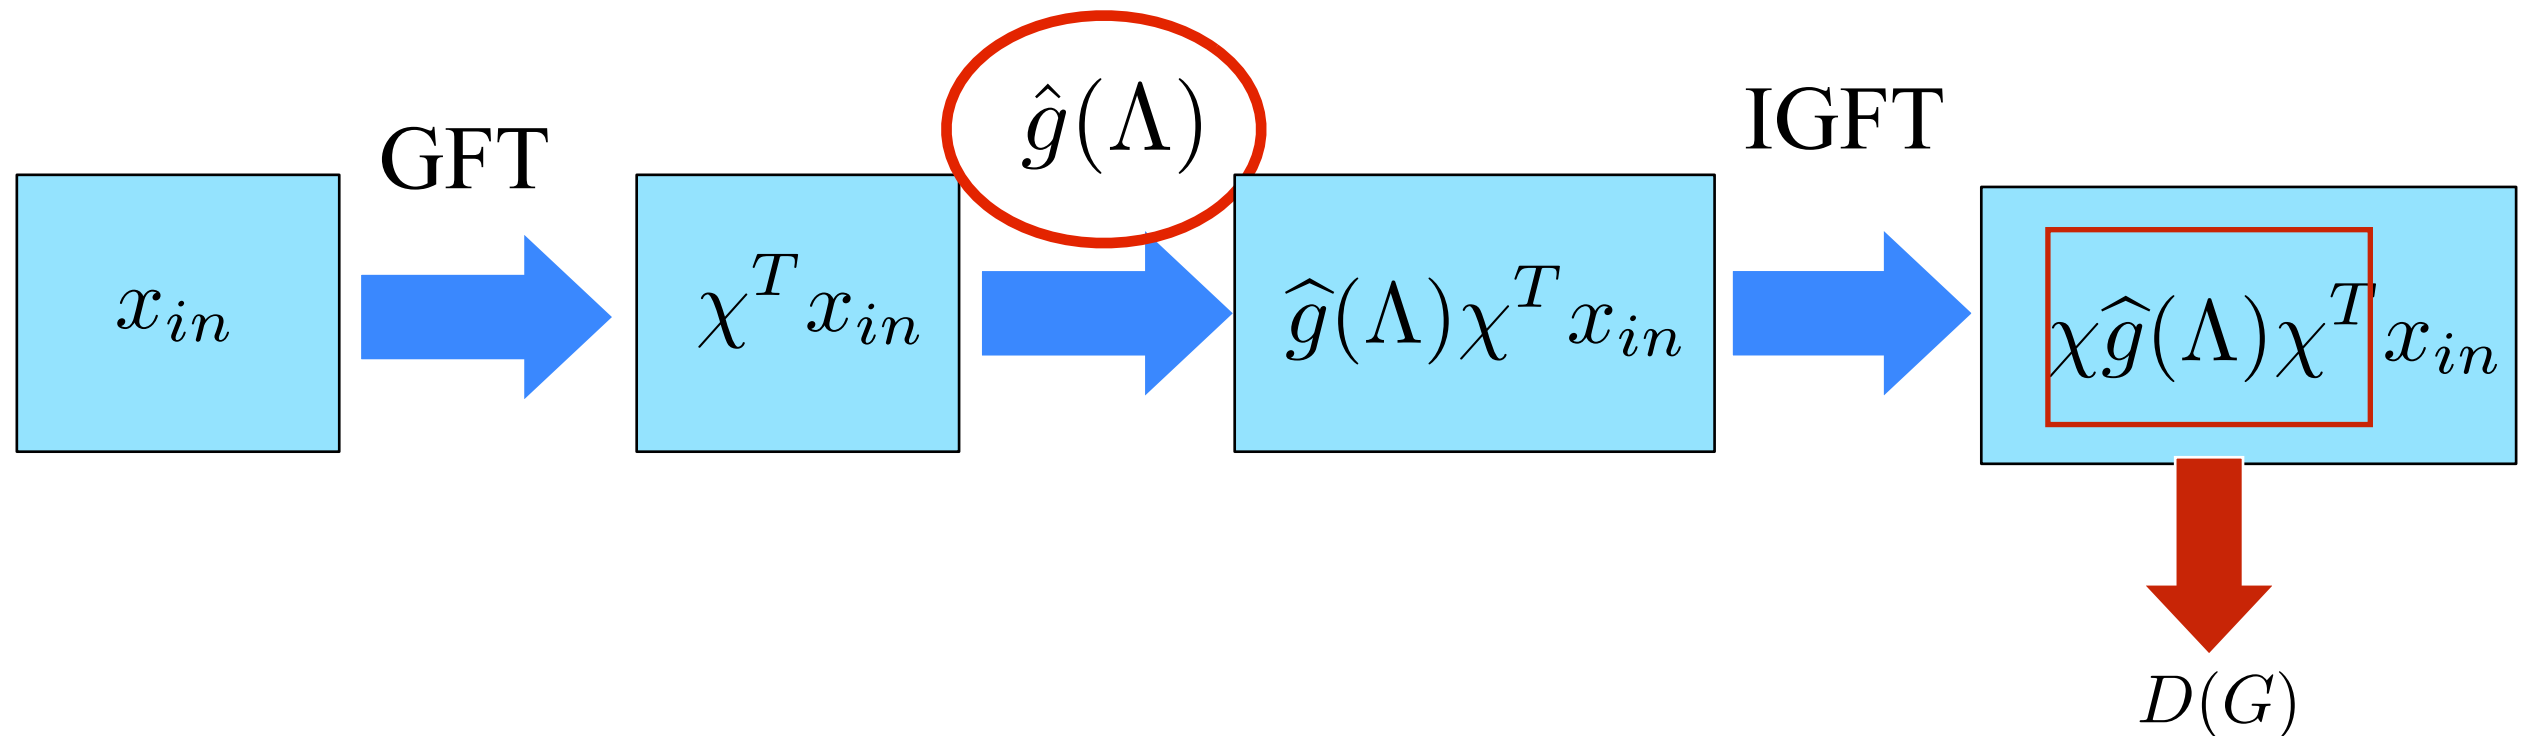
Filtering on graphs

- Filtering in the spectral domain with a transfer function $\hat{g}(\cdot)$



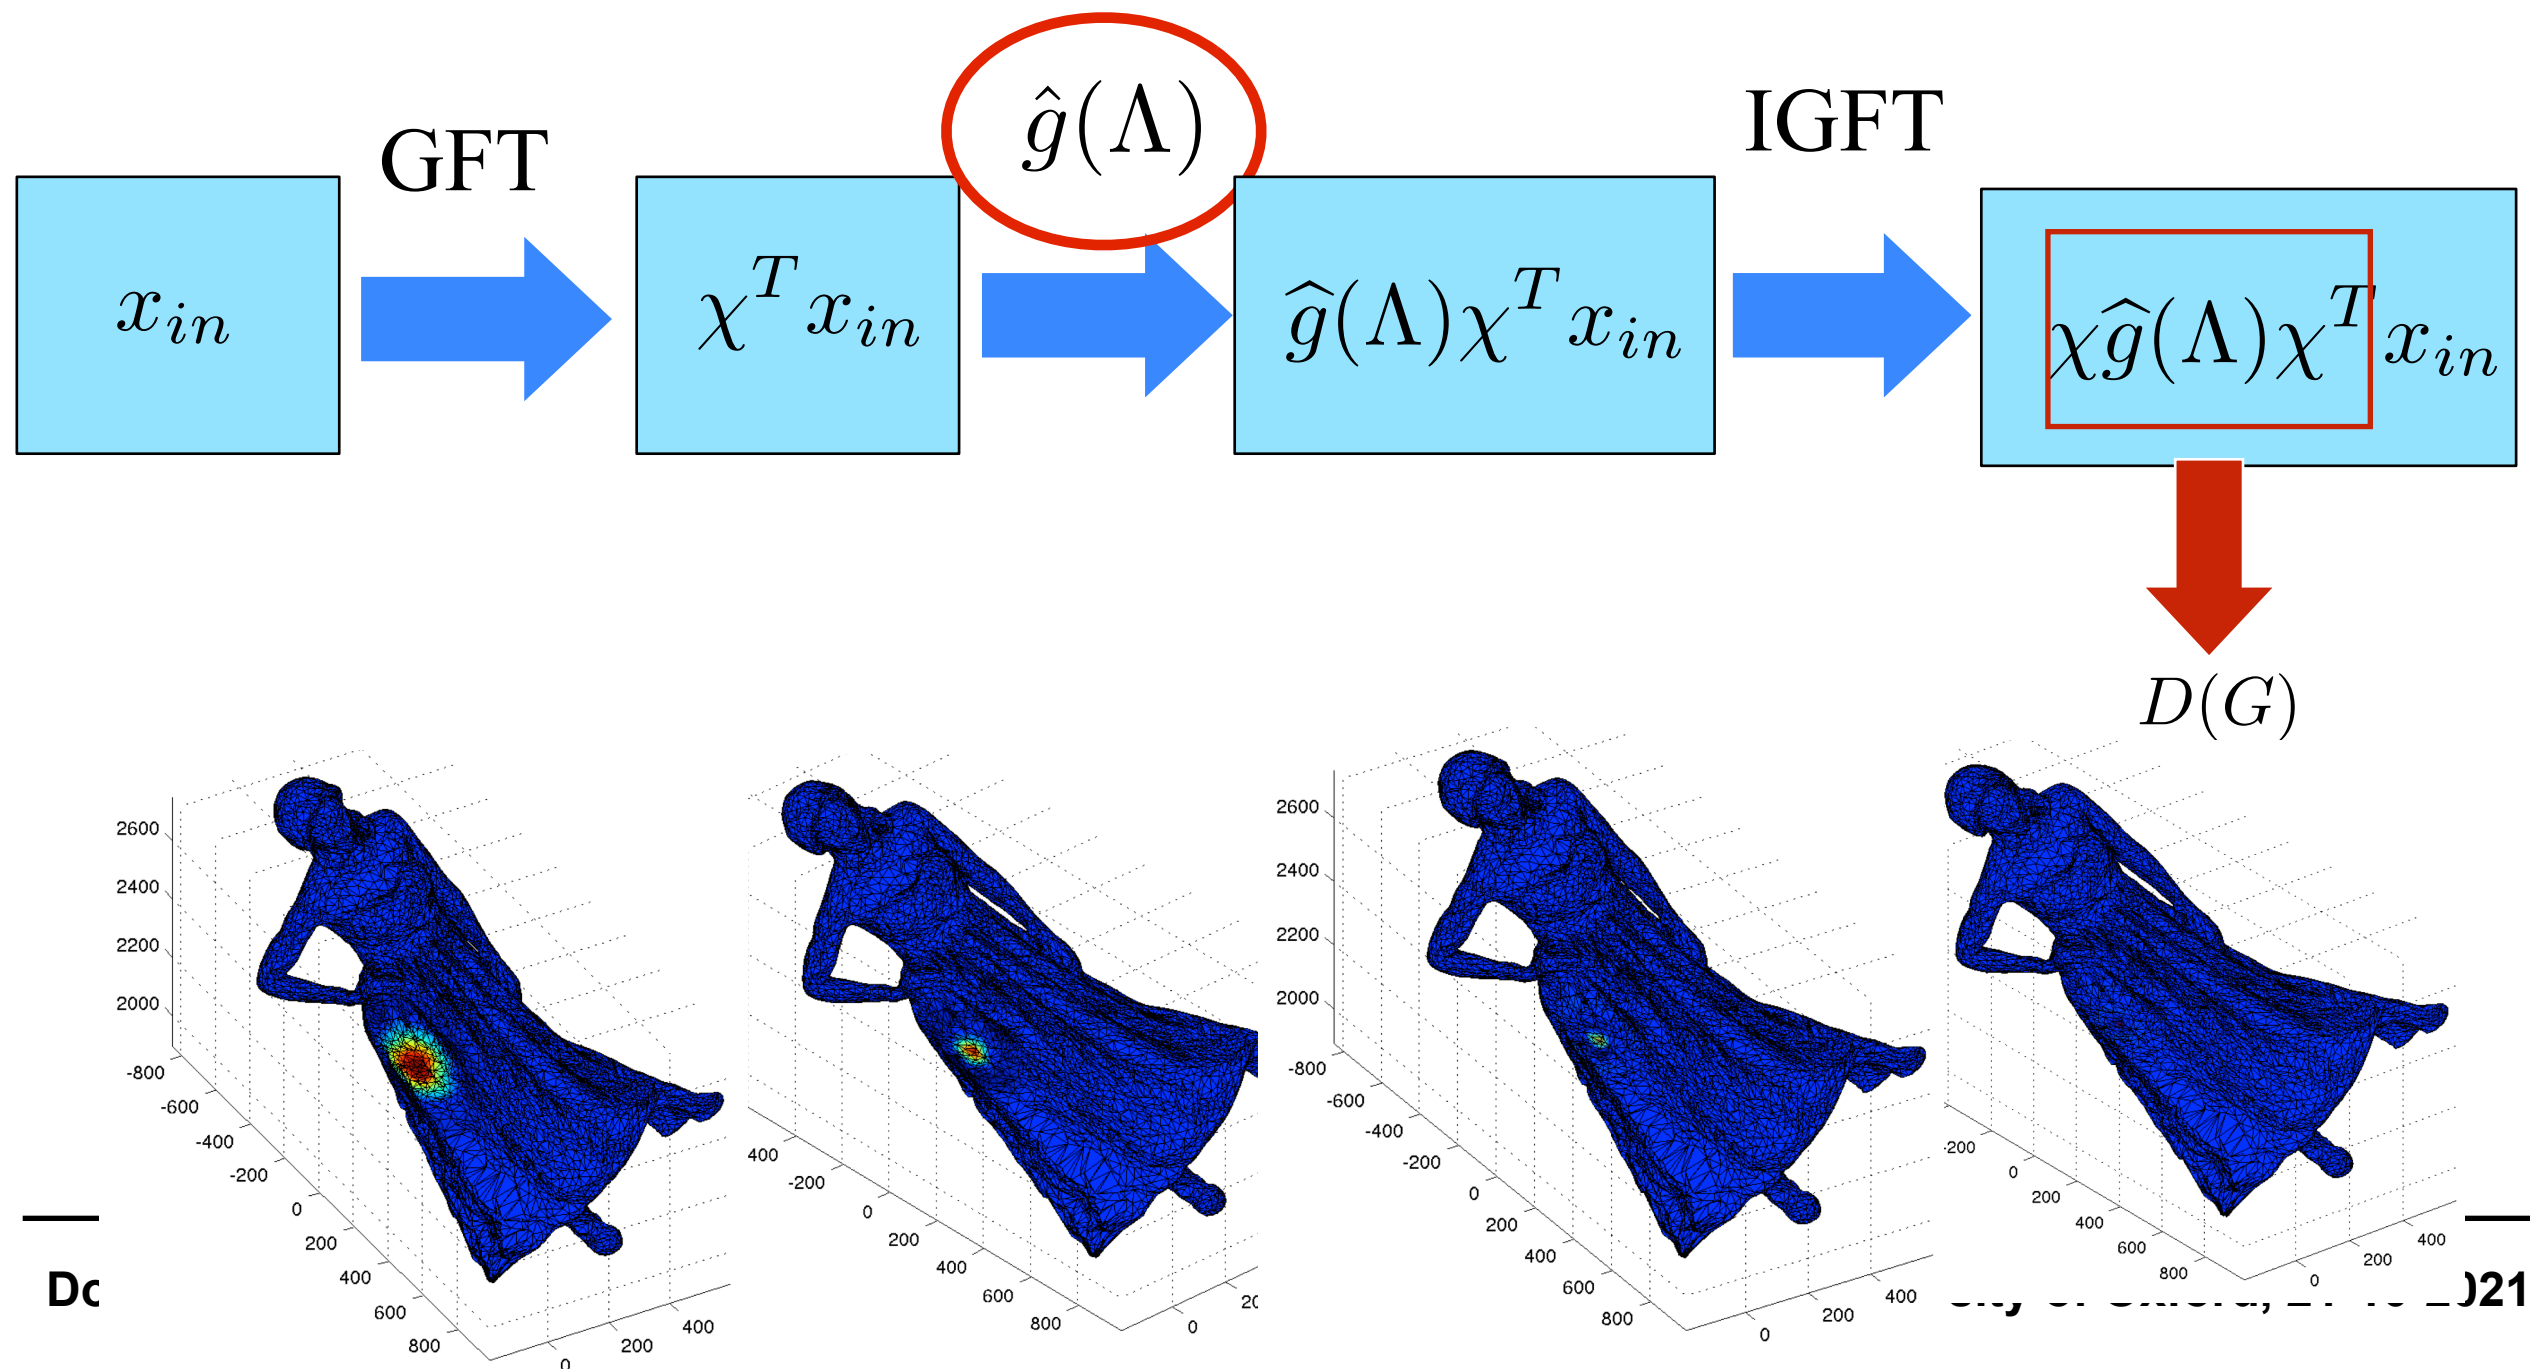
Filtering on graphs

- Filtering in the spectral domain with a transfer function $\hat{g}(\cdot)$



Filtering on graphs

- Filtering in the spectral domain with a transfer function $\hat{g}(\cdot)$



Spectral convolution on graphs

Classical convolution

$$(f * g)(t) = \int_{-\infty}^{\infty} f(t - \tau)g(\tau)d\tau$$

Convolution on graphs

Spectral convolution on graphs

Classical convolution

$$(f * g)(t) = \int_{-\infty}^{\infty} f(t - \tau)g(\tau)d\tau$$



Convolution on graphs

Spectral convolution on graphs

Classical convolution

$$(f * g)(t) = \int_{-\infty}^{\infty} f(t - \tau)g(\tau)d\tau$$



$$\widehat{(f * g)}(\omega) = \hat{f}(\omega) \cdot \hat{g}(\omega)$$

Convolution on graphs

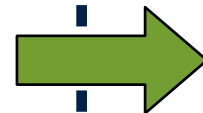
Spectral convolution on graphs

Classical convolution

$$(f * g)(t) = \int_{-\infty}^{\infty} f(t - \tau)g(\tau)d\tau$$



$$\widehat{(f * g)}(\omega) = \hat{f}(\omega) \cdot \hat{g}(\omega)$$



Convolution on graphs

Spectral convolution on graphs

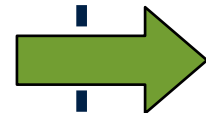
Classical convolution

$$(f * g)(t) = \int_{-\infty}^{\infty} f(t - \tau)g(\tau)d\tau$$



$$\widehat{(f * g)}(\omega) = \hat{f}(\omega) \cdot \hat{g}(\omega)$$

Convolution on graphs



$$\widehat{(f * g)}(\lambda) = ((\chi^T f) \circ \hat{g})(\lambda)$$

Spectral convolution on graphs

Classical convolution

$$(f * g)(t) = \int_{-\infty}^{\infty} f(t - \tau)g(\tau)d\tau$$

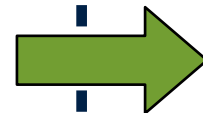


$$\widehat{(f * g)}(\omega) = \hat{f}(\omega) \cdot \hat{g}(\omega)$$

Convolution on graphs



$$\widehat{(f * g)}(\lambda) = ((\chi^T f) \circ \hat{g})(\lambda)$$



Spectral convolution on graphs

Classical convolution

$$(f * g)(t) = \int_{-\infty}^{\infty} f(t - \tau)g(\tau)d\tau$$



$$\widehat{(f * g)}(\omega) = \hat{f}(\omega) \cdot \hat{g}(\omega)$$

Convolution on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$



$$\widehat{(f * g)}(\lambda) = ((\chi^T f) \circ \hat{g})(\lambda)$$



K-hops localization on graphs

K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

$$\hat{g}_\theta(\lambda) = \sum_{j=0}^K \theta_j \lambda^j, \quad \theta \in \mathbb{R}^{K+1}$$

K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

$$\hat{g}_\theta(\lambda) = \sum_{j=0}^K \theta_j \lambda^j, \quad \theta \in \mathbb{R}^{K+1} \quad \rightarrow$$


K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

$$\hat{g}_\theta(\lambda) = \sum_{j=0}^K \theta_j \lambda^j, \quad \theta \in \mathbb{R}^{K+1} \quad \longrightarrow \quad \hat{g}_\theta(L) = \sum_{j=0}^K \theta_j L^j$$

K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

$$\hat{g}_\theta(\lambda) = \sum_{j=0}^K \theta_j \lambda^j, \quad \theta \in \mathbb{R}^{K+1} \quad \xrightarrow{\quad} \quad \hat{g}_\theta(L) = \sum_{j=0}^K \theta_j L^j$$


K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

$$\hat{g}_\theta(\lambda) = \sum_{j=0}^K \theta_j \lambda^j, \quad \theta \in \mathbb{R}^{K+1} \quad \longrightarrow \quad \hat{g}_\theta(L) = \sum_{j=0}^K \theta_j L^j$$

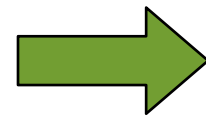


Localization within K-hop
neighborhood

K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

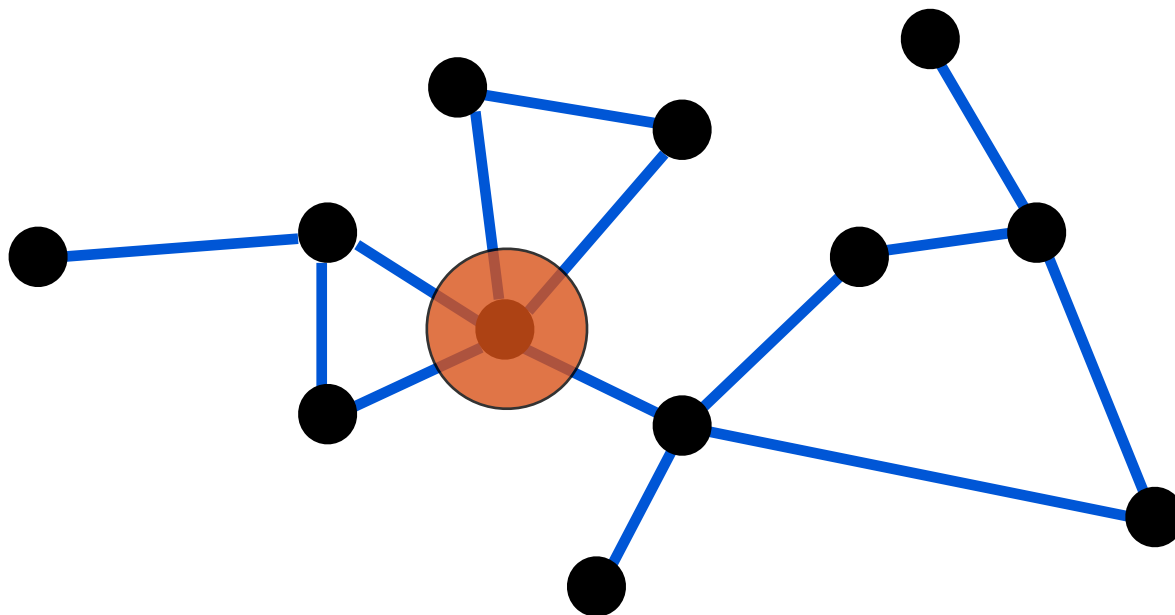
$$\hat{g}_\theta(\lambda) = \sum_{j=0}^K \theta_j \lambda^j, \quad \theta \in \mathbb{R}^{K+1}$$



$$\hat{g}_\theta(L) = \sum_{j=0}^K \theta_j L^j$$



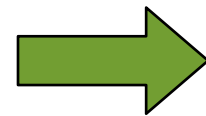
Localization within K-hop neighborhood



K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

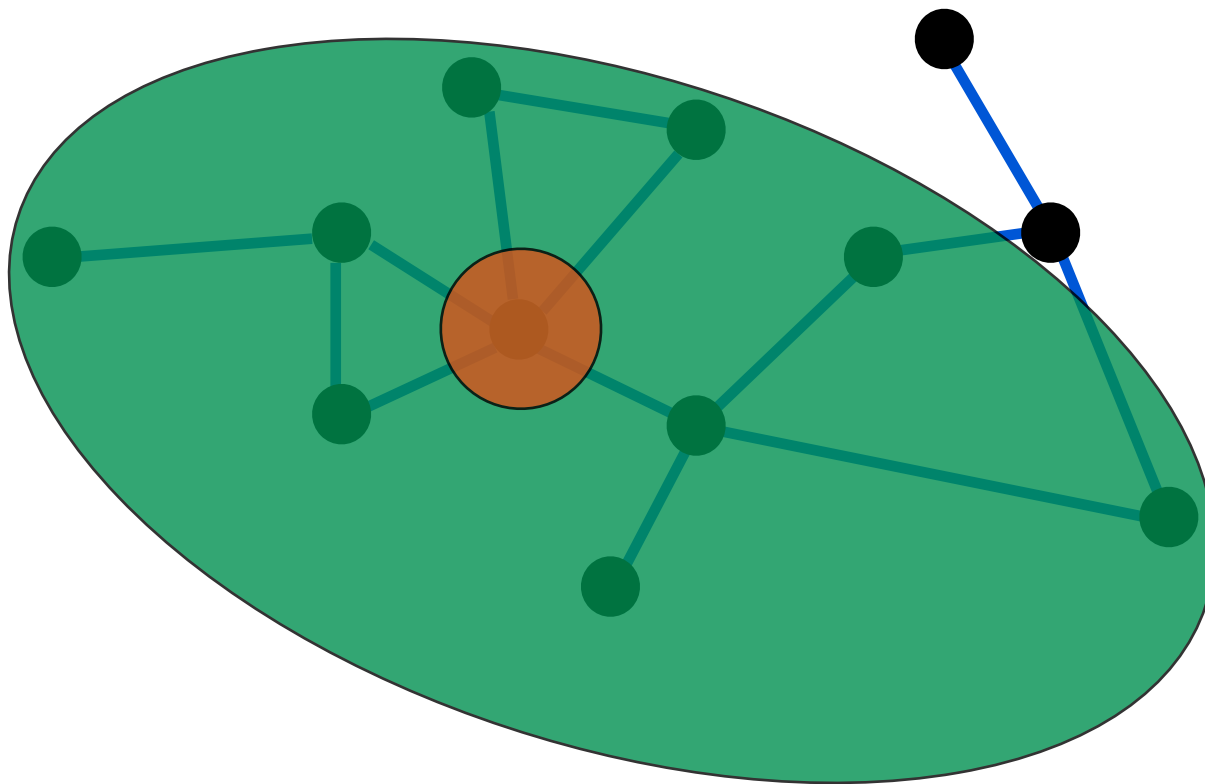
$$\hat{g}_\theta(\lambda) = \sum_{j=0}^K \theta_j \lambda^j, \quad \theta \in \mathbb{R}^{K+1}$$



$$\hat{g}_\theta(L) = \sum_{j=0}^K \theta_j L^j$$



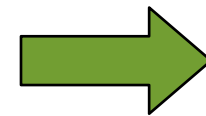
Localization within K-hop neighborhood



K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

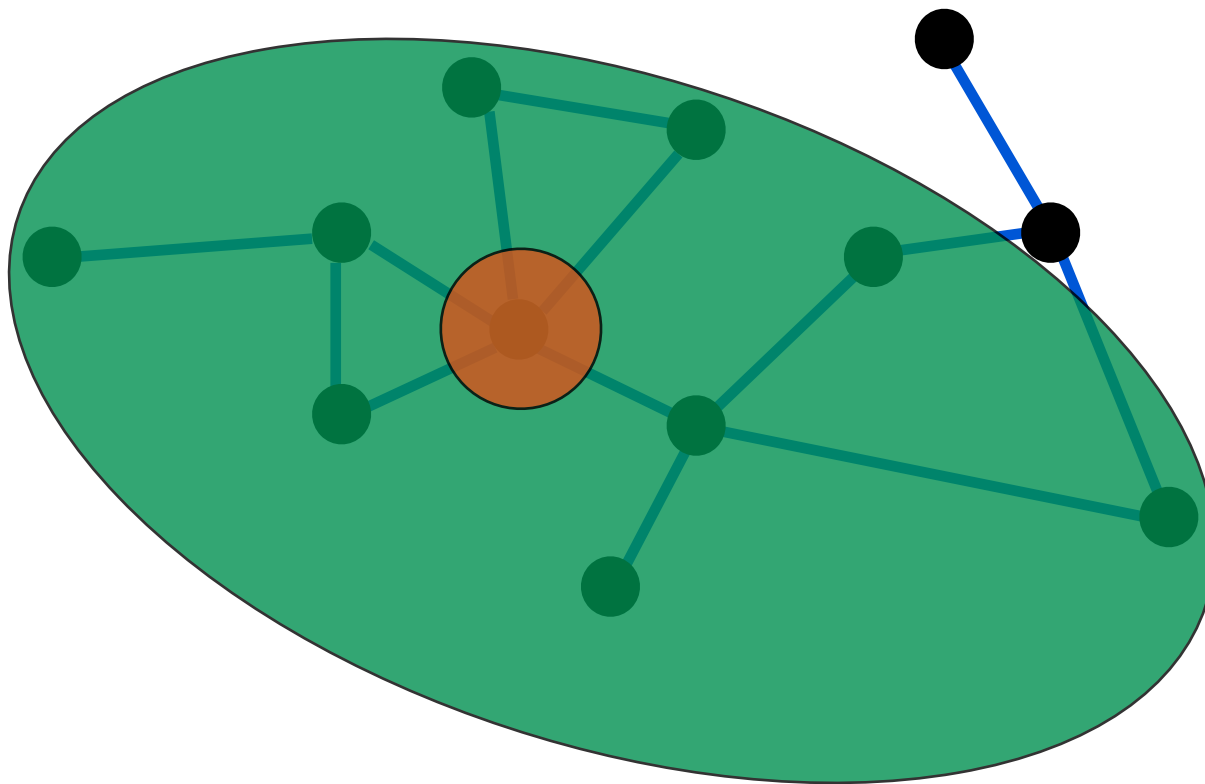
$$\hat{g}_\theta(\lambda) = \sum_{j=0}^K \theta_j \lambda^j, \quad \theta \in \mathbb{R}^{K+1}$$



$$\hat{g}_\theta(L) = \sum_{j=0}^K \theta_j L^j$$



Localization within K-hop neighborhood



- **Convolution on graphs:**
 - spectrally motivated
 - spatially implemented

K-hops localization on graphs

$$f * g = \chi \hat{g}(\Lambda) \chi^T f = \hat{g}(L) f$$

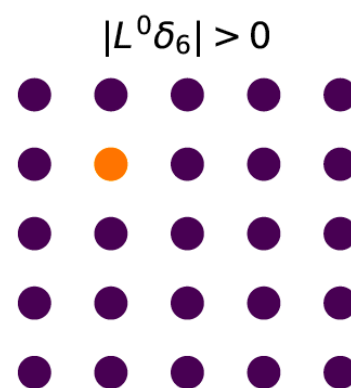
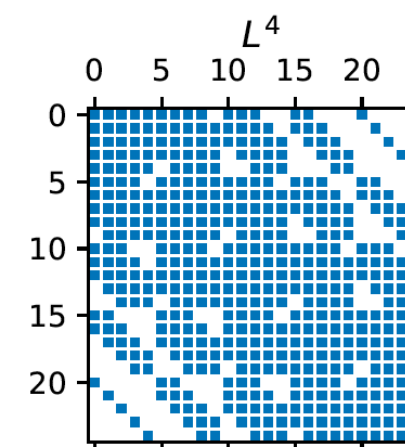
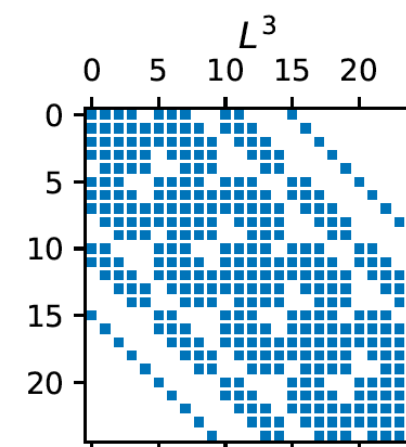
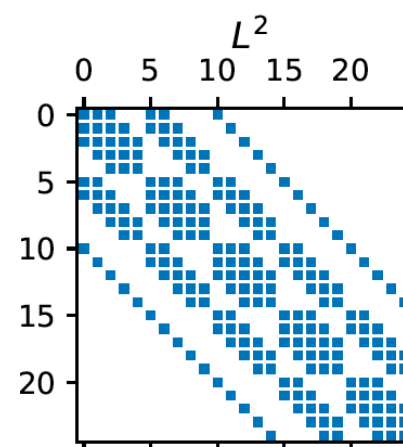
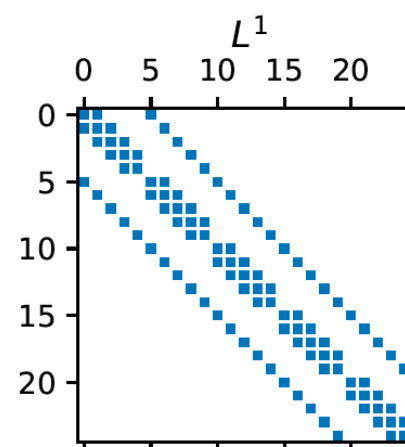
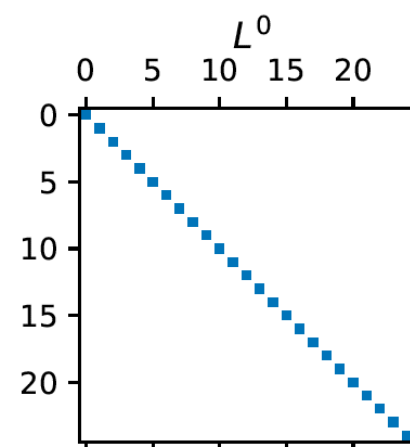
$$\hat{g}_\theta(\lambda) = \sum_{j=0}^K \theta_j \lambda^j, \quad \theta \in \mathbb{R}^{K+1} \quad \longrightarrow \quad \hat{g}_\theta(L) = \sum_{j=0}^K \theta_j L^j$$

Localization within K-hop neighborhood

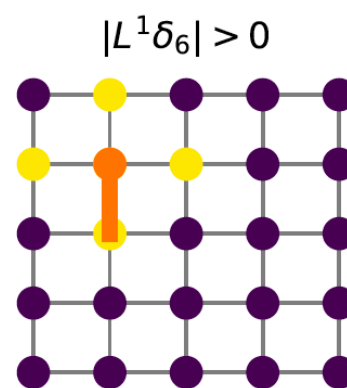
- **Convolution on graphs:**
 - spectrally motivated
 - spatially implemented

Powers of the graph Laplacian

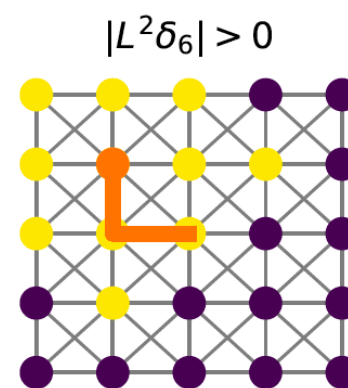
L^k defines the k -neighborhood



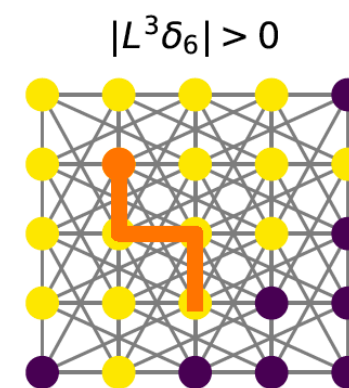
$\|W^0\|_0 = 0$ edges



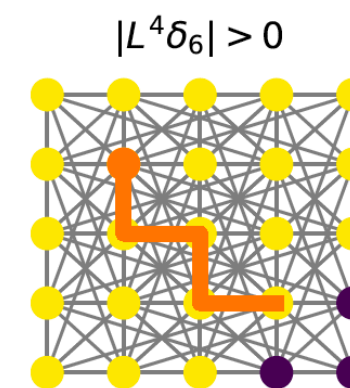
$\|W^1\|_0 = 40$ edges



$\|W^2\|_0 = 62$ edges



$\|W^3\|_0 = 108$ edges



$\|W^4\|_0 = 122$ edges

Localization: $d_G(v_i, v_j) > K$ implies $(L^K)_{ij} = 0$

Slide made by M. Deferrard

Spatial convolution on graphs

Classical convolution

Convolution on graphs

$$(f * g)(t) = \int_{-\infty}^{\infty} f(t - \tau) g(\tau) d\tau \quad \longrightarrow \quad (f * g)_n = \sum_{m \in \mathcal{V}} f(m) g(m, n)$$

function value at neighbors

- Defined as a weighted sum of function values at neighboring nodes

Spatial convolution on graphs

Classical convolution

Convolution on graphs

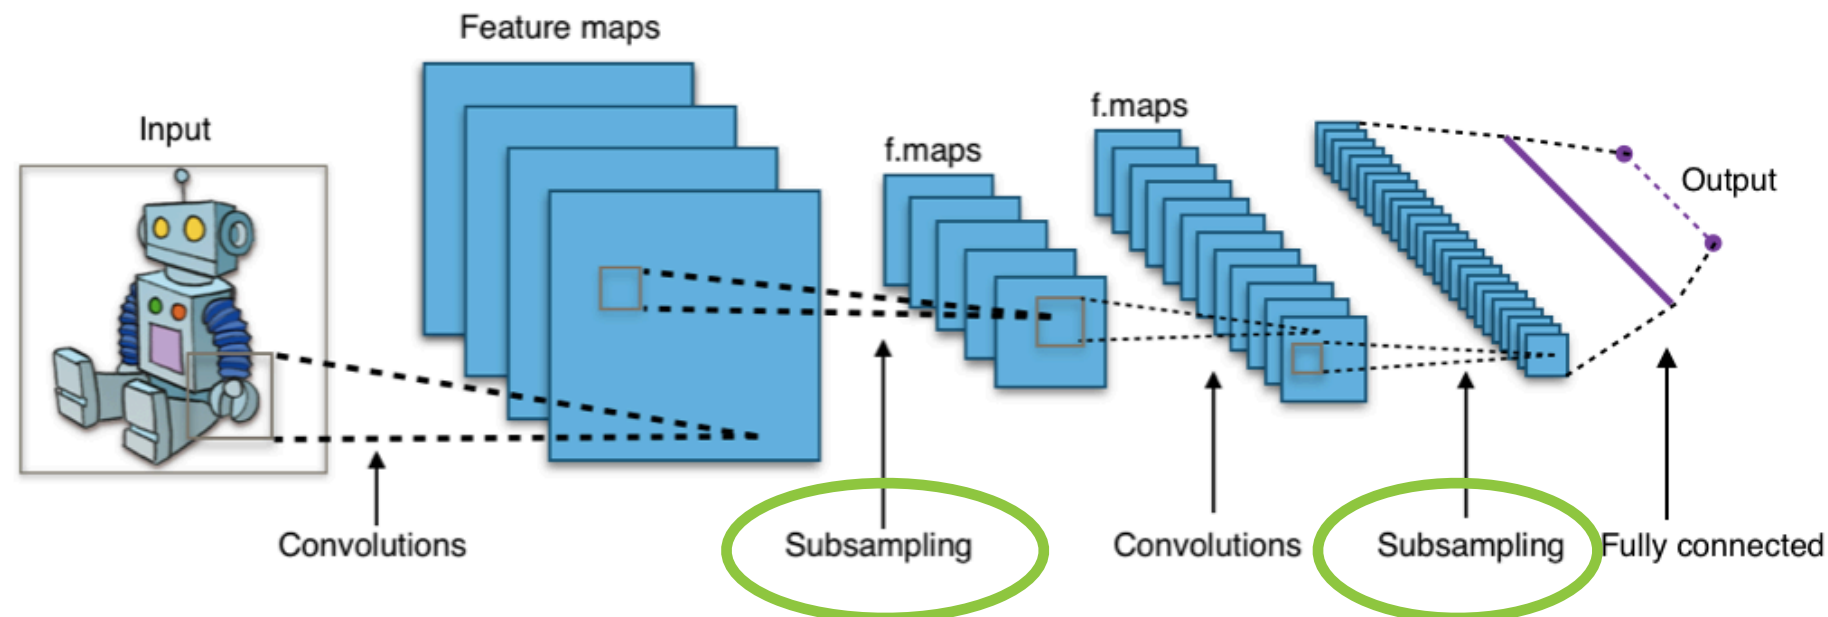
$$(f * g)(t) = \int_{-\infty}^{\infty} f(t - \tau) g(\tau) d\tau \quad \longrightarrow \quad (f * g)_n = \sum_{m \in \mathcal{V}} f(m) g(m, n)$$

function value at neighbors

weighting function/nodes' similarity

- Defined as a weighted sum of function values at neighboring nodes

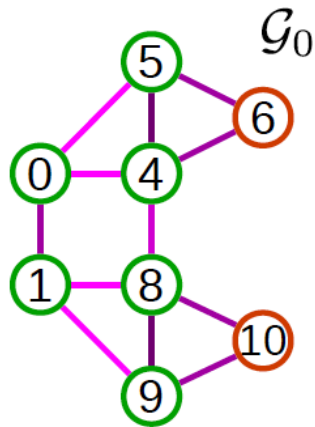
How to define pooling on graphs?



- A relatively open question, with ongoing research
- Methods can be grouped in three main categories:
 - topology based pooling
 - global pooling
 - hierarchical pooling

Topology based pooling

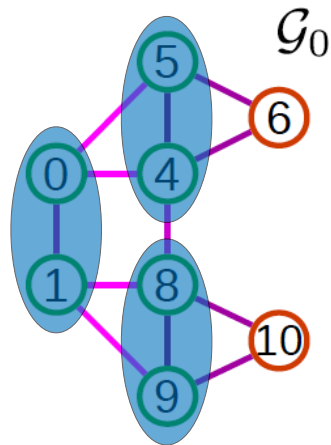
- Multi-scale graph coarsening: no features involved



- Graclus algorithm (Dhillon et al. 2007)

Topology based pooling

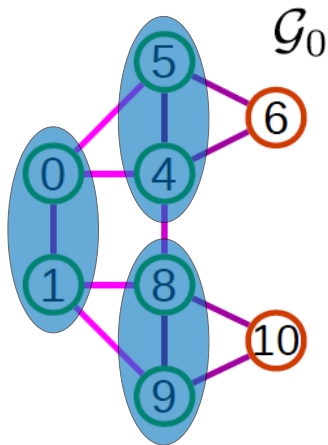
- Multi-scale graph coarsening: no features involved



- Graclus algorithm (Dhillon et al. 2007)

Topology based pooling

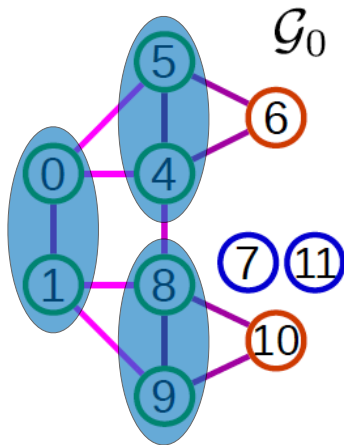
- Multi-scale graph coarsening: no features involved



- Graclus algorithm (Dhillon et al. 2007)
 - Local greedy way of merging vertices that minimises the normalised cut

Topology based pooling

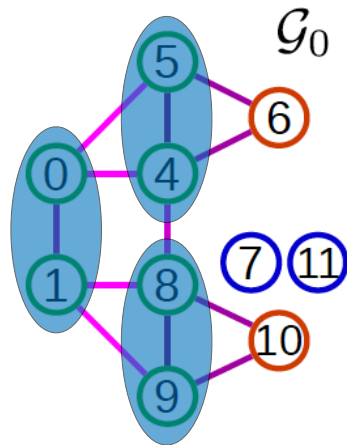
- Multi-scale graph coarsening: no features involved



- Graclus algorithm (Dhillon et al. 2007)
 - Local greedy way of merging vertices that minimises the normalised cut

Topology based pooling

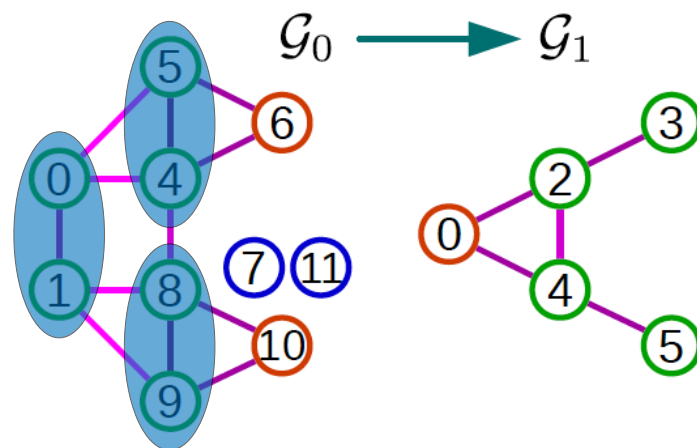
- Multi-scale graph coarsening: no features involved



- Graclus algorithm (Dhillon et al. 2007)
 - Local greedy way of merging vertices that minimises the normalised cut
 - Add artificial nodes to ensure two children for each vertex

Topology based pooling

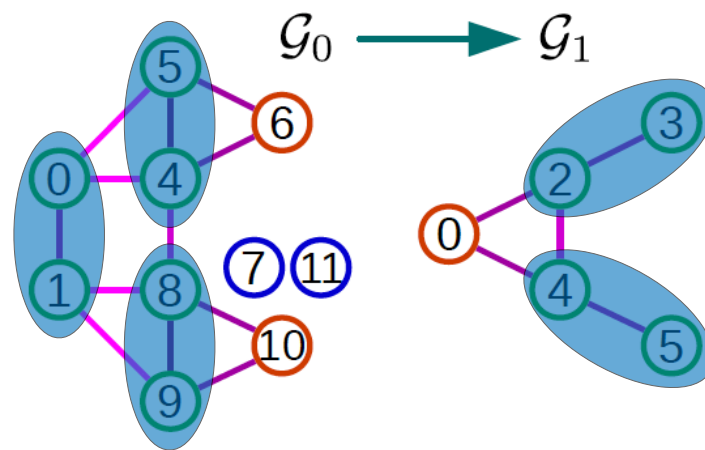
- Multi-scale graph coarsening: no features involved



- Graclus algorithm (Dhillon et al. 2007)
 - Local greedy way of merging vertices that minimises the normalised cut
 - Add artificial nodes to ensure two children for each vertex

Topology based pooling

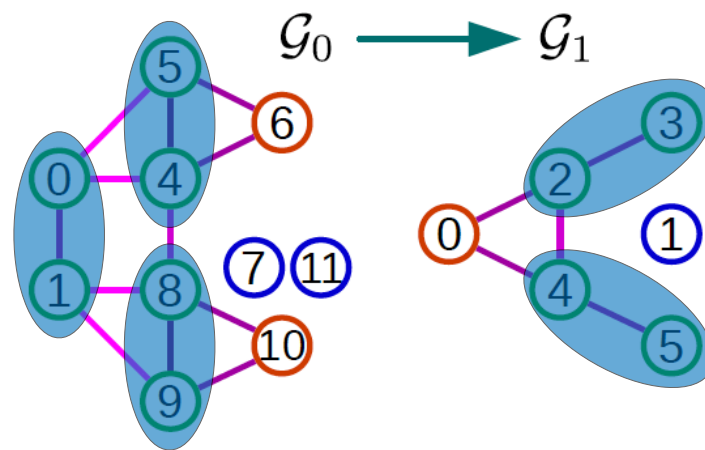
- Multi-scale graph coarsening: no features involved



- Graclus algorithm (Dhillon et al. 2007)
 - Local greedy way of merging vertices that minimises the normalised cut
 - Add artificial nodes to ensure two children for each vertex

Topology based pooling

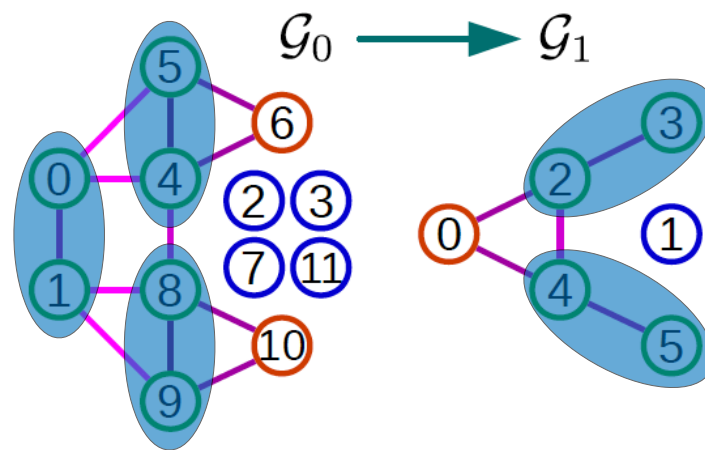
- Multi-scale graph coarsening: no features involved



- Graclus algorithm (Dhillon et al. 2007)
 - Local greedy way of merging vertices that minimises the normalised cut
 - Add artificial nodes to ensure two children for each vertex

Topology based pooling

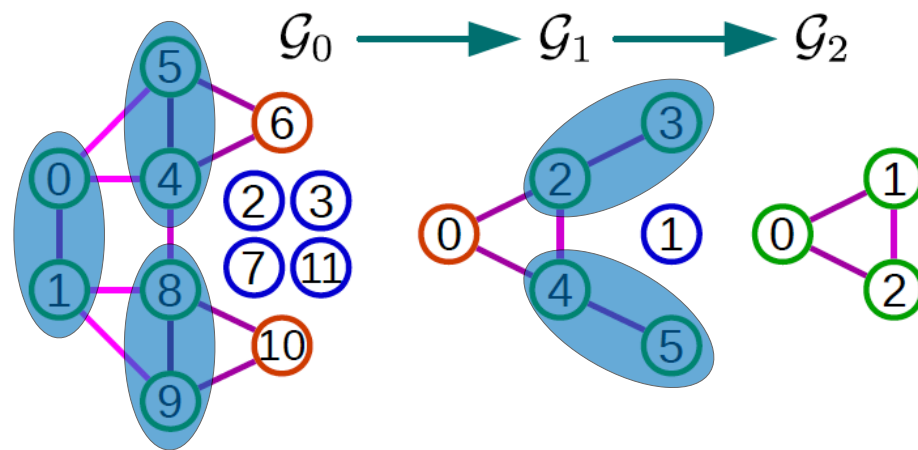
- Multi-scale graph coarsening: no features involved



- Grclus algorithm (Dhillon et al. 2007)
 - Local greedy way of merging vertices that minimises the normalised cut
 - Add artificial nodes to ensure two children for each vertex

Topology based pooling

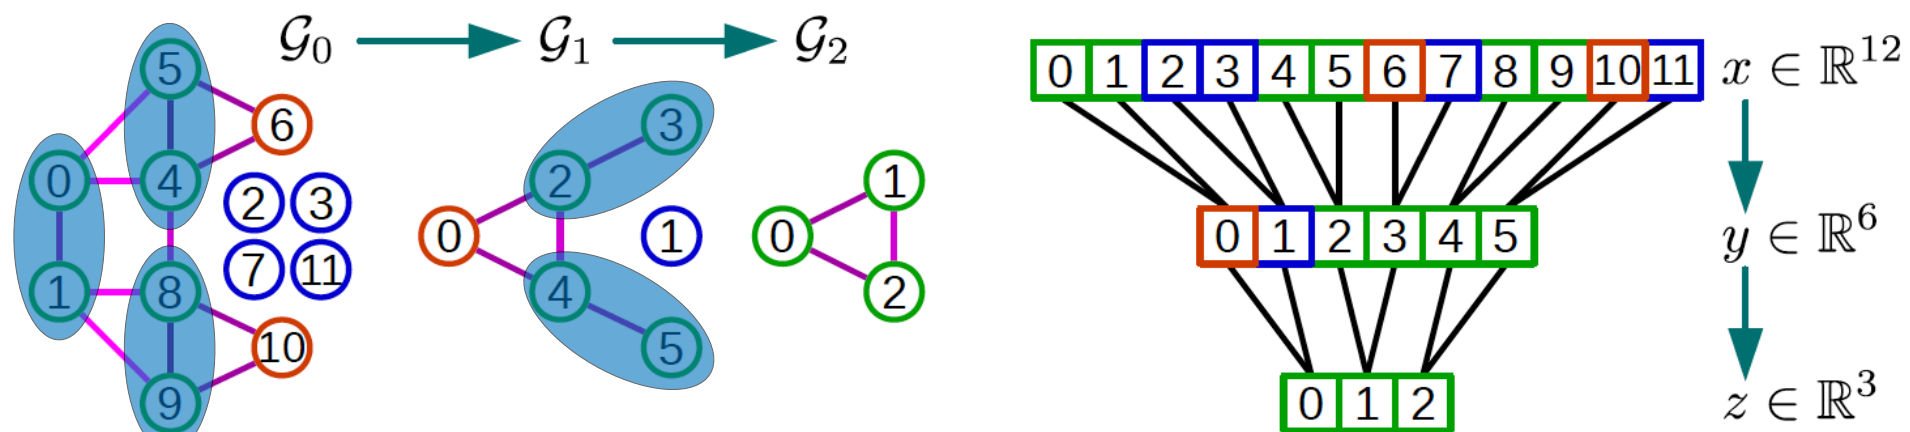
- Multi-scale graph coarsening: no features involved



- Grclus algorithm (Dhillon et al. 2007)
 - Local greedy way of merging vertices that minimises the normalised cut
 - Add artificial nodes to ensure two children for each vertex

Topology based pooling

- Multi-scale graph coarsening: no features involved



Defferrard et al. 2016

- Graculus algorithm (Dhillon et al. 2007)
 - Local greedy way of merging vertices that minimises the normalised cut
 - Add artificial nodes to ensure two children for each vertex
 - 1D grid pooling: $[\max(0,1), \max(4,5,6), \max(8,9,10)]$

Global pooling

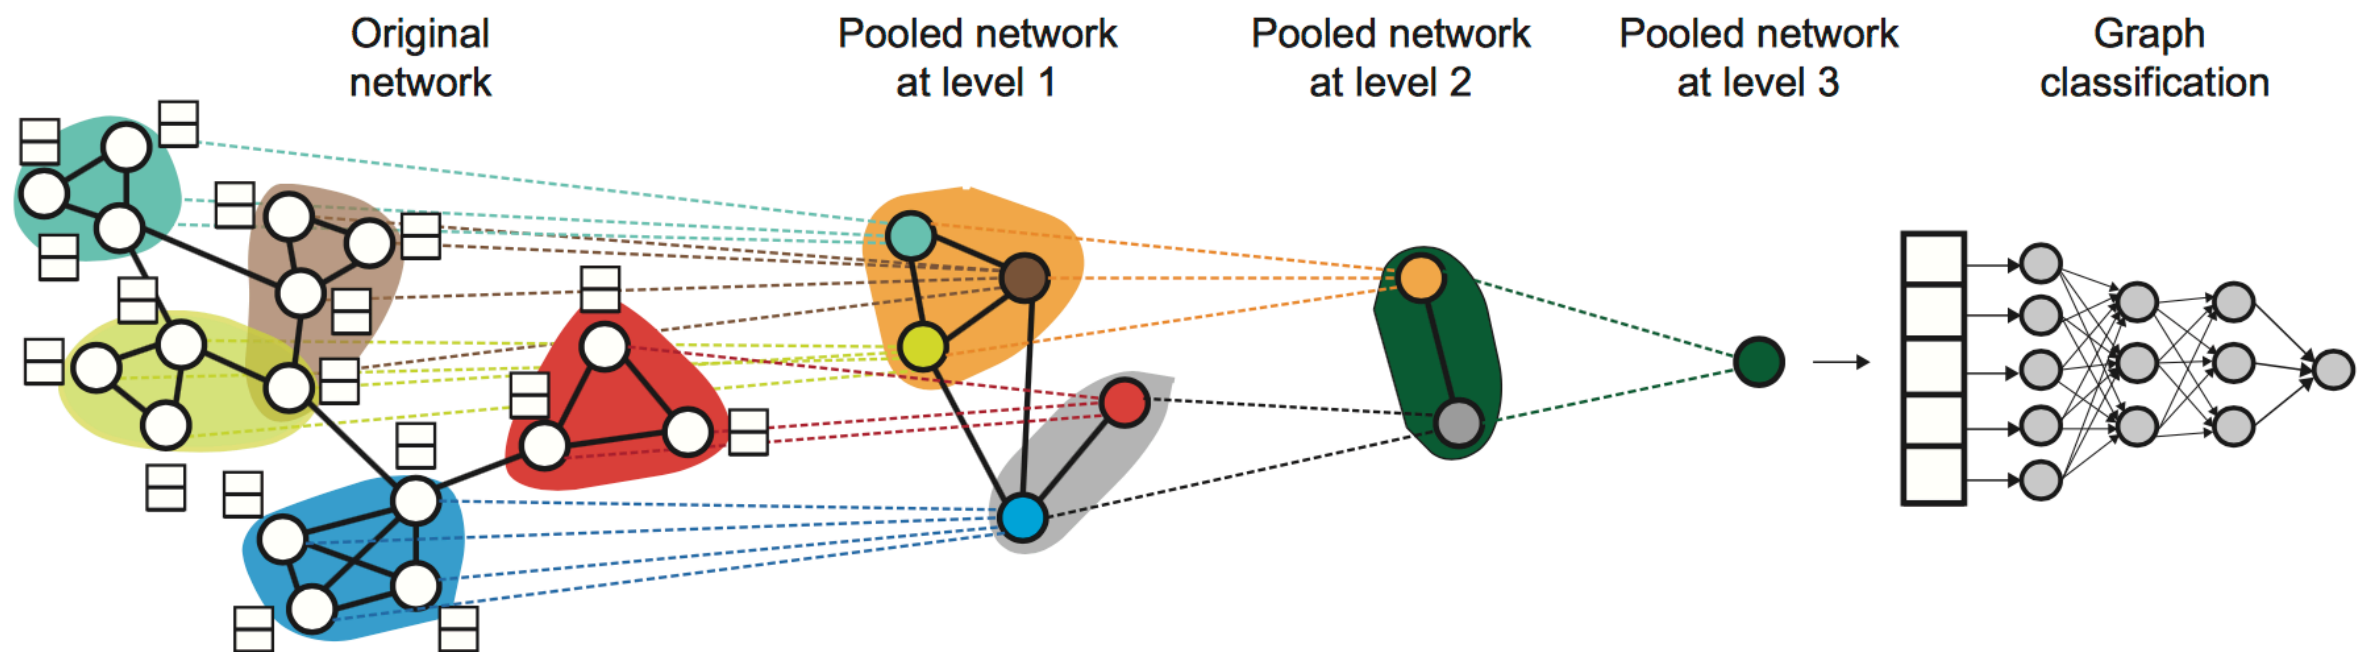
- Involves node features
- Uses sum/max or neural networks to pool all representation of nodes

$$h_G = \text{mean/max/sum}(h_1^{(K)}, h_2^{(K)}, \dots, h_N^{(K)})$$

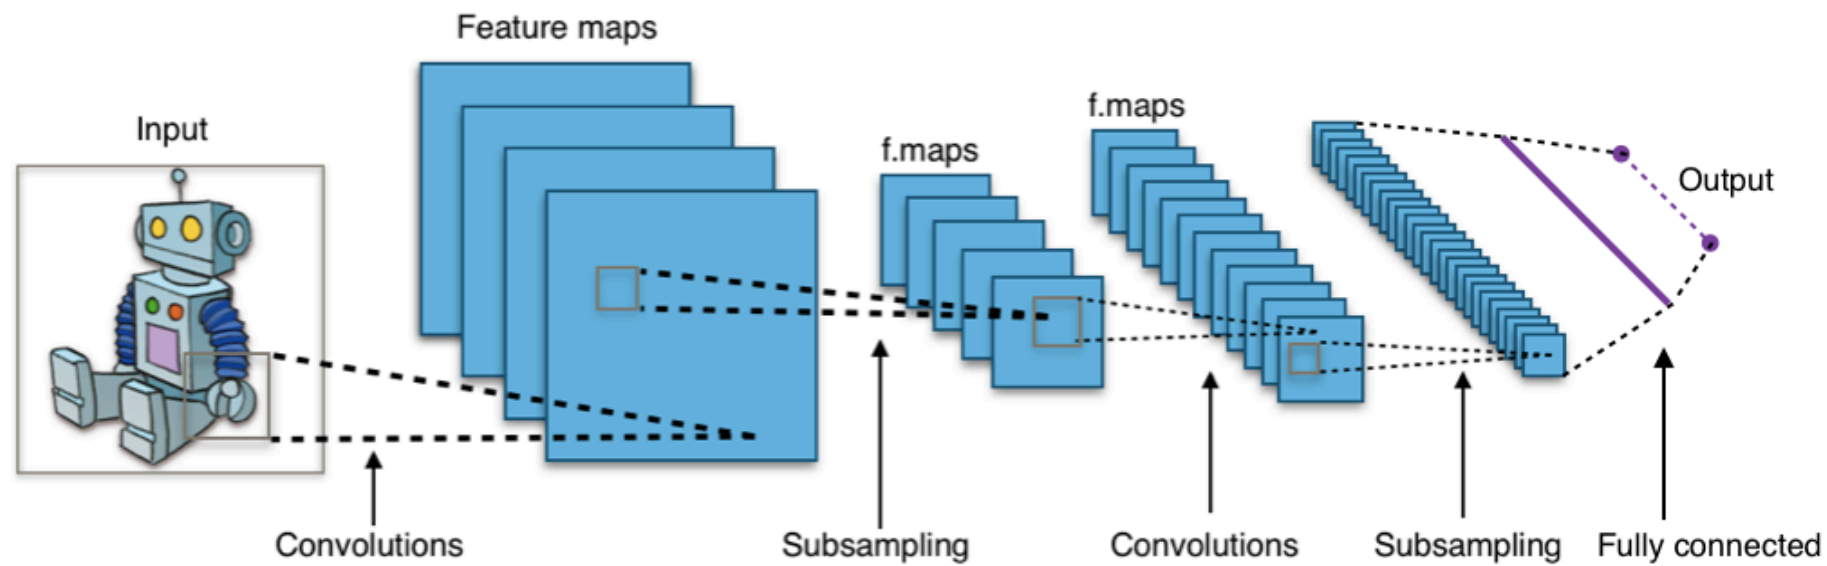
- Also known as READOUT
- Example: SortPool (Zhang et al. 2018)
 - sorts embeddings for nodes according to the structural roles of a graph

Hierarchical pooling

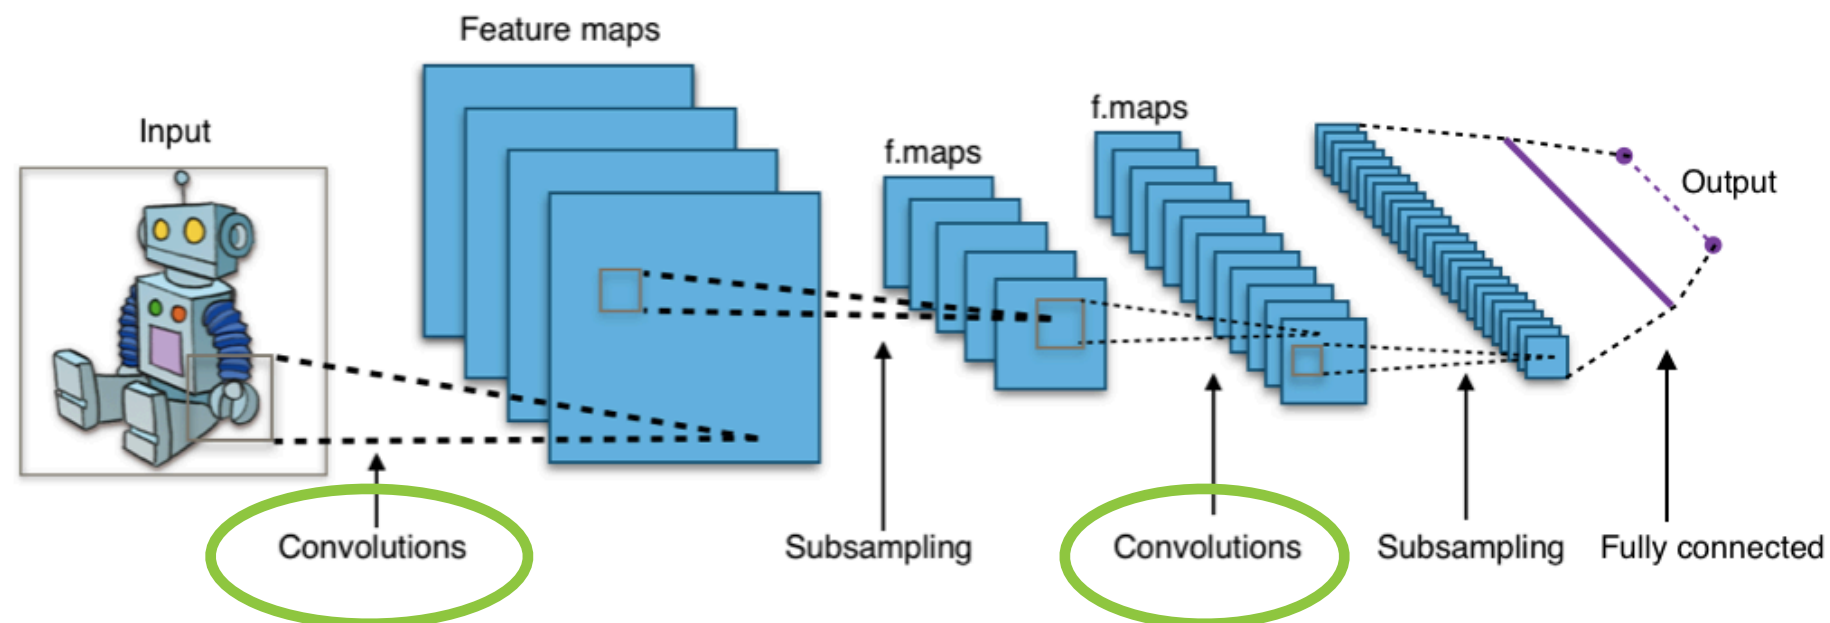
- Involved nodes features
- Aggregate information in a hierarchical way that respects the graph structure
- Results in cluster selection
- Example: DiffPool [Ying et al. 2019]



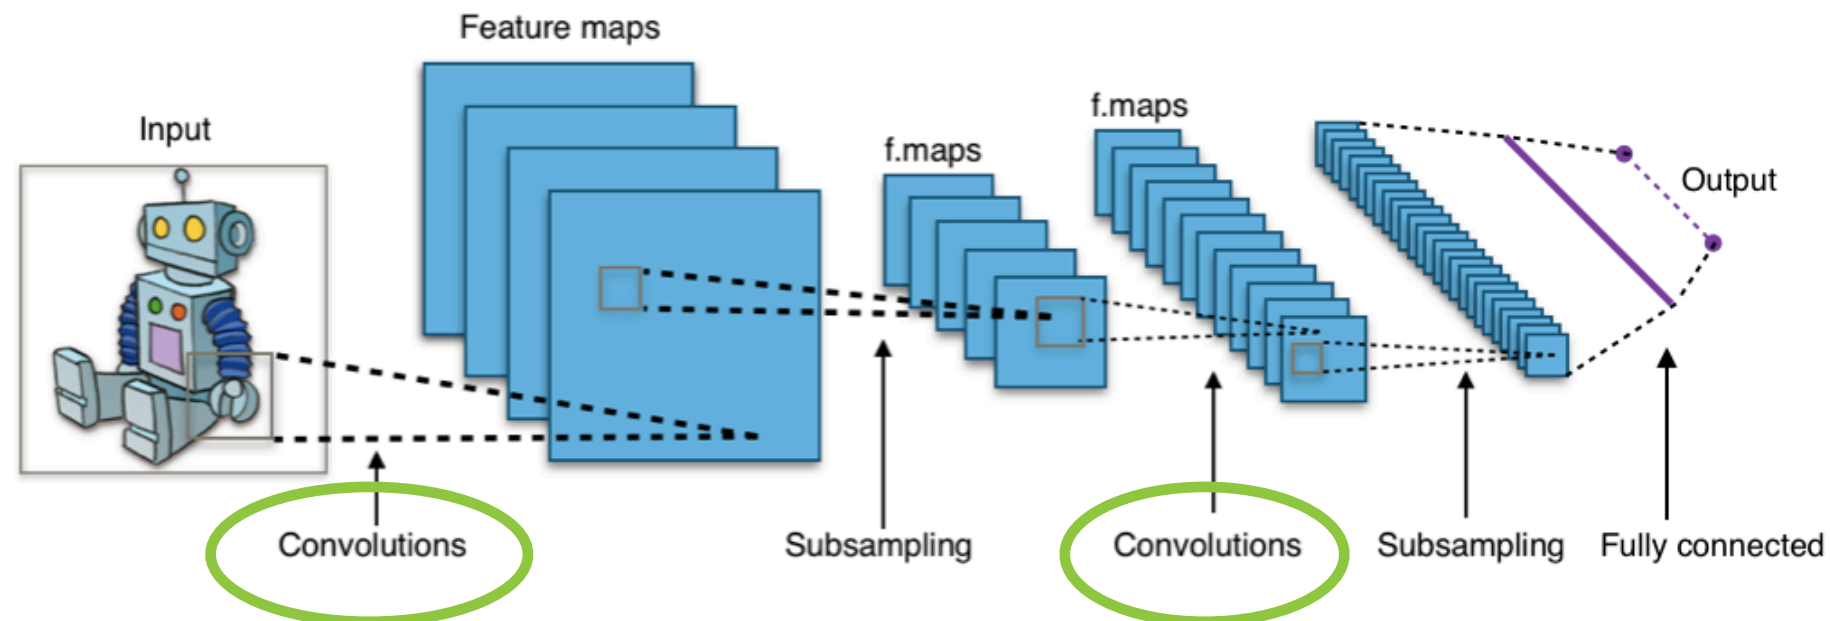
Summary so far



Summary so far

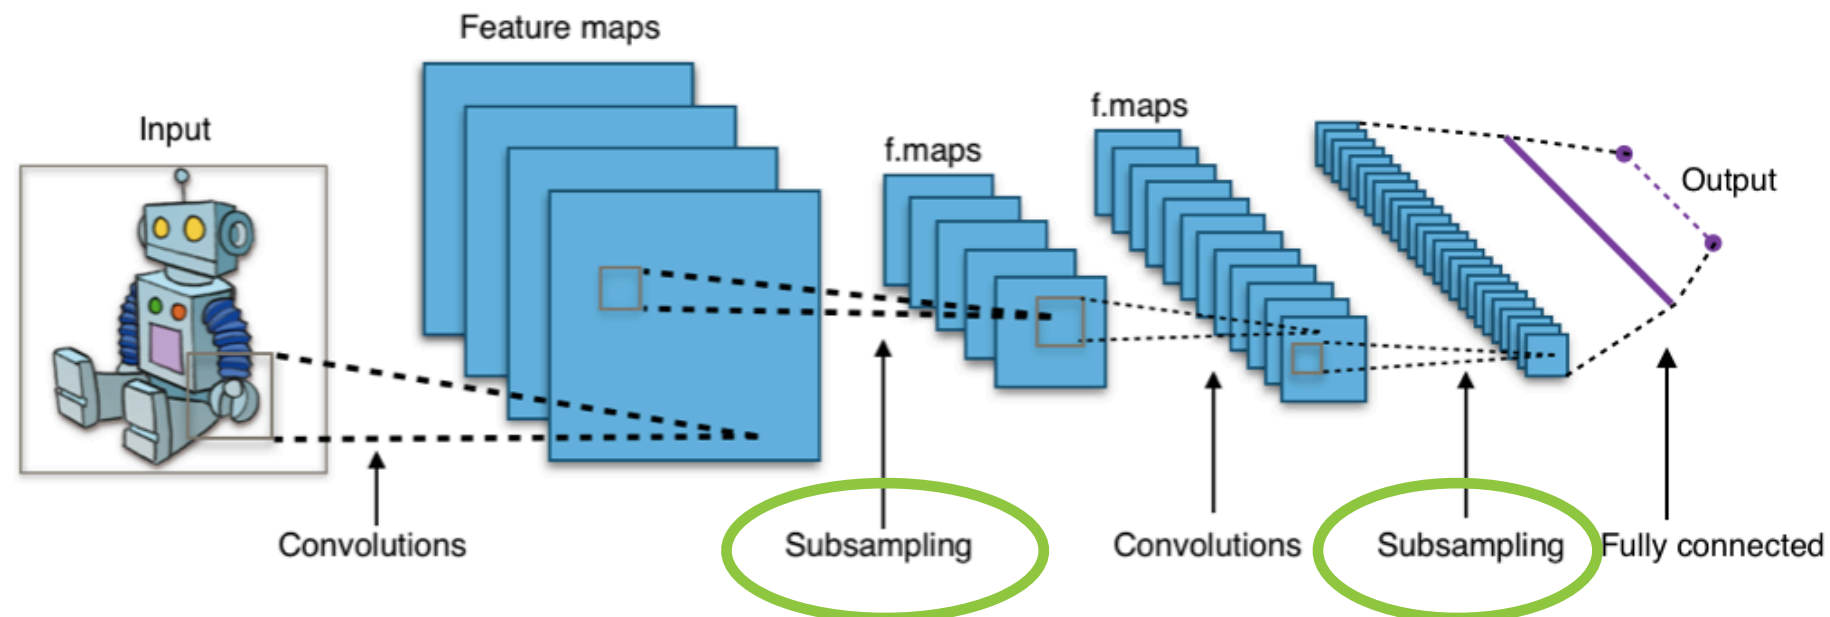


Summary so far



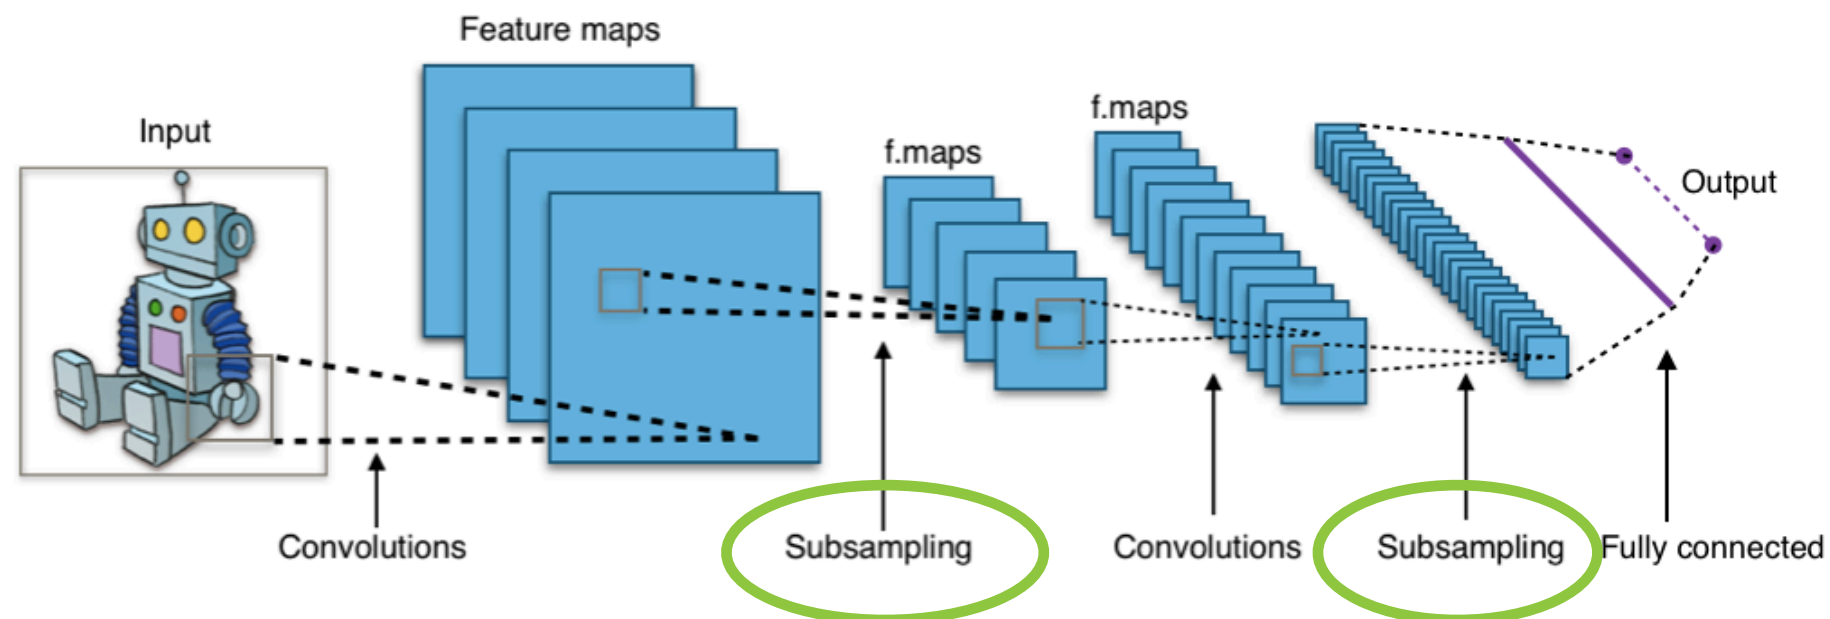
- Graph convolution: spectral or spatial

Summary so far



- Graph convolution: spectral or spatial

Summary so far

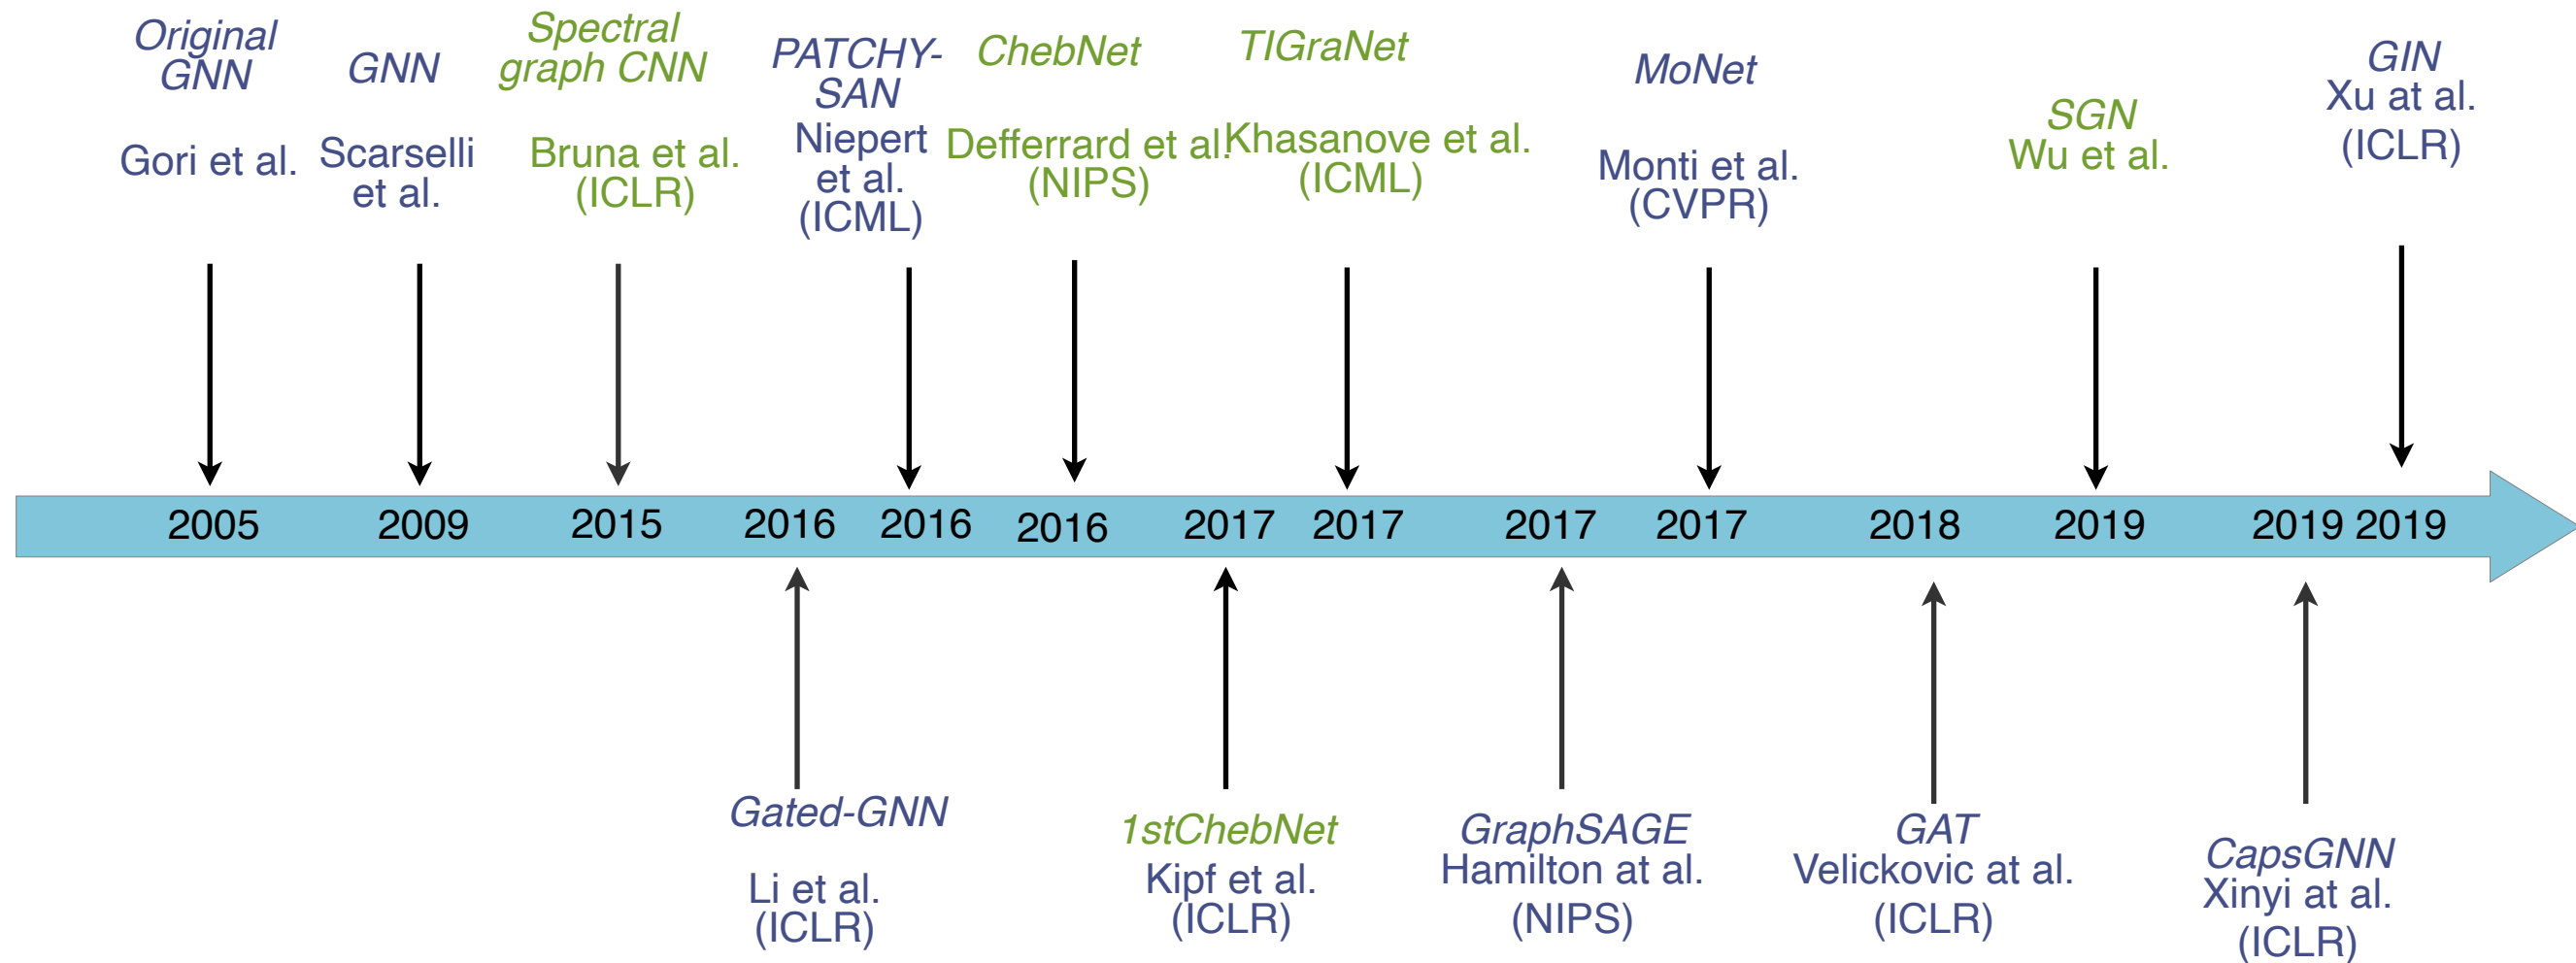


- Graph convolution: spectral or spatial
- Graph pooling: structural, global, or hierarchical

Outline

- Motivation: Why Graph Neural Networks?
- Basic definitions on graphs
- Partial historical overview
 - **Graph CNN (main focus)**
 - Graph autoencoders (briefly)
- Applications
 - Naturally graph-structured data
 - Images

Partial historical overview



Spatial-based methods

Spectral-based methods

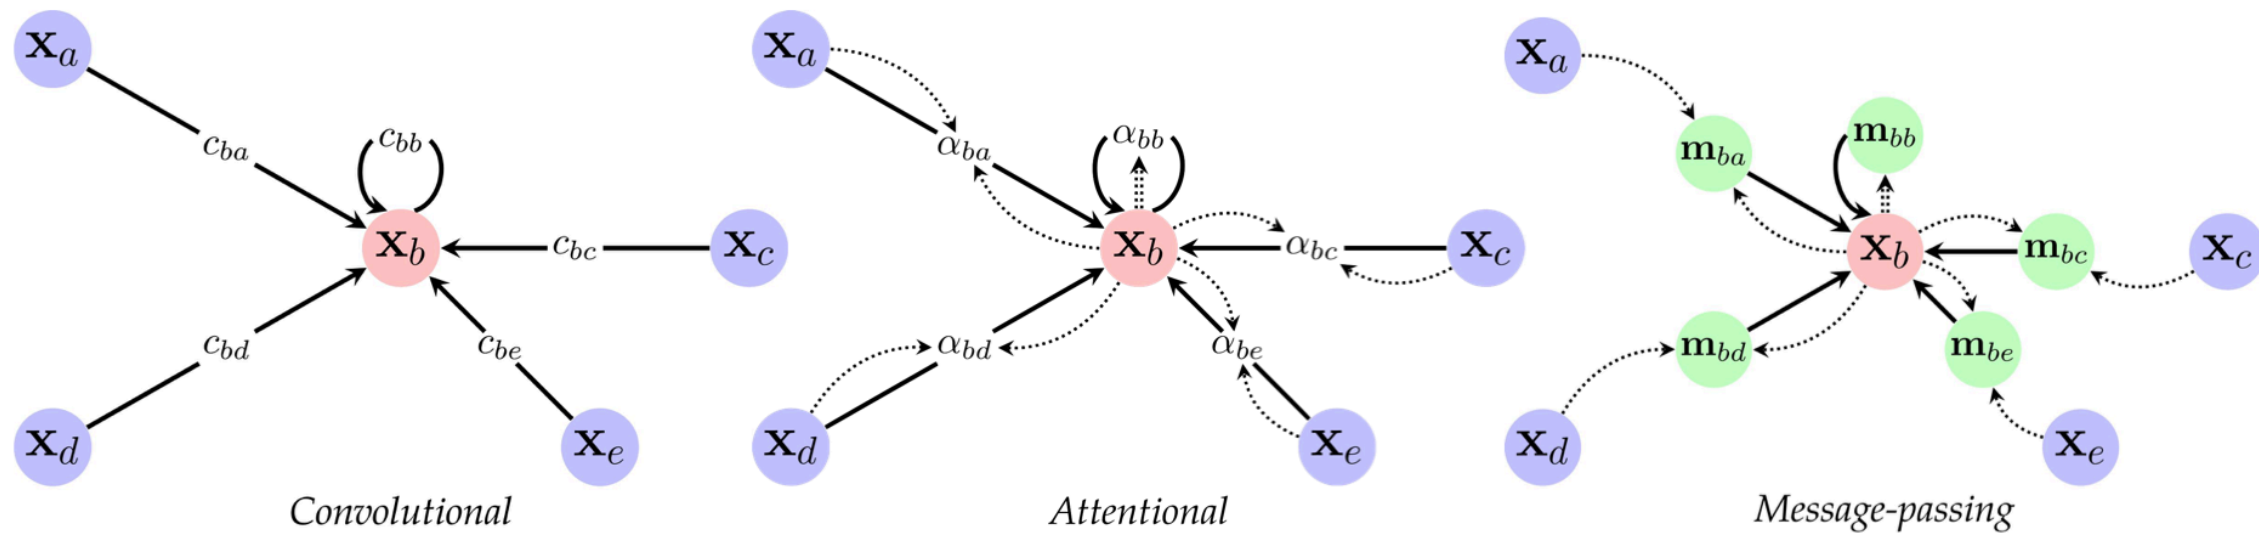
• Recent trends in introducing GSP-inspired architectures

- GraphHeat (Xu'19), GWNN (Xu'19), SIGN (Frasca'20), DGN (Beaini'20), Spectral GNs (Stachenfeld'20), Framelets (Zheng'21), FAGCN (Bo'21) ...

Balcilar et al., "Analyzing the expressive power of graph neural networks in a spectral perspective," ICLR, 2021

Bronstein et al., Geometric deep learning: Grids, Groups, Graphs, Geodesics, and Gauges, arXiv, 2021

GNNs in one slide



$$\mathbf{h}_i = \phi \left(\mathbf{x}_i, \bigoplus_{j \in \mathcal{N}_i} c_{ij} \psi(\mathbf{x}_j) \right)$$

ChebyNet
GCN
SGC

$$\mathbf{h}_i = \phi \left(\mathbf{x}_i, \bigoplus_{j \in \mathcal{N}_i} a(\mathbf{x}_i, \mathbf{x}_j) \psi(\mathbf{x}_j) \right)$$

MoNet
GAT

$$\mathbf{h}_i = \phi \left(\mathbf{x}_i, \bigoplus_{j \in \mathcal{N}_i} \psi(\mathbf{x}_i, \mathbf{x}_j) \right)$$

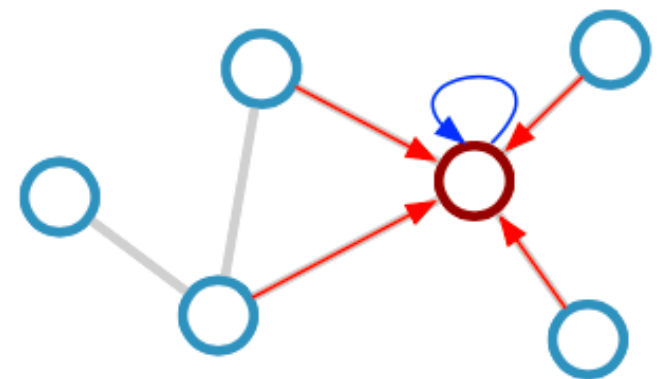
MPNN
GraphNets

Slide taken from P. Veličković

Spatial approaches in one slide

- Generate node embeddings based on local neighborhoods
 - nodes aggregate information from their neighbors using neural networks
- Feed the embeddings into a loss function
- **Key difference:** how nodes aggregate information across layers

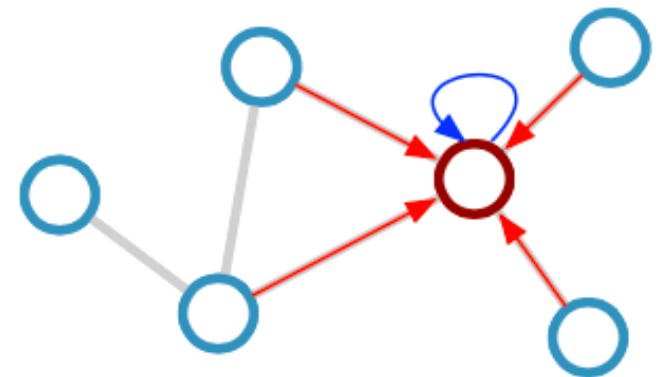
$$h_i^{(k+1)} = \sigma(W^{(k)} h_i^{(k)} + Q^{(k)} \sum_{n \in \mathcal{N}_i} h_n^{(k)})$$



Spatial approaches in one slide

- Generate node embeddings based on local neighborhoods
 - nodes aggregate information from their neighbors using neural networks
- Feed the embeddings into a loss function
- **Key difference:** how nodes aggregate information across layers

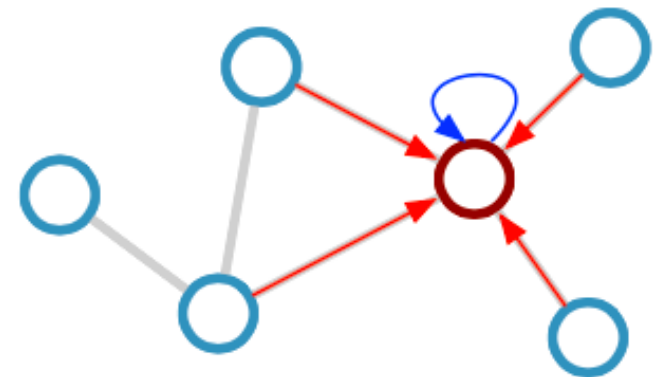
$$h_i^{(k+1)} = \sigma \left(W^{(k)} h_i^{(k)} + Q^{(k)} \sum_{n \in \mathcal{N}_i} h_n^{(k)} \right)$$



Spatial approaches in one slide

- Generate node embeddings based on local neighborhoods
 - nodes aggregate information from their neighbors using neural networks
- Feed the embeddings into a loss function
- **Key difference:** how nodes aggregate information across layers

$$h_i^{(k+1)} = \sigma \left(W^{(k)} h_i^{(k)} + Q^{(k)} \sum_{n \in \mathcal{N}_i} h_n^{(k)} \right)$$

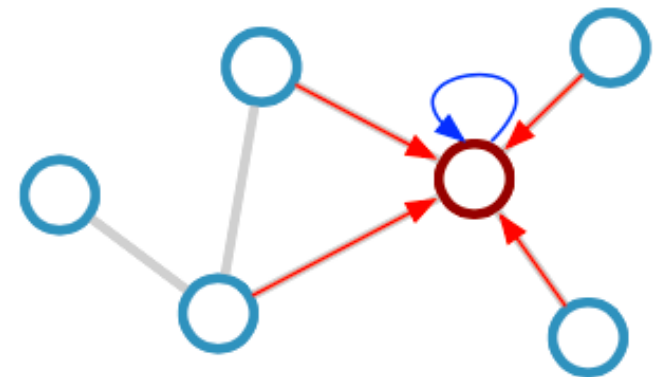


Spatial approaches in one slide

- Generate node embeddings based on local neighborhoods
 - nodes aggregate information from their neighbors using neural networks
- Feed the embeddings into a loss function
- **Key difference:** how nodes aggregate information across layers

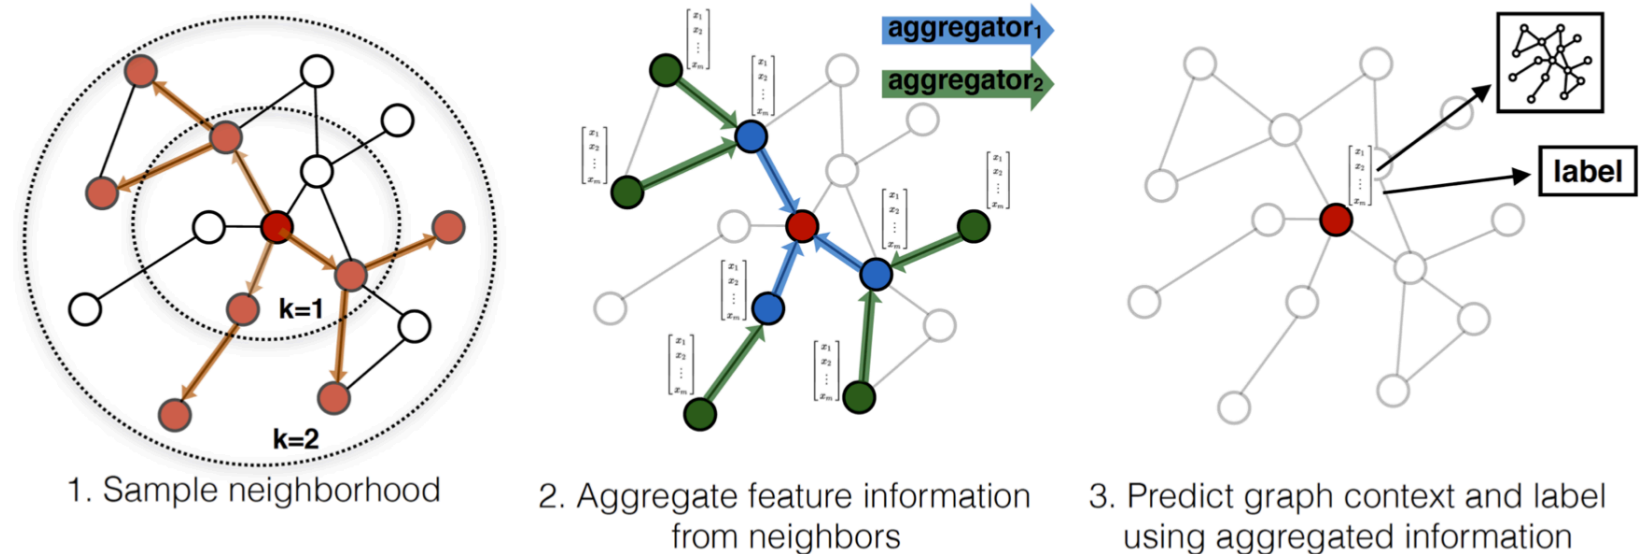
$$h_i^{(k+1)} = \sigma \left(W^{(k)} h_i^{(k)} + Q^{(k)} \sum_{n \in \mathcal{N}_i} h_n^{(k)} \right)$$

Trainable parameters



GraphSAGE

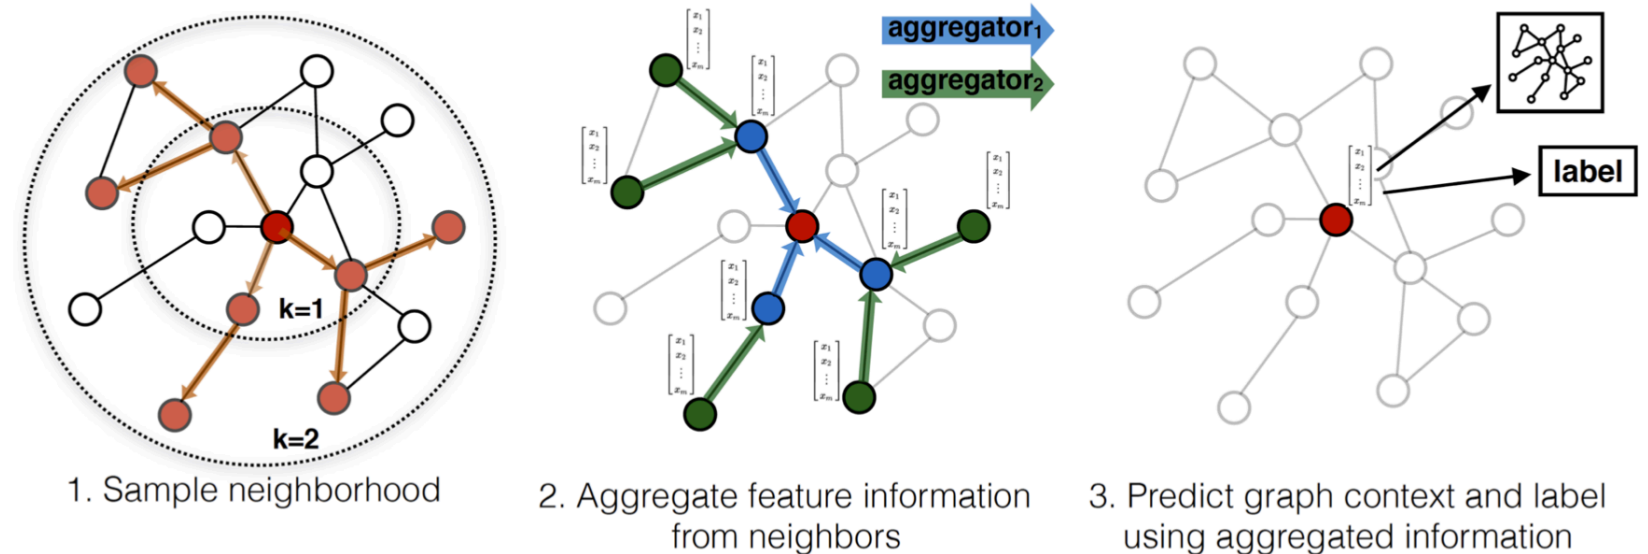
- Node's neighborhood defines a computational graph
- Each edge in this graph is a transformation/aggregator function
- Cross-entropy loss
- 2-3 layers deep



$$h_i^{(k+1)} = \text{Relu}(W^{(k)} h_i^{(k)}, \sum_{n \in \mathcal{N}_i} (\text{Relu}(Q^{(k)} h_n^{(k)})))$$

GraphSAGE

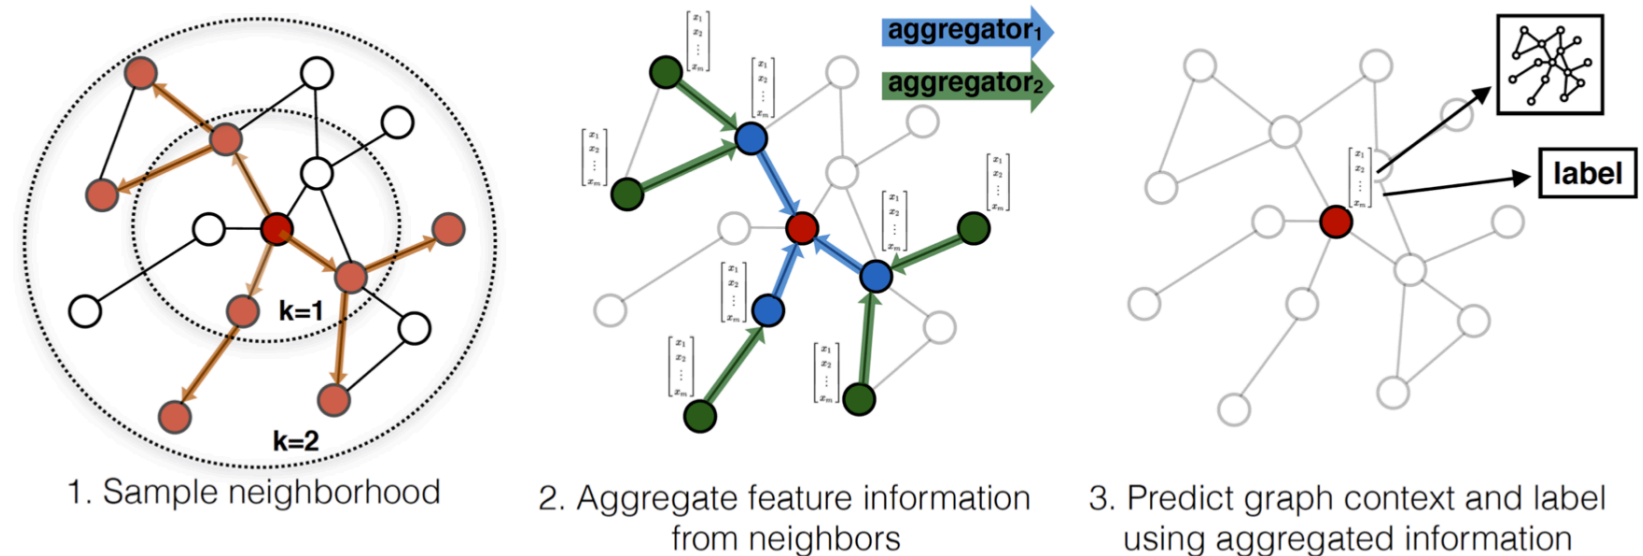
- Node's neighborhood defines a computational graph
- Each edge in this graph is a transformation/aggregator function
- Cross-entropy loss
- 2-3 layers deep



$$h_i^{(k+1)} = Relu(W^{(k)} h_i^{(k)}, \sum_{n \in \mathcal{N}_i} (Relu(Q^{(k)} h_n^{(k)})))$$

GraphSAGE

- Node's neighborhood defines a computational graph
- Each edge in this graph is a transformation/aggregator function
- Cross-entropy loss
- 2-3 layers deep

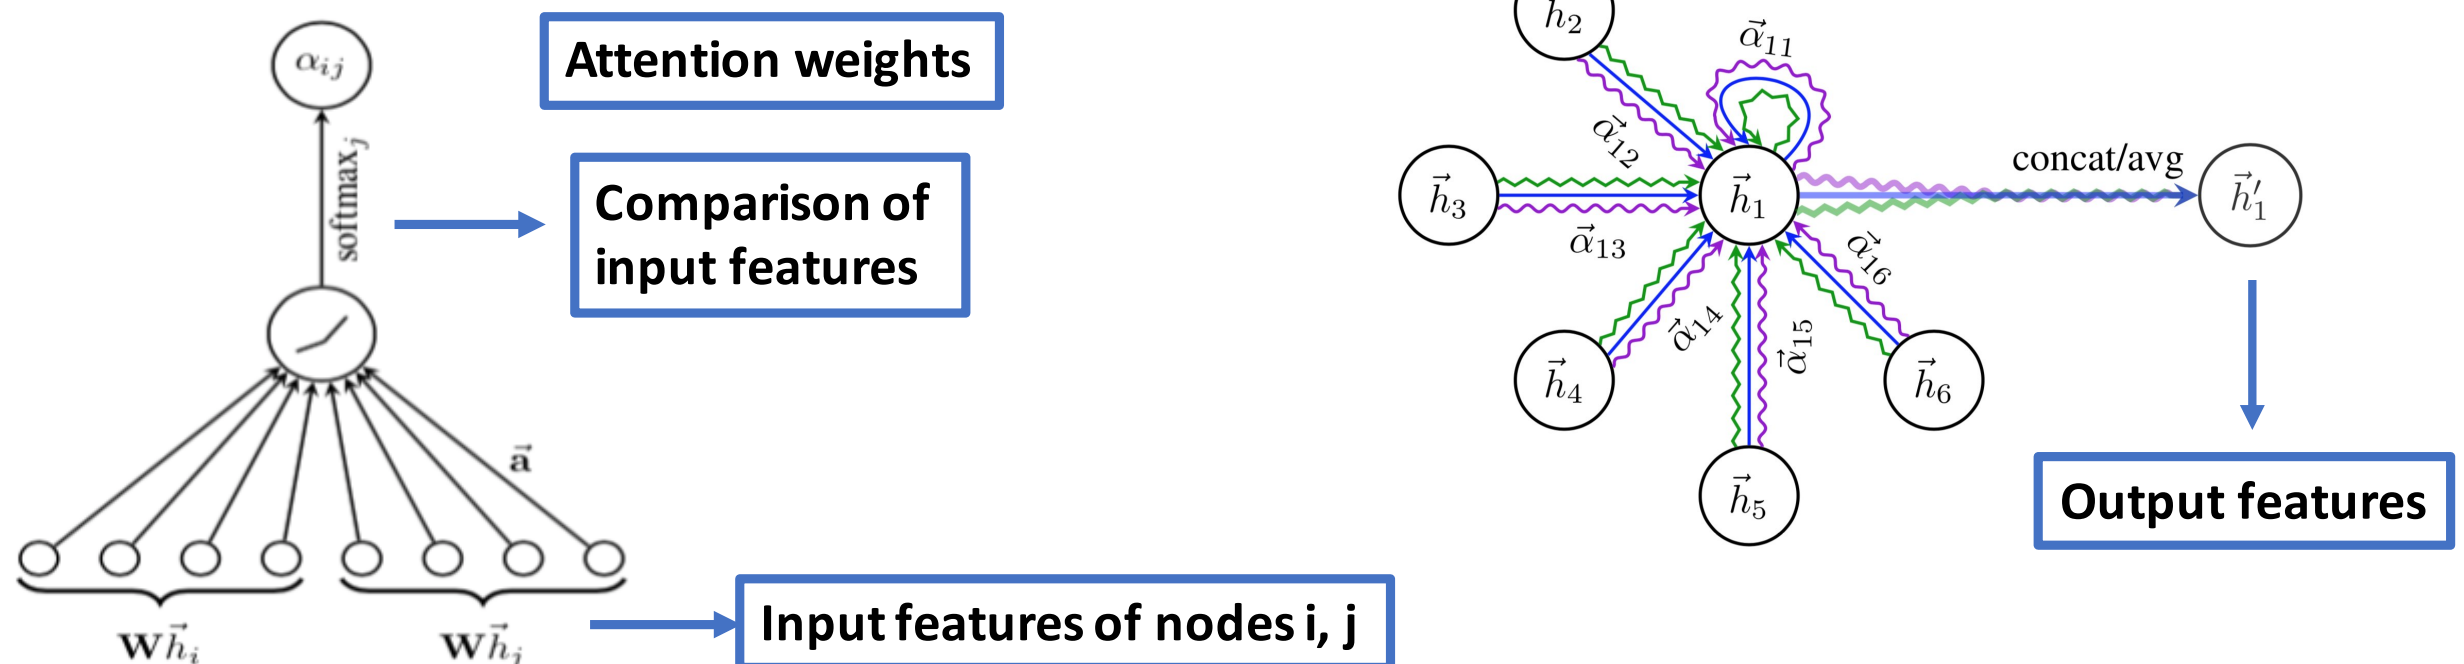


$$h_i^{(k+1)} = \text{Relu}(W^{(k)} h_i^{(k)}, \sum_{n \in \mathcal{N}_i} (\text{Relu}(Q^{(k)} h_n^{(k)})))$$

Averaging/max pooling/LSTM

Graph attention networks (GAT)

- Nodes attend over their neighborhood's features
 - Different weights to different nodes in a neighbourhood
 - Remove dependence on the global graph structure



$$\vec{h}'_i = \sigma \left(\frac{1}{K} \sum_{k=1}^K \sum_{j \in \mathcal{N}_i} \alpha_{ij}^k \mathbf{W}^k \vec{h}_j \right)$$

Summary of spatial methods

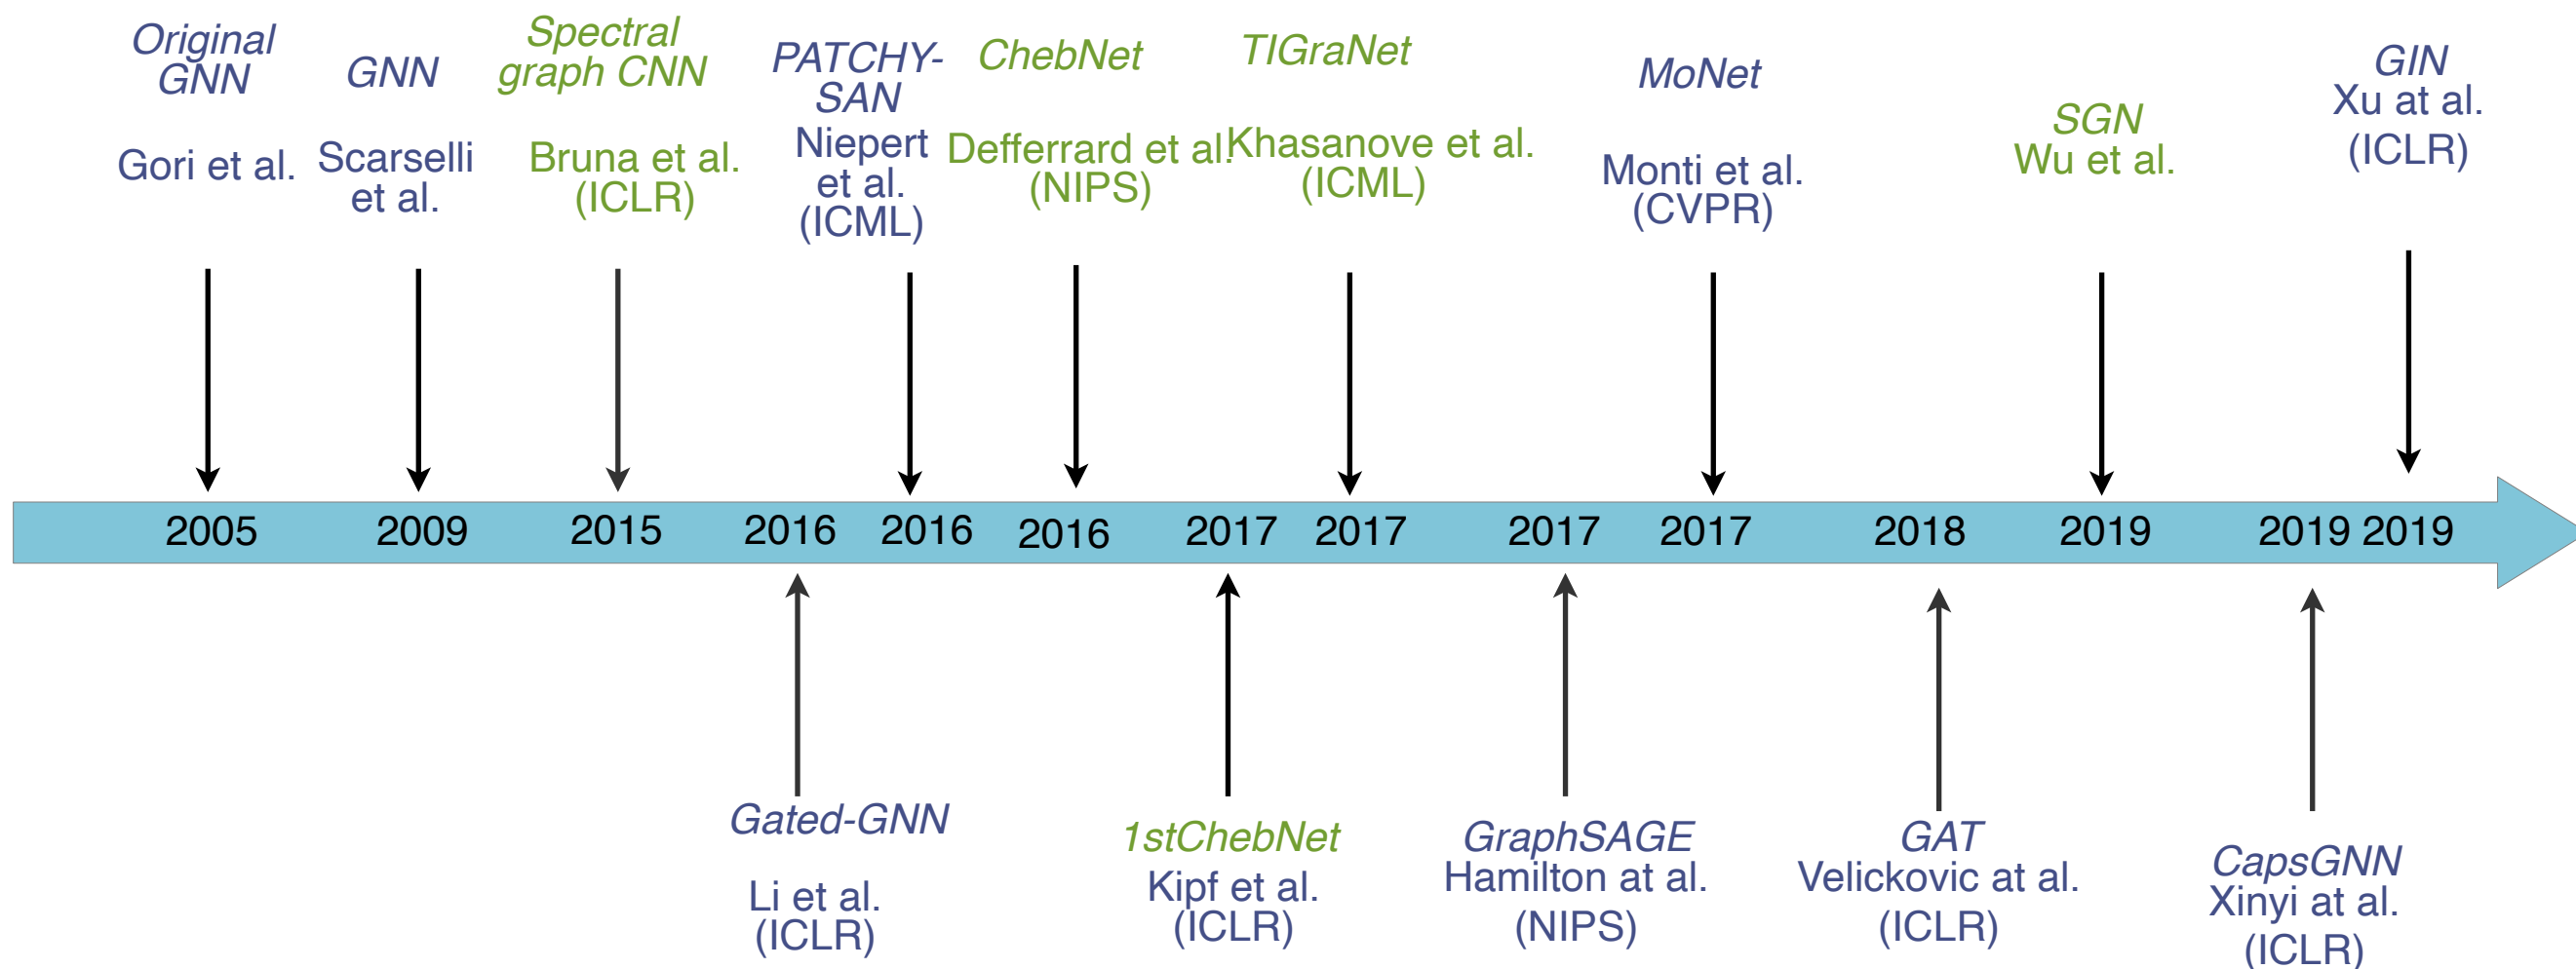
- **Pros:**

- Intuitive
- Easy to implement
- Generalized to inductive settings

- **Cons:**

- Lack of interpretation in the spectral domain
- Requires many message passing iterations if the size of the graph is large

Partial historical overview



Spatial-based methods

Spectral-based methods

Spectral approaches in one slide

- Convolution is defined on the graph Fourier domain

$$x *_\theta g = \chi \hat{g}(\theta) \chi^T x$$

- Spectral GCNN:** $\hat{g}(\theta) = \Theta \quad \longrightarrow \quad x *_\theta g = \chi \Theta \chi^T x$

- ChebNet:** $\hat{g}(\theta) = \sum_{i=1}^K \theta_i T_i(\Lambda) \quad \longrightarrow \quad x *_\theta g = \sum_{i=1}^K \theta_i T_i(L) x$

- GCN:** $K = 1 \quad \longrightarrow \quad x *_\theta g = (\theta_0 I - \theta_1 D^{-1/2} A D^{-1/2}) x$

- Parameters Θ are learned through the network

Spectral approaches in one slide

- Convolution is defined on the graph Fourier domain

$$x *_\theta g = \chi \hat{g}(\theta) \chi^T x$$

- Spectral GCNN:** $\hat{g}(\theta) = \Theta \quad \longrightarrow \quad x *_\theta g = \chi \Theta \chi^T x$

- ChebNet:** $\hat{g}(\theta) = \sum_{i=1}^K \theta_i T_i(\Lambda) \quad \longrightarrow \quad x *_\theta g = \sum_{i=1}^K \theta_i T_i(L) x$

- GCN:** $K = 1 \quad \longrightarrow \quad x *_\theta g = (\theta_0 I - \theta_1 D^{-1/2} A D^{-1/2}) x$

- Parameters Θ are learned through the network

Spectral approaches in one slide

- Convolution is defined on the graph Fourier domain

$$x *_\theta g = \chi \hat{g}(\theta) \chi^T x$$

- Spectral GCNN:** $\hat{g}(\theta) = \Theta \quad \longrightarrow \quad x *_\theta g = \chi \Theta \chi^T x$

- ChebNet:** $\hat{g}(\theta) = \sum_{i=1}^K \theta_i T_i(\Lambda) \quad \longrightarrow \quad x *_\theta g = \sum_{i=1}^K \theta_i T_i(L) x$

- GCN:** $K = 1 \quad \longrightarrow \quad x *_\theta g = (\theta_0 I - \theta_1 D^{-1/2} A D^{-1/2}) x$

- Parameters Θ are learned through the network

Spectral approaches in one slide

- Convolution is defined on the graph Fourier domain

$$x *_\theta g = \chi \hat{g}(\theta) \chi^T x$$

- Spectral GCNN:** $\hat{g}(\theta) = \Theta \quad \longrightarrow \quad x *_\theta g = \chi \Theta \chi^T x$

- ChebNet:** $\hat{g}(\theta) = \sum_{i=1}^K \theta_i T_i(\Lambda) \quad \longrightarrow \quad x *_\theta g = \sum_{i=1}^K \theta_i T_i(L) x$

- GCN:** $K = 1 \quad \longrightarrow \quad x *_\theta g = (\theta_0 I - \theta_1 D^{-1/2} A D^{-1/2}) x$

- Parameters Θ are learned through the network

Spectral approaches in one slide

- Convolution is defined on the graph Fourier domain

$$x *_\theta g = \chi \hat{g}(\theta) \chi^T x$$

- Spectral GCNN:** $\hat{g}(\theta) = \Theta \quad \longrightarrow \quad x *_\theta g = \chi \Theta \chi^T x$

- ChebNet:** $\hat{g}(\theta) = \sum_{i=1}^K \theta_i T_i(\Lambda) \quad \longrightarrow \quad x *_\theta g = \sum_{i=1}^K \theta_i T_i(L) x$

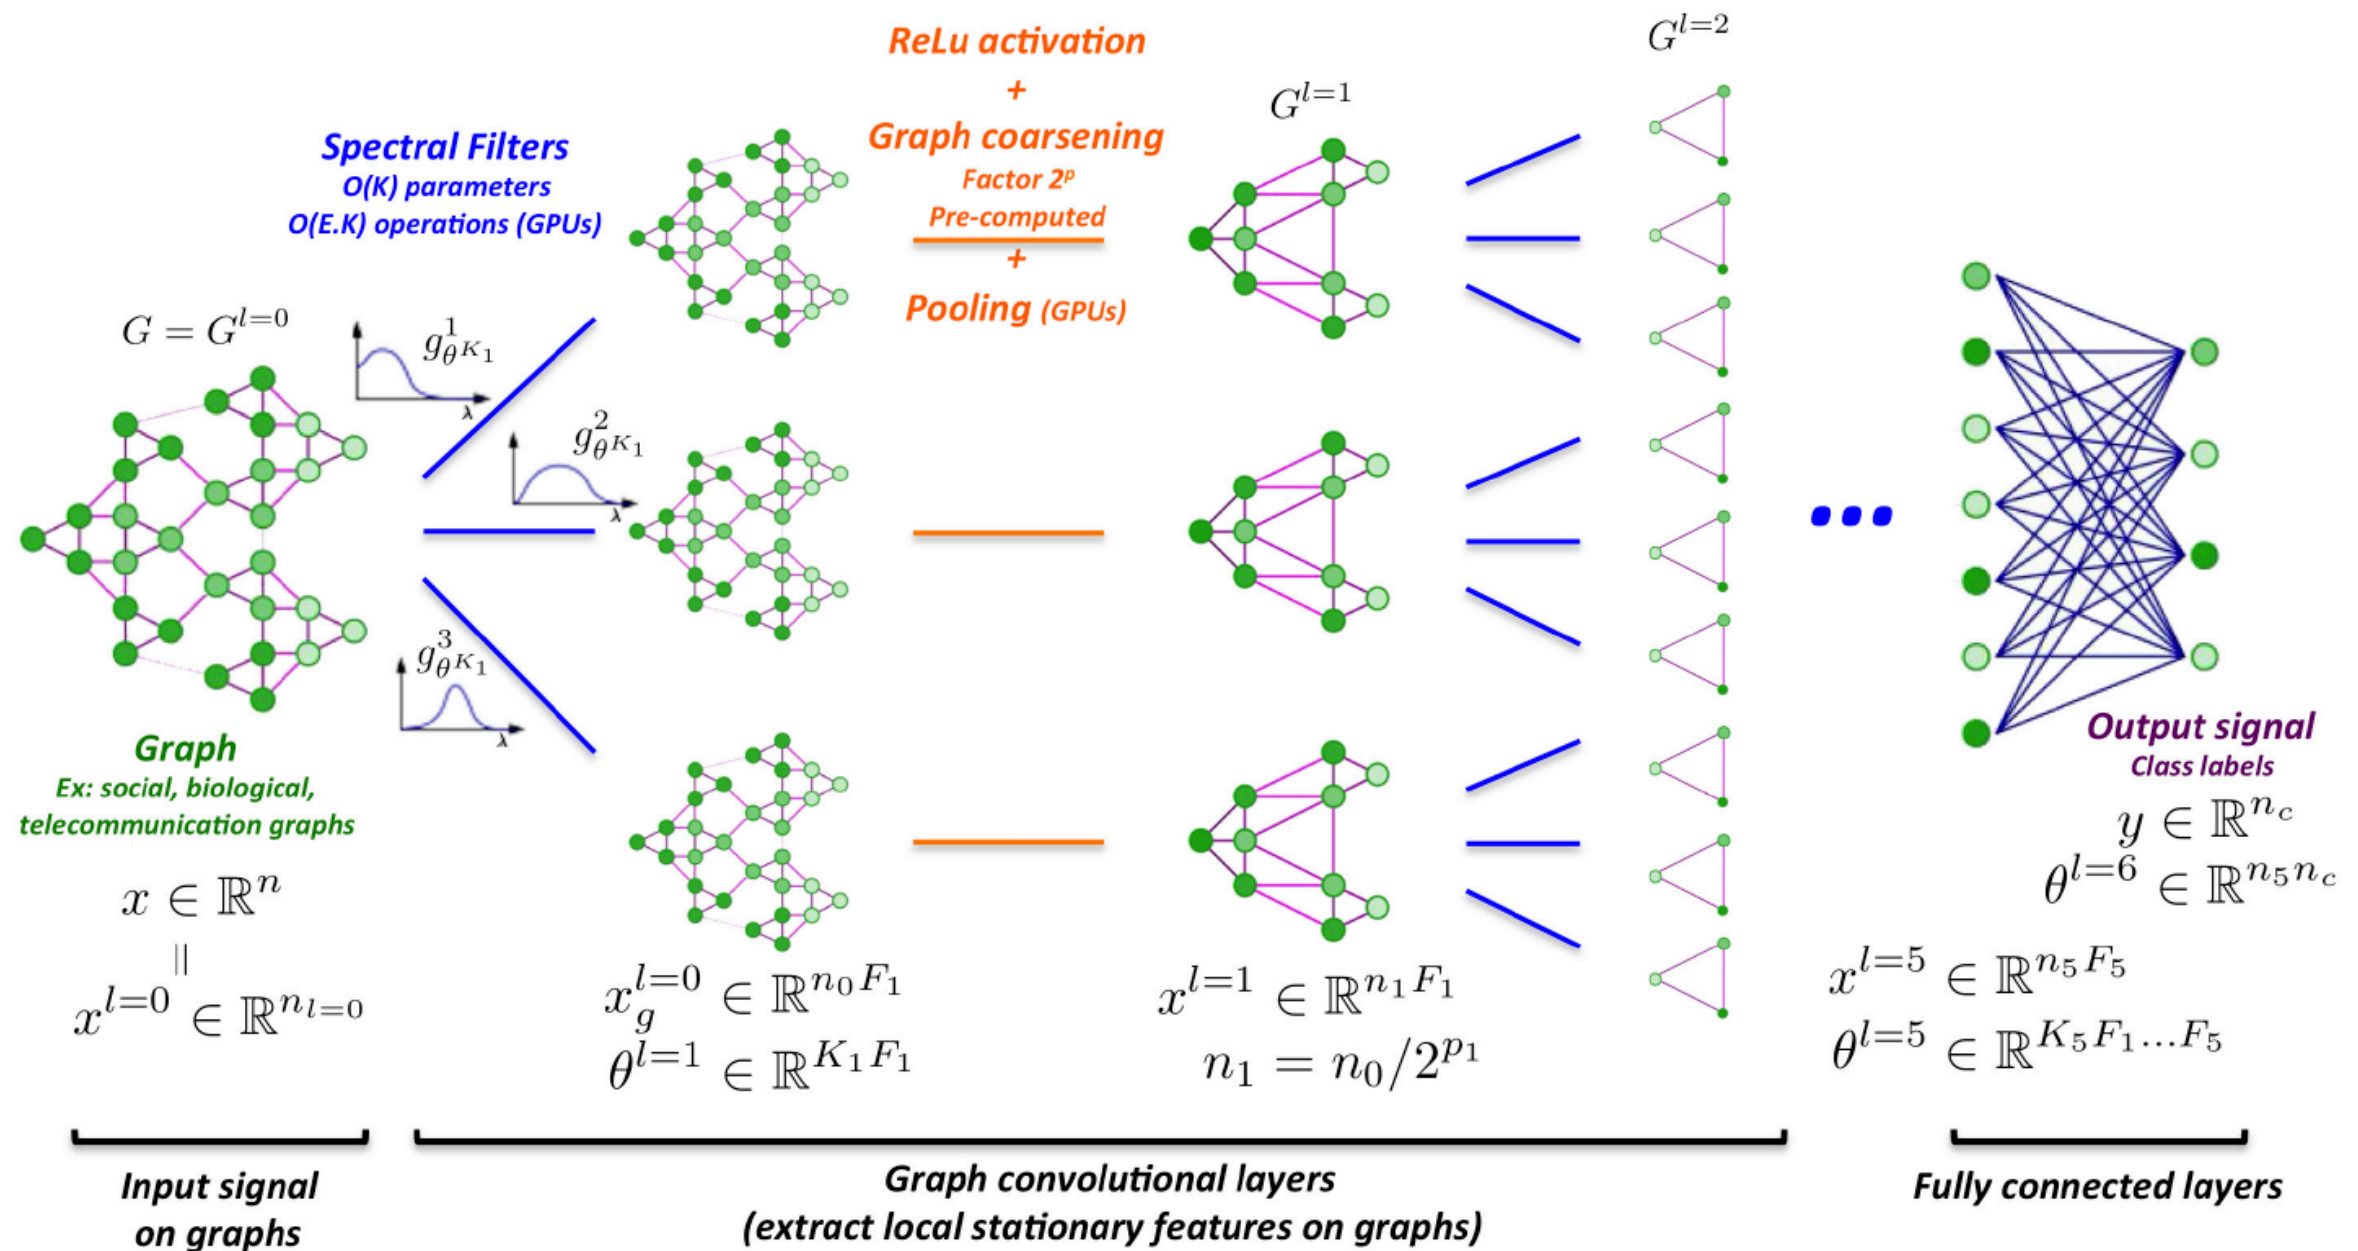
- GCN:** $K = 1 \quad \longrightarrow \quad x *_\theta g = (\theta_0 I - \theta_1 D^{-1/2} A D^{-1/2}) x$

Neighborhood aggregation

- Parameters Θ are learned through the network

ChebNet

- Graph filters: Chebyshev polynomials



GCN

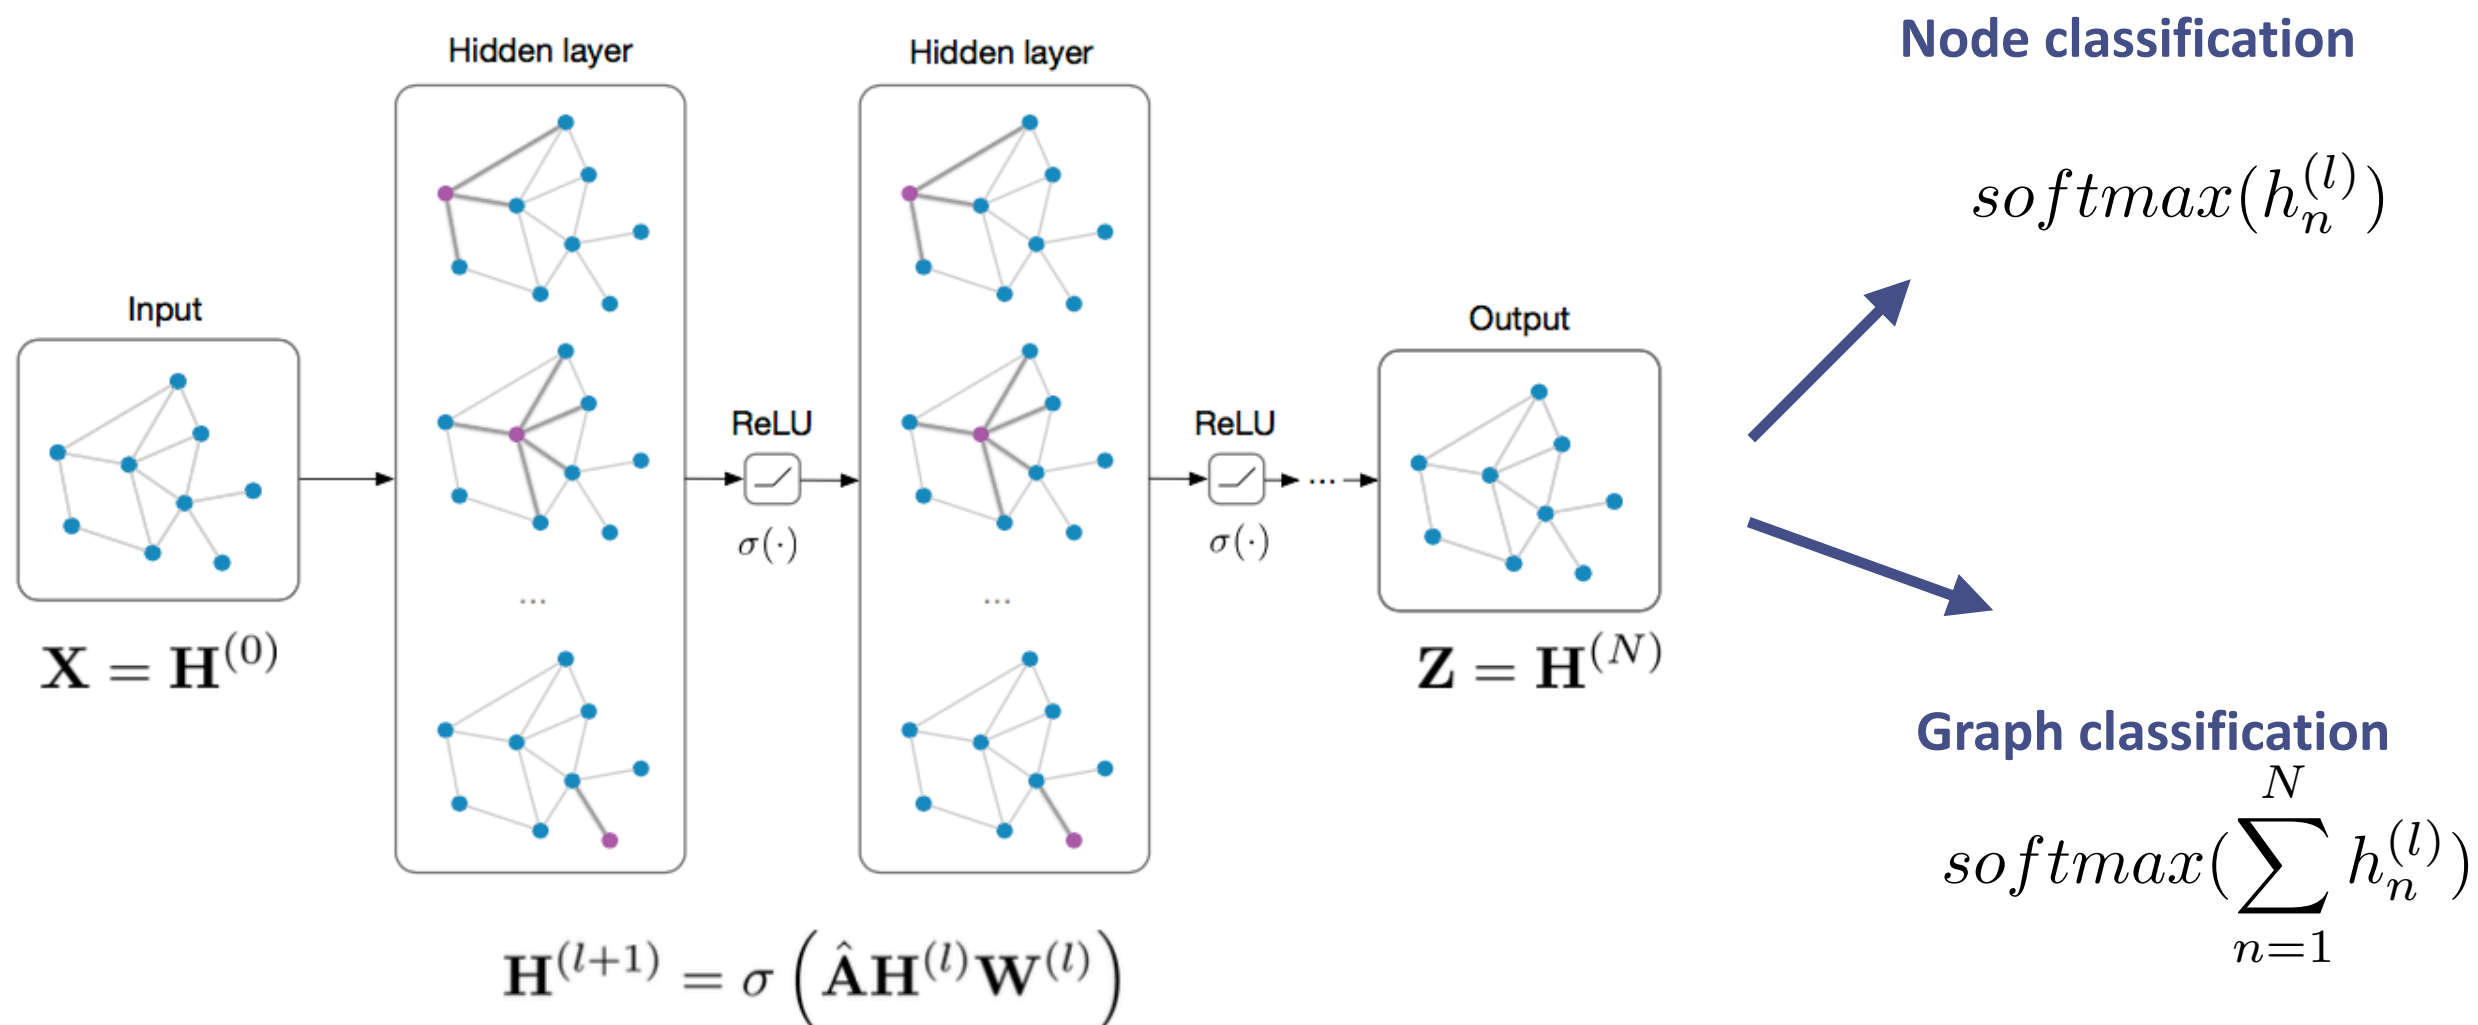


Fig. from T. Kirf

GCN

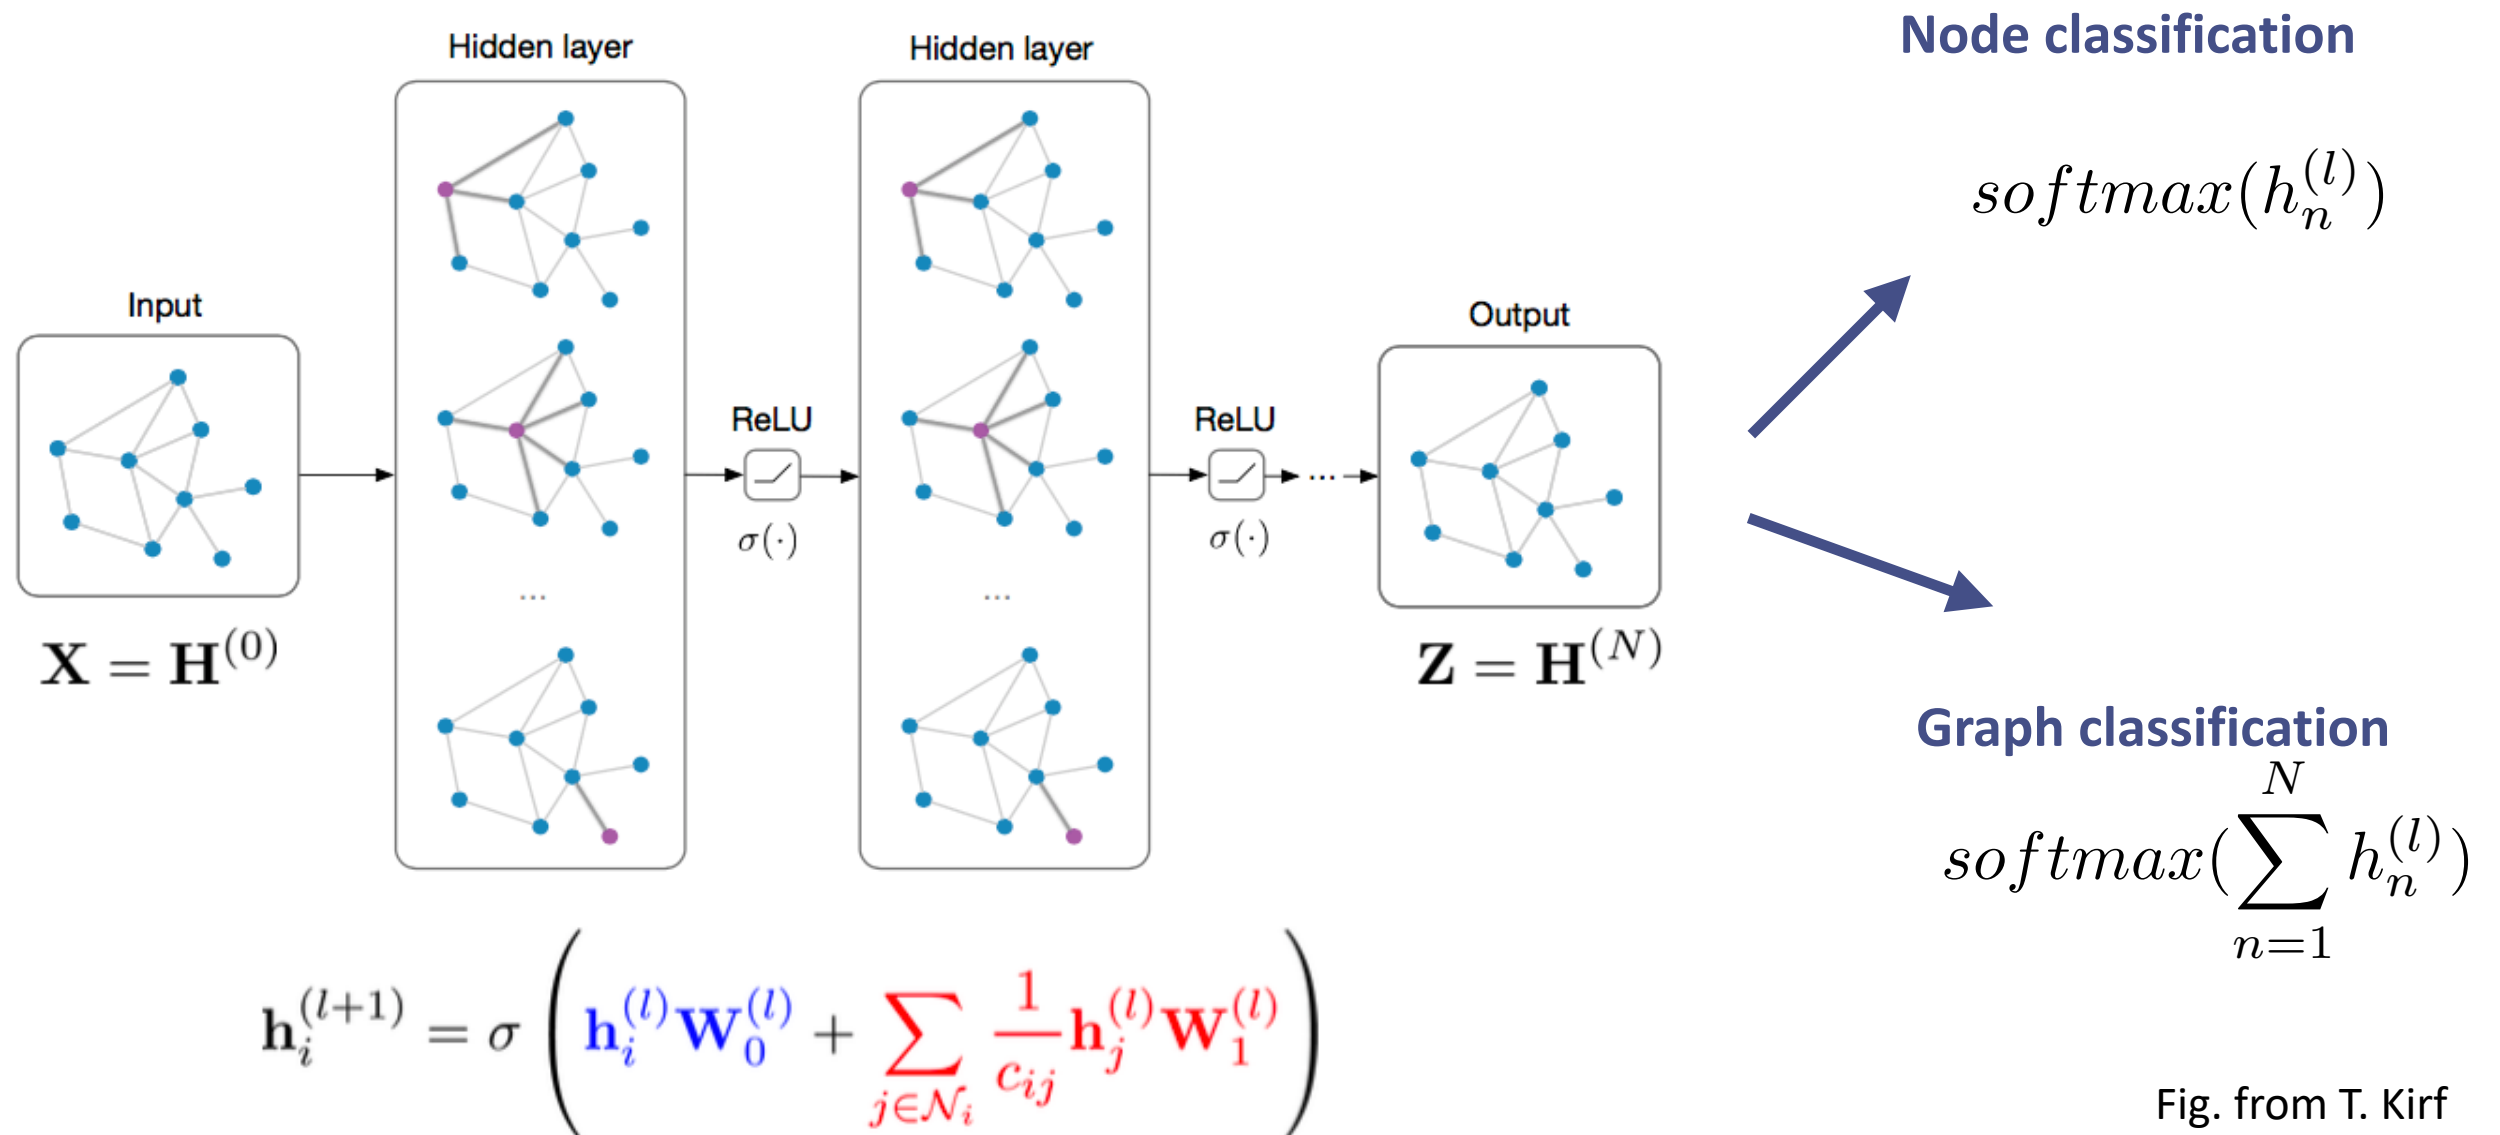


Fig. from T. Kirf

Simple graph convolution (SGC)

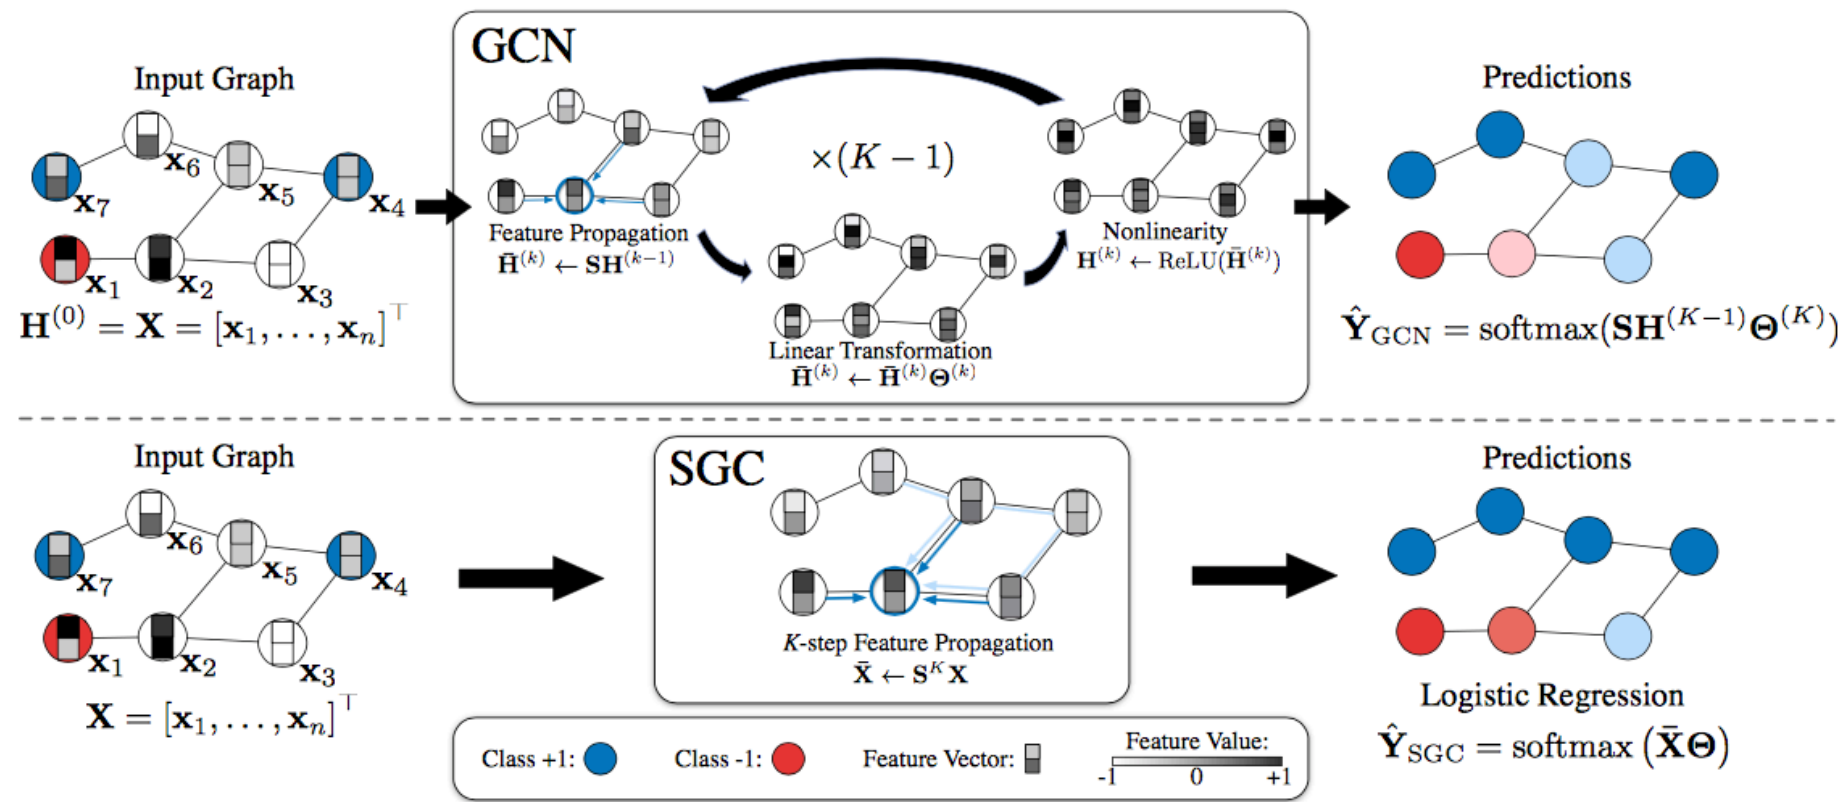


Fig. from Wu et al.

- Reduces the procedure to a simple feature propagation step followed by standard logistic regression

Summary of spectral methods

- **Pros:**

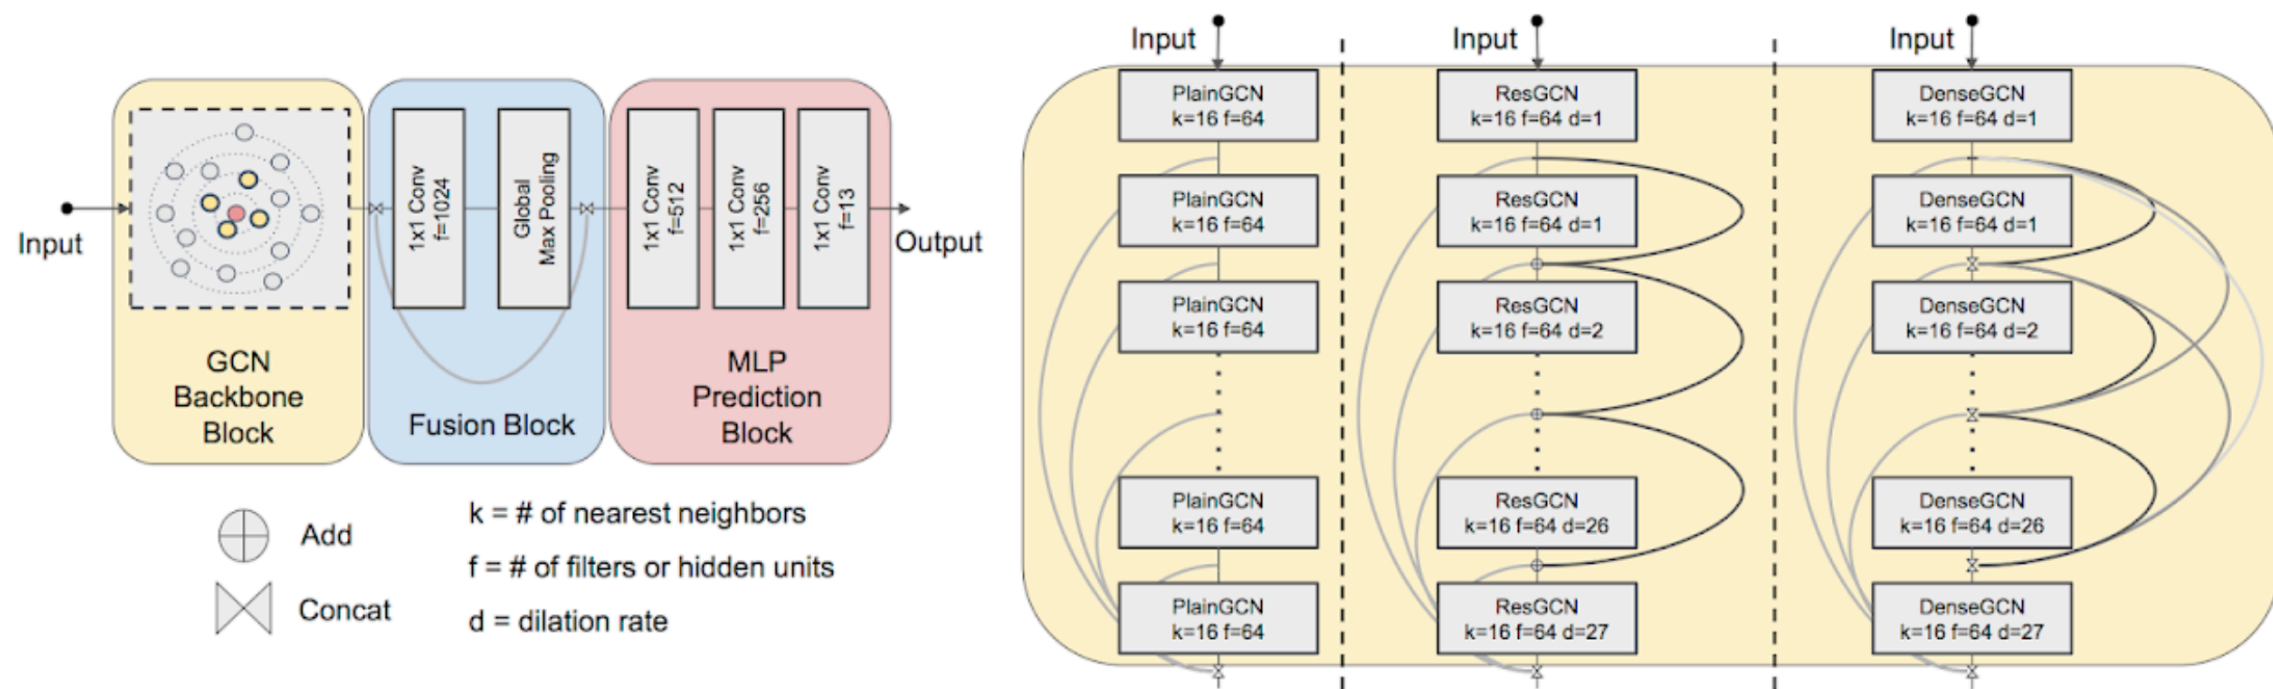
- More interpretable
- Takes into account the global structure of the graph

- **Cons:**

- Generalization is an open research question

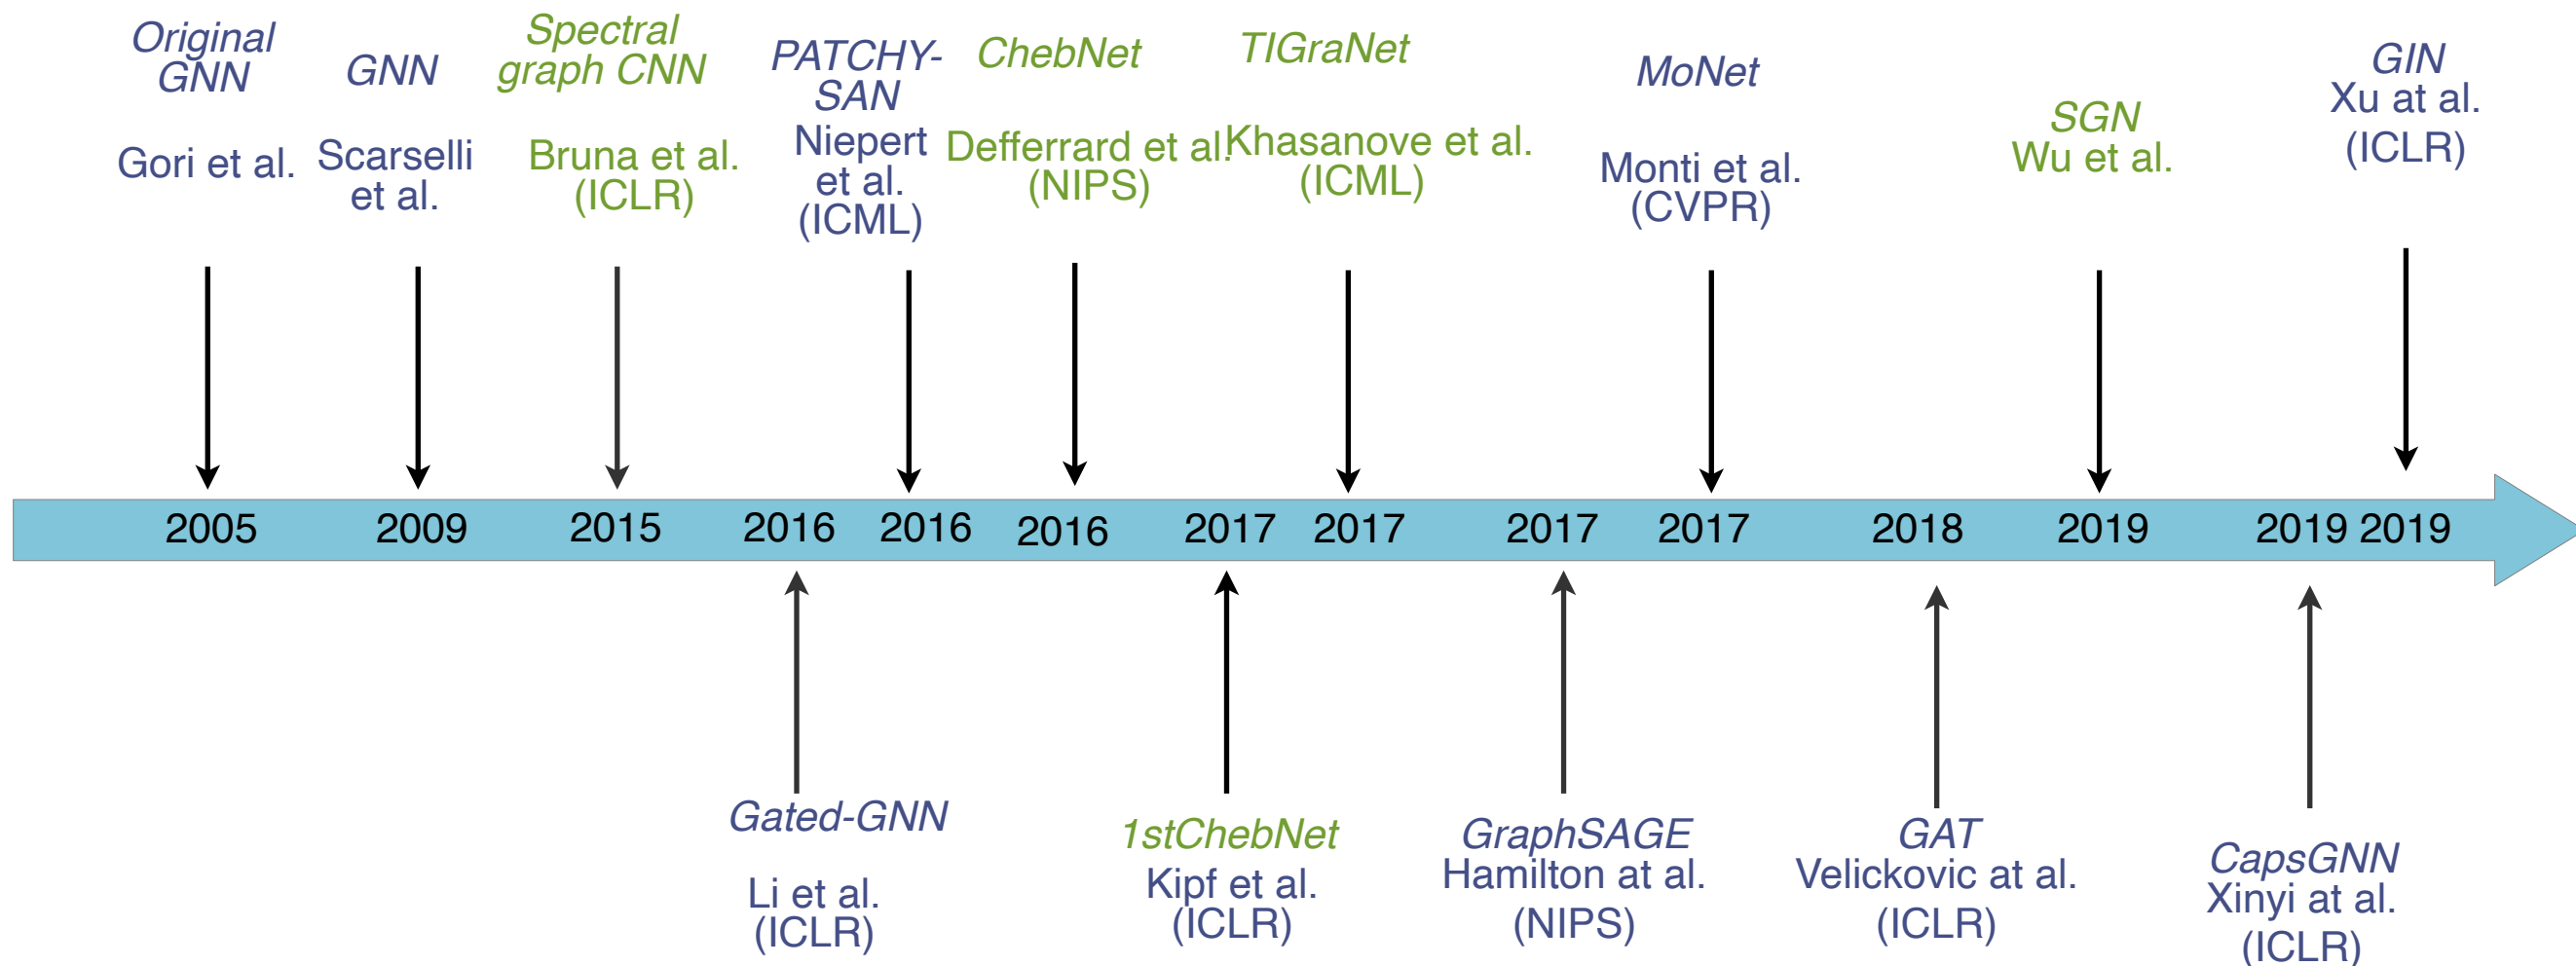
DeepGCNs

- Design deeper networks by using intuitions from ResNet, DenseNet



Li et. al, 2019

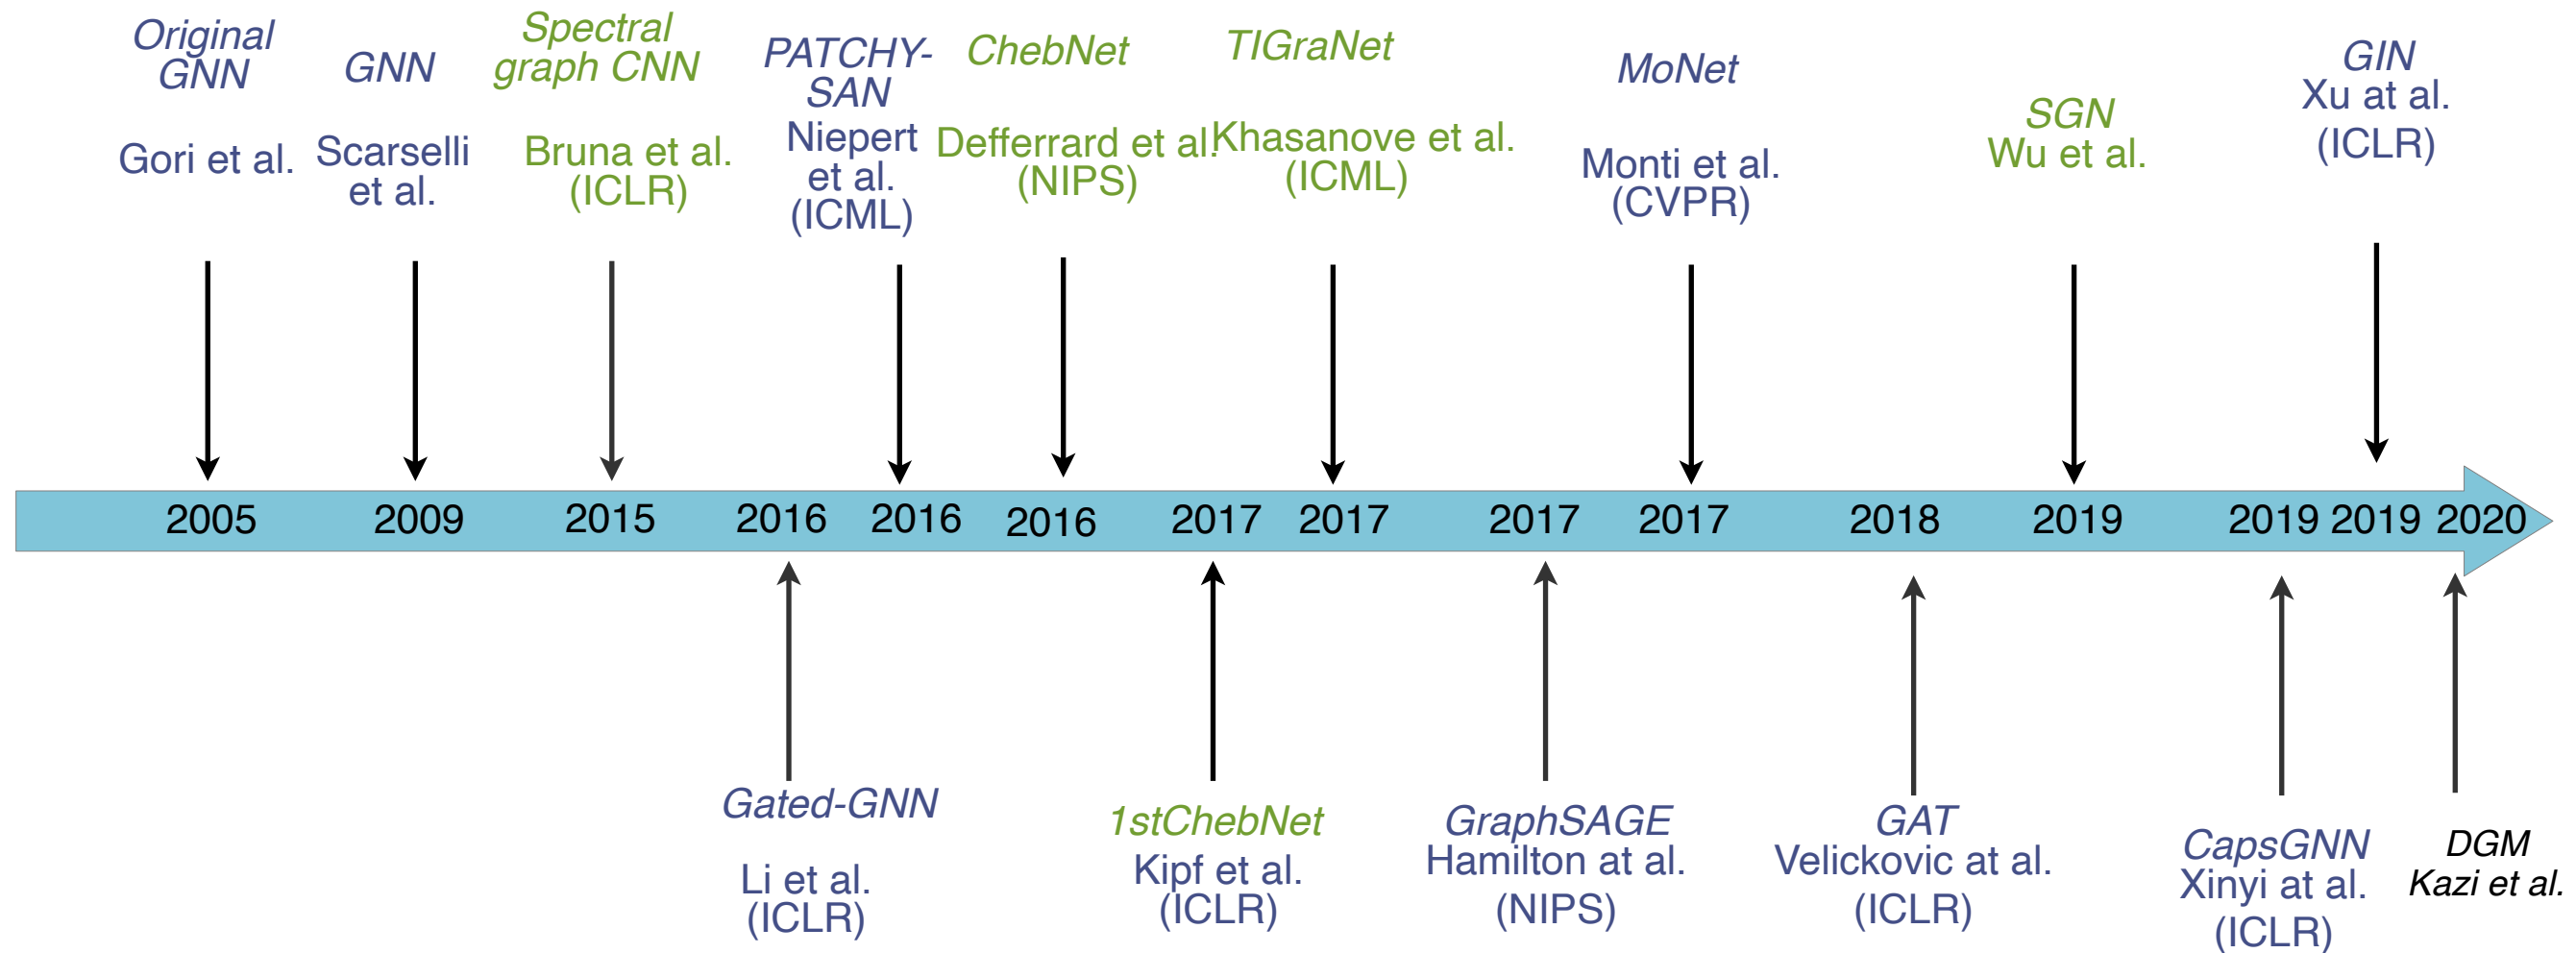
Partial historical overview



Spatial-based methods

Spectral-based methods

Partial historical overview

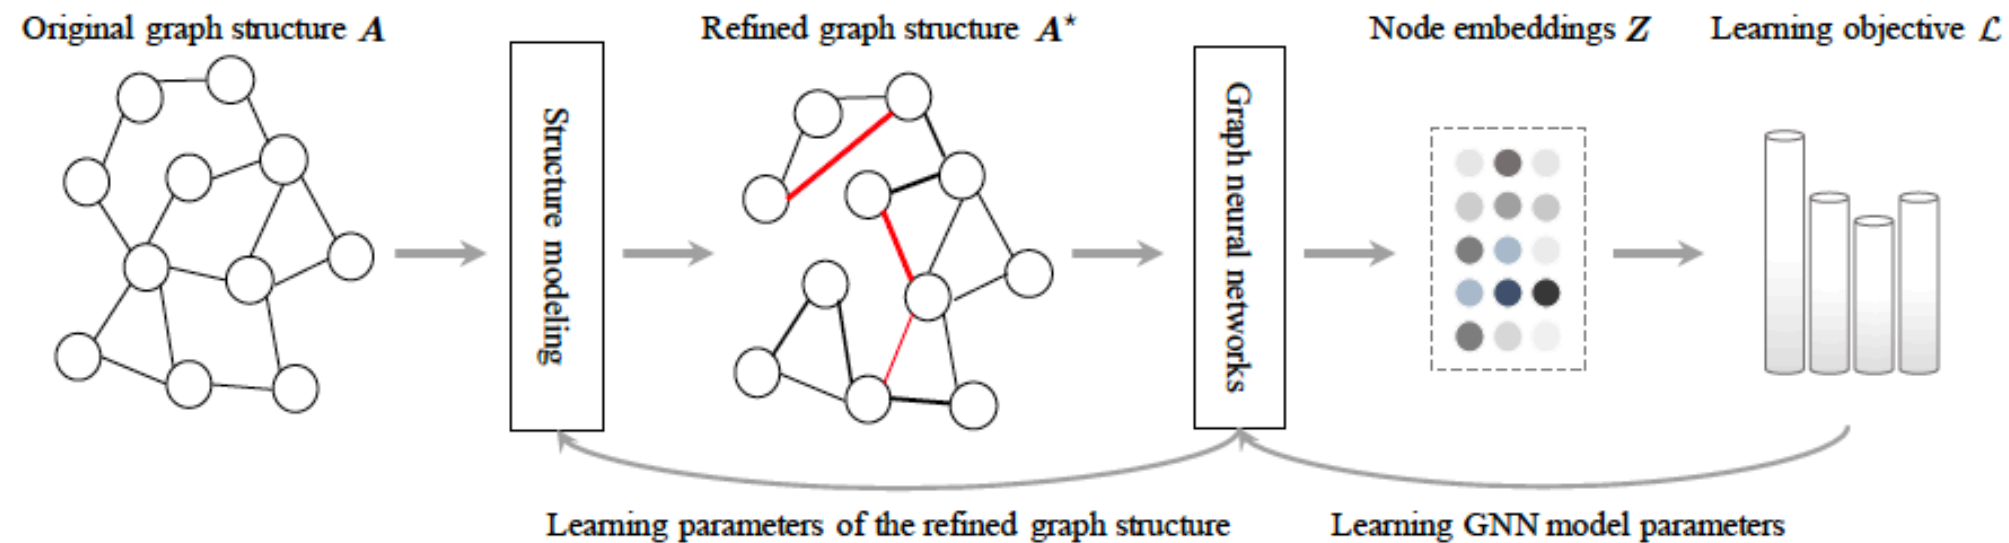


Spatial-based methods

Spectral-based methods

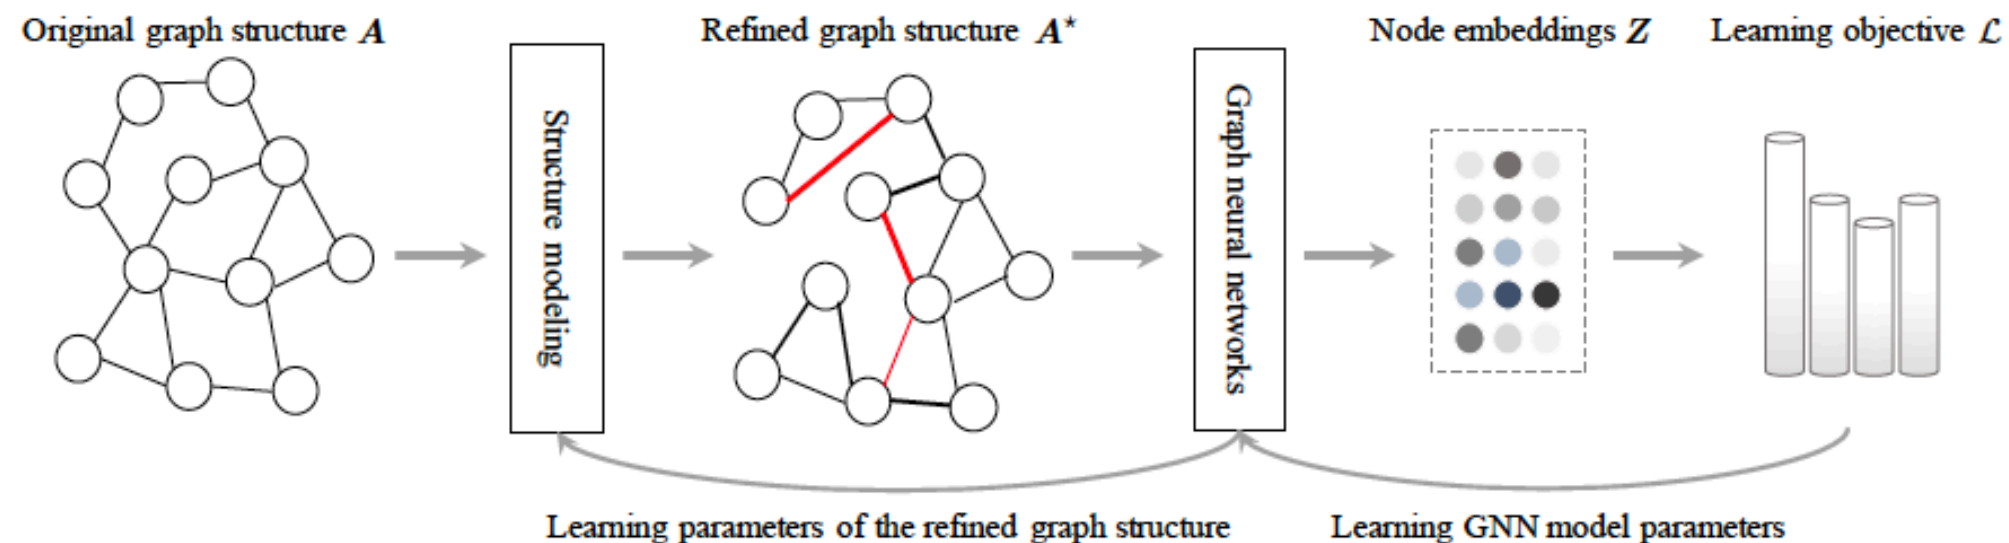
Joint graph learning

- Infers simultaneously the graph structure



Joint graph learning

- Infers simultaneously the graph structure

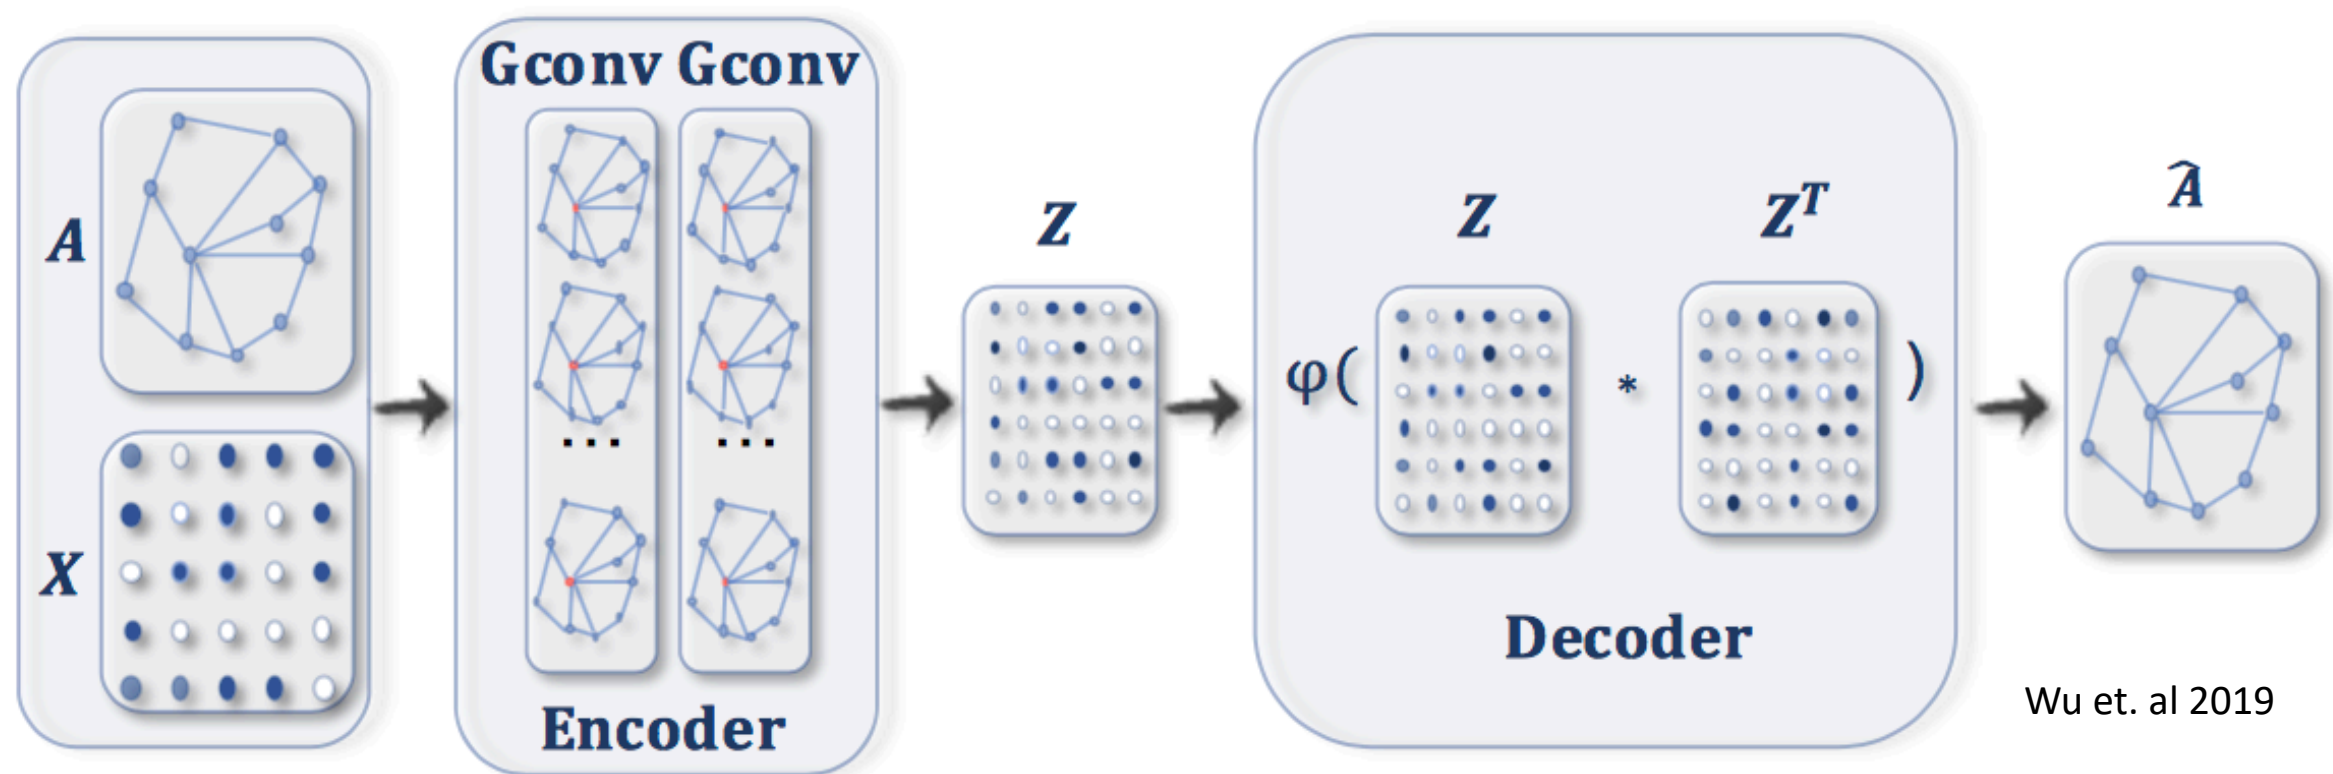


Outline

- Motivation: Why Graph Neural Networks?
- Basic definitions on graphs
- Partial historical overview
 - Graph CNN (main focus)
 - **Graph autoencoders (briefly)**
- Applications
 - Naturally graph-structured data
 - Images

Graph autoencoders (GAEs)

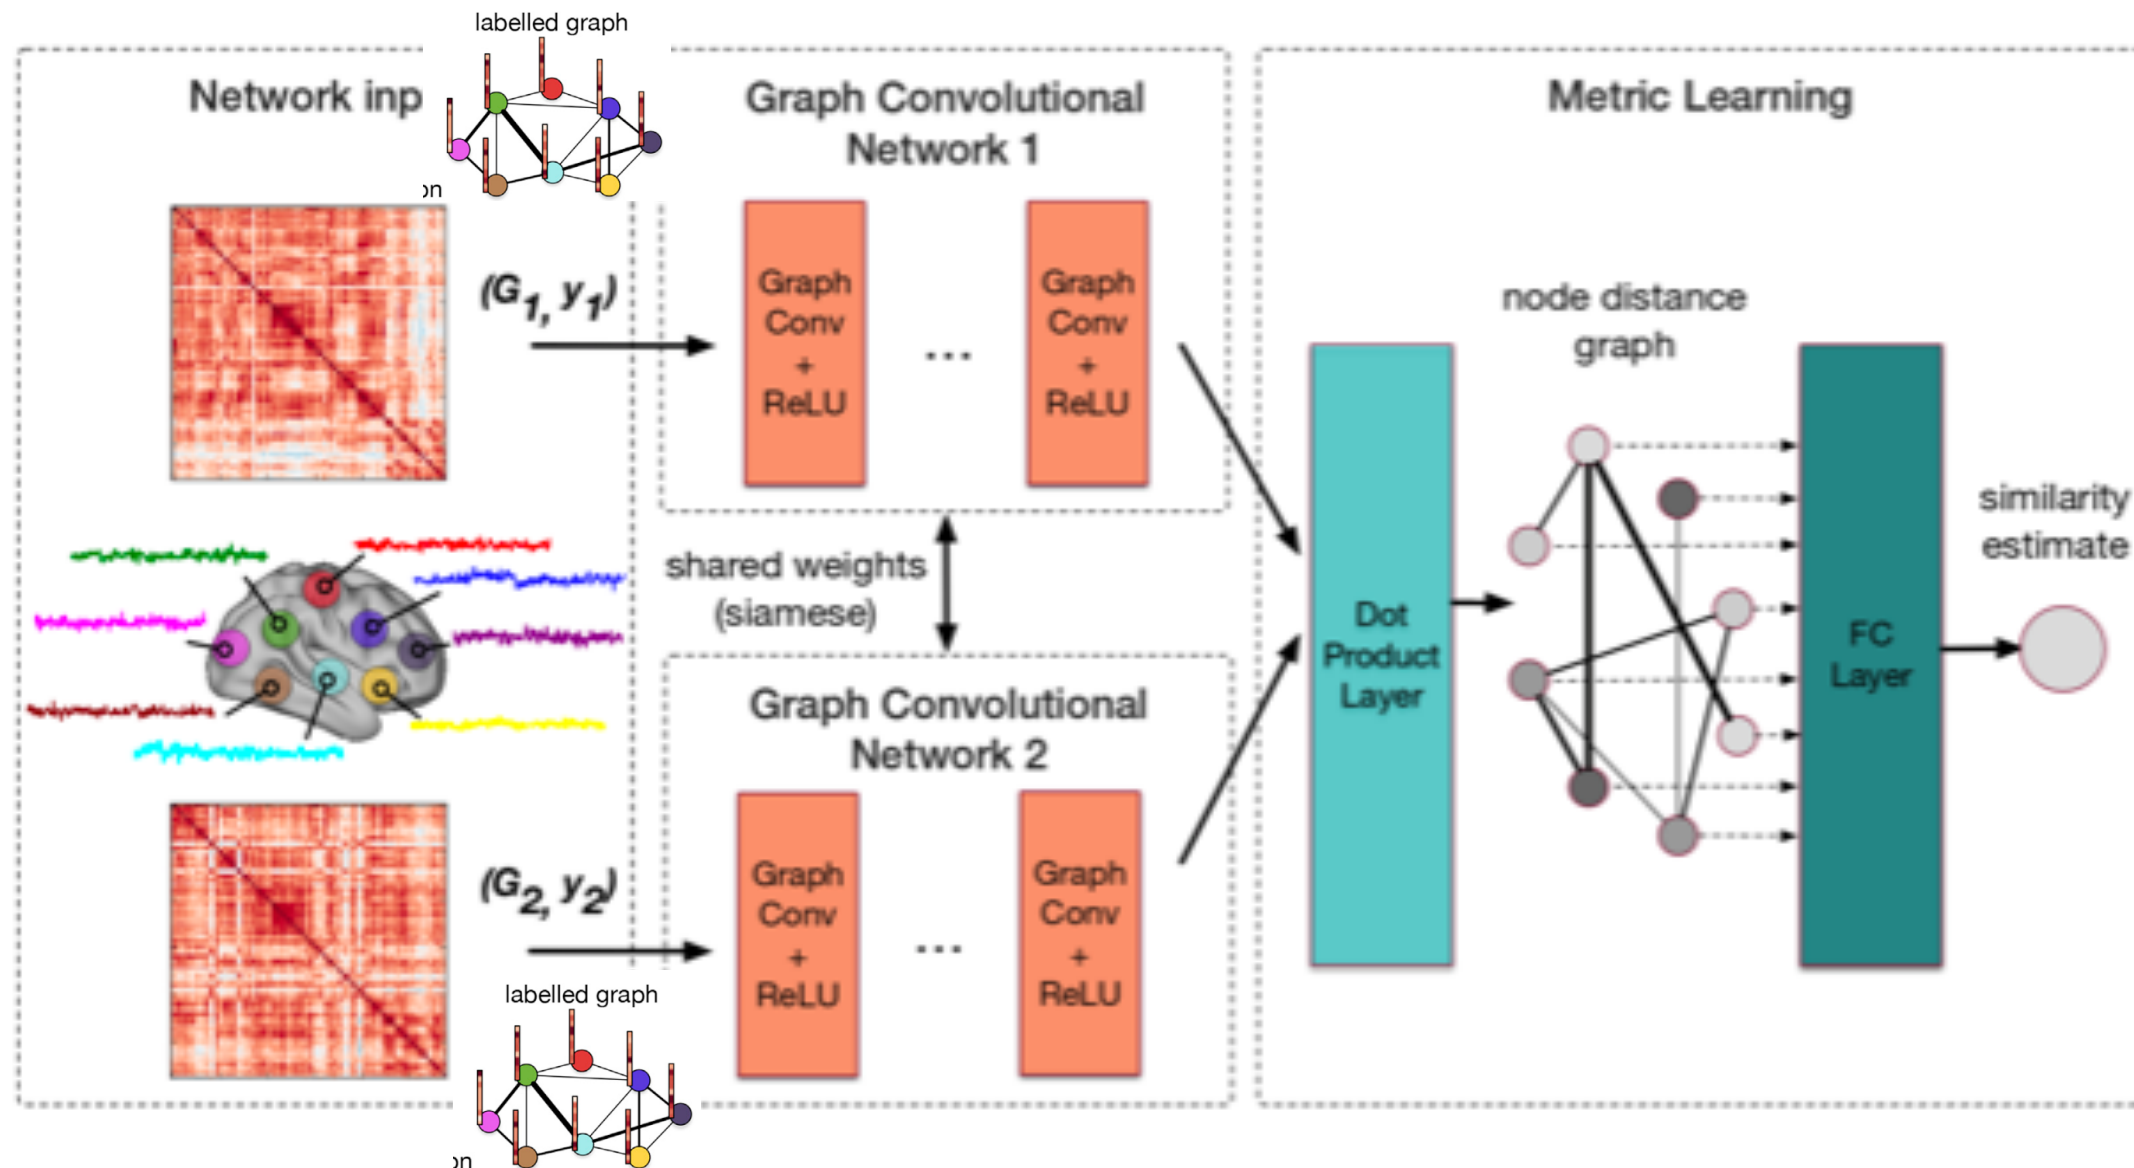
- Unsupervised learning: encode graphs into a latent space and reconstruct the graph from that space



Outline

- Motivation: Why Graph Neural Networks?
- Basic definitions on graphs
- Partial historical overview
 - Graph CNN (main focus)
 - Graph autoencoders (briefly)
- **Applications**
 - Naturally graph-structured data
 - Images

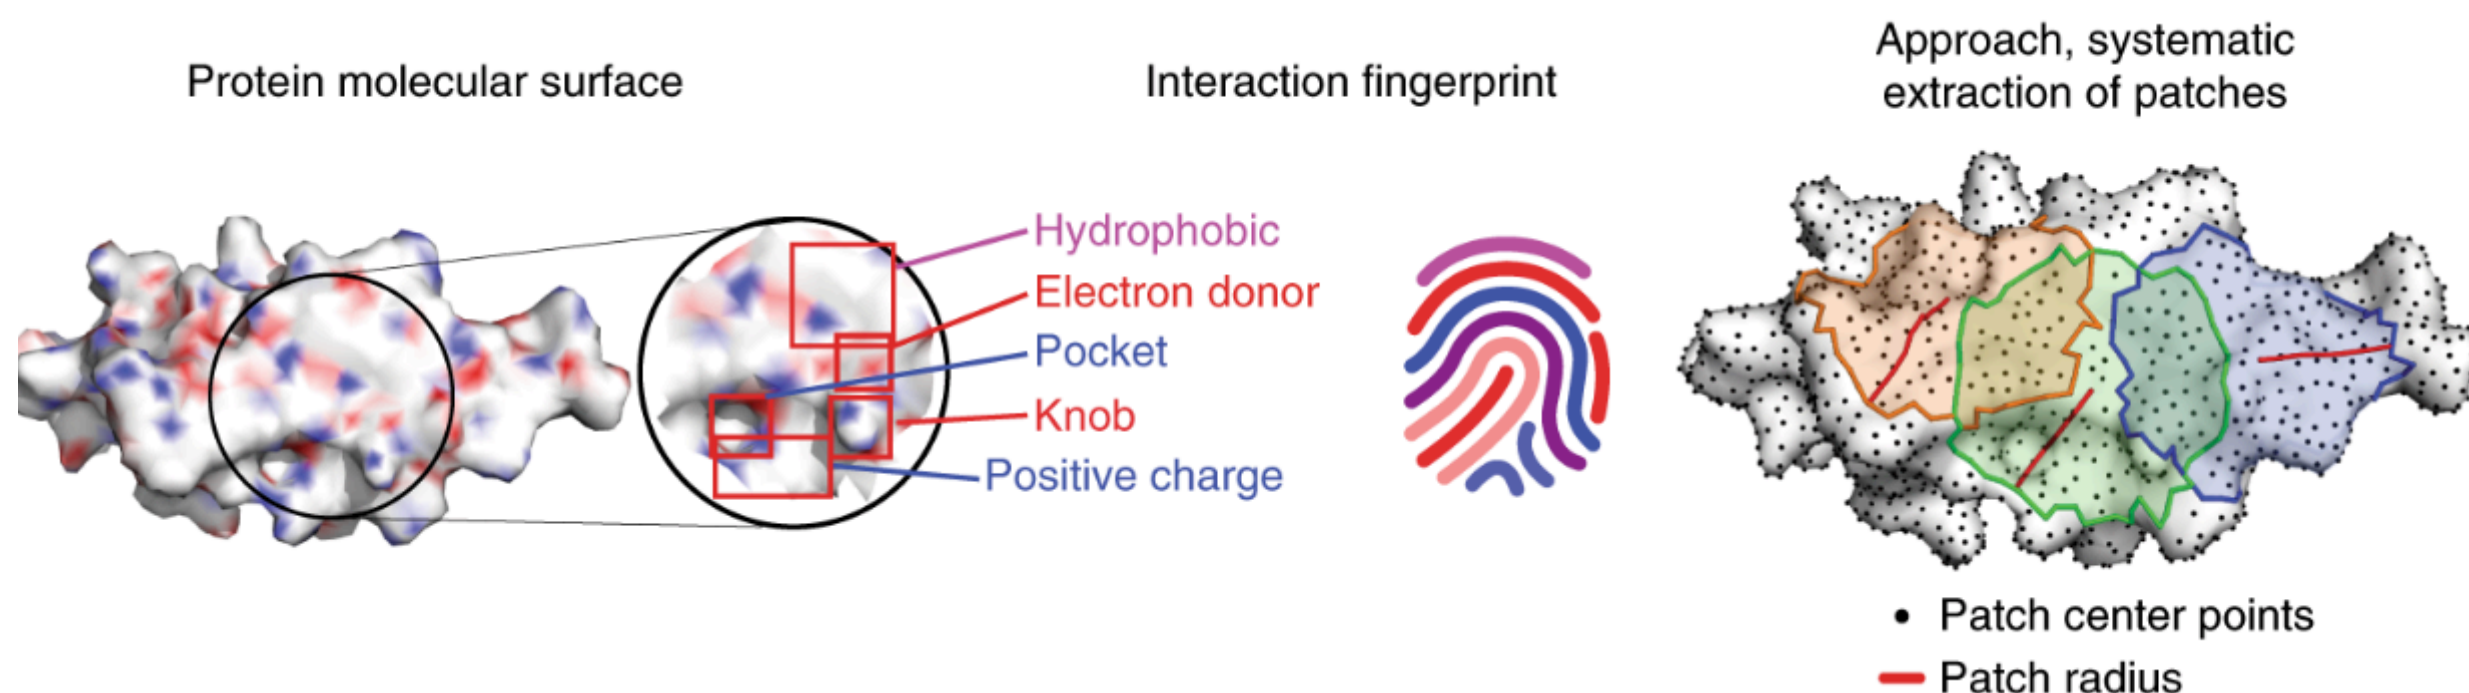
Neuroscience: learn to compare



Ktena et al., NeuroImage, 2018

Protein-protein interactions

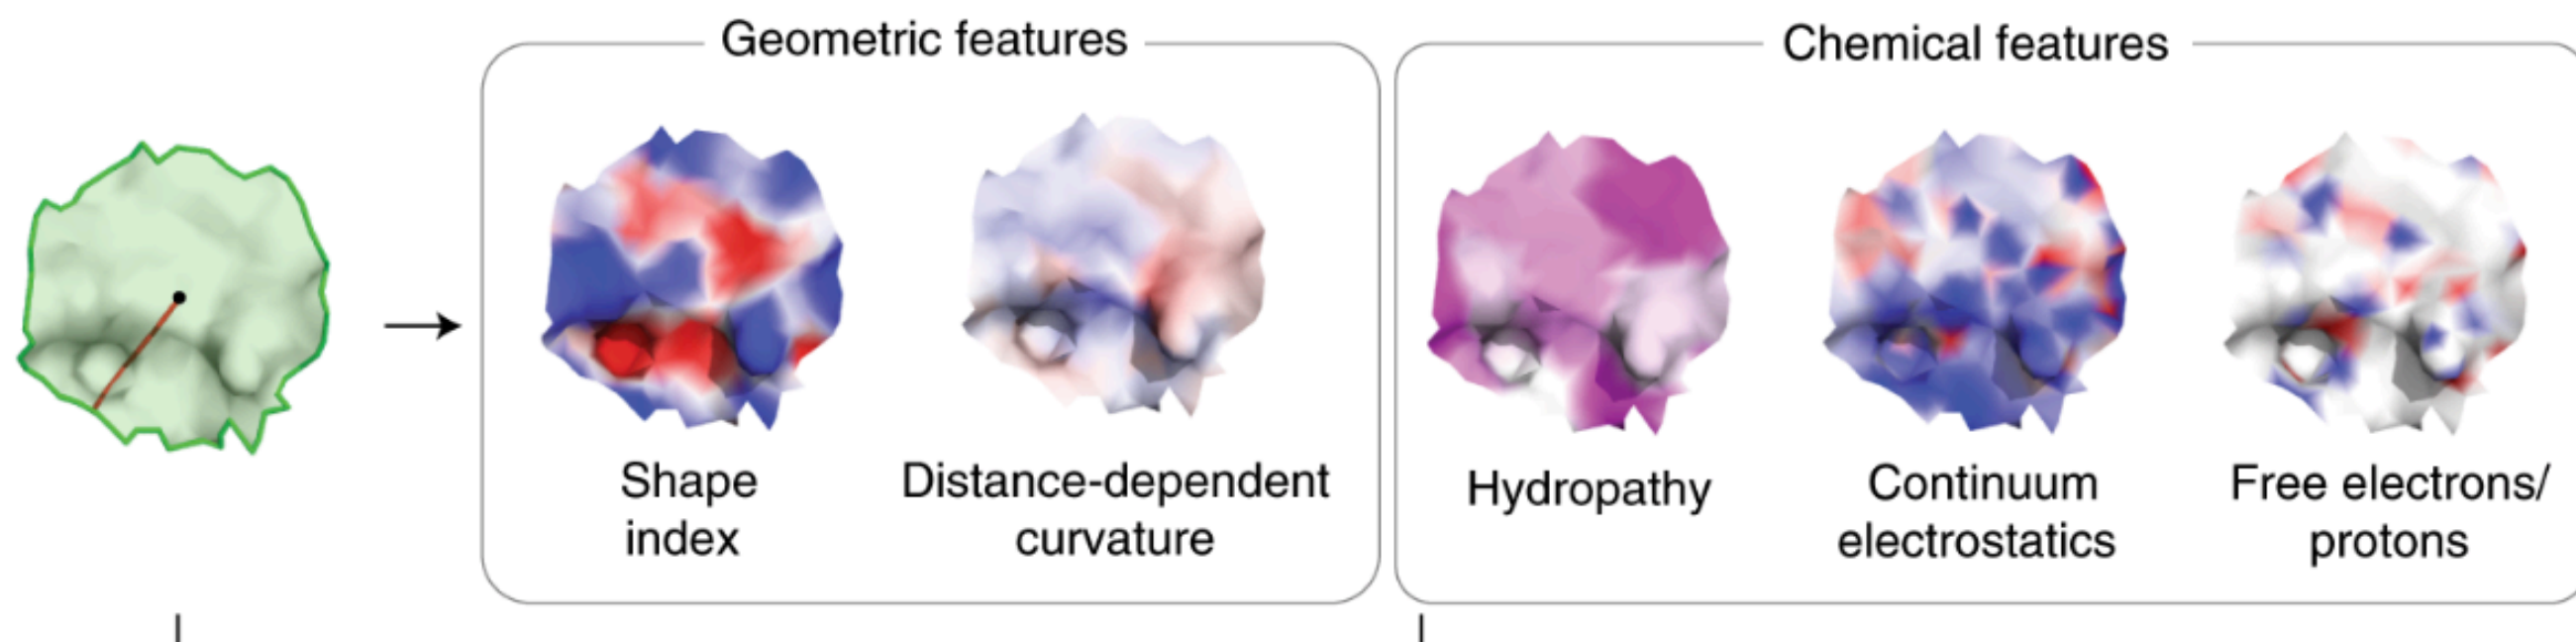
- Exploit GNNs to learn interaction fingerprints in protein molecular surfaces



Gainza et al, Nature methods, 2019

Protein-protein interactions

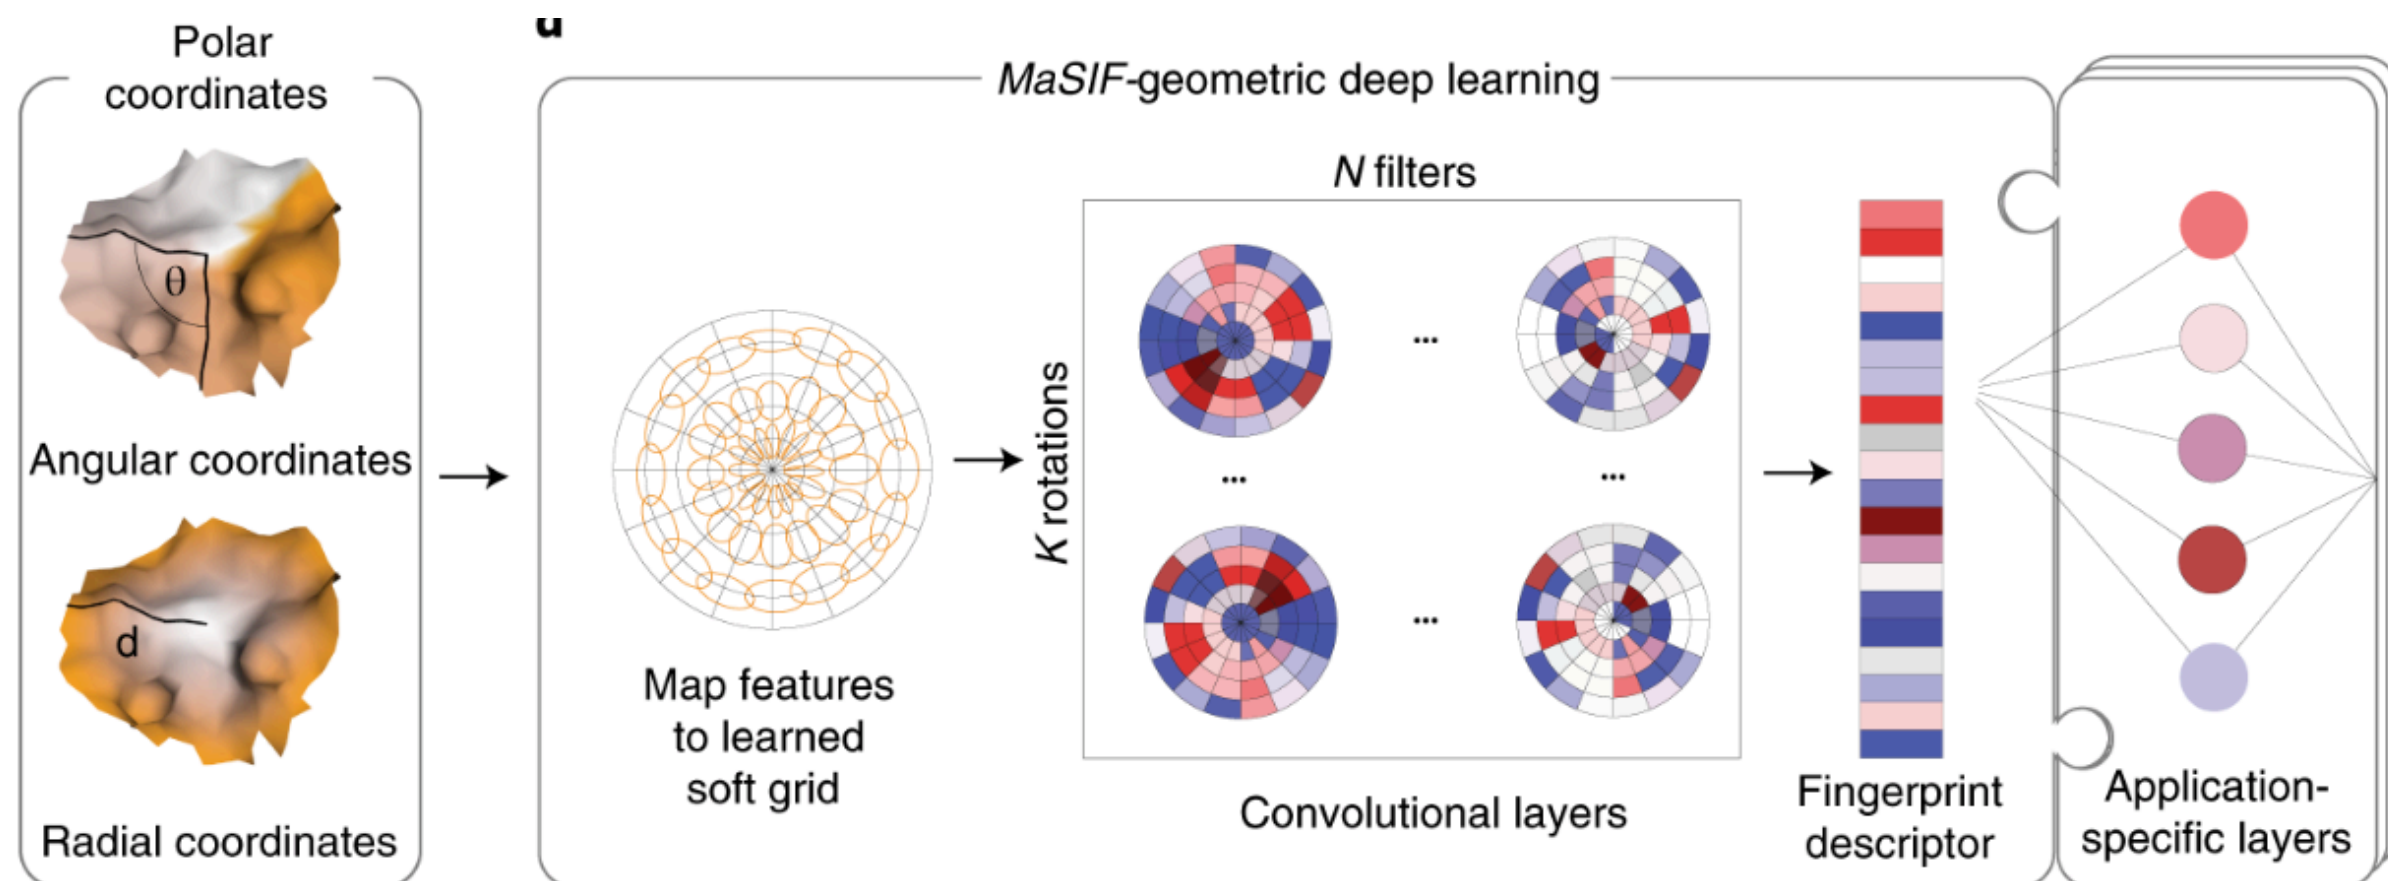
- Exploit GNNs to learn interaction fingerprints in protein molecular surfaces



Gainza et al, Nature methods, 2019

Protein-protein interactions

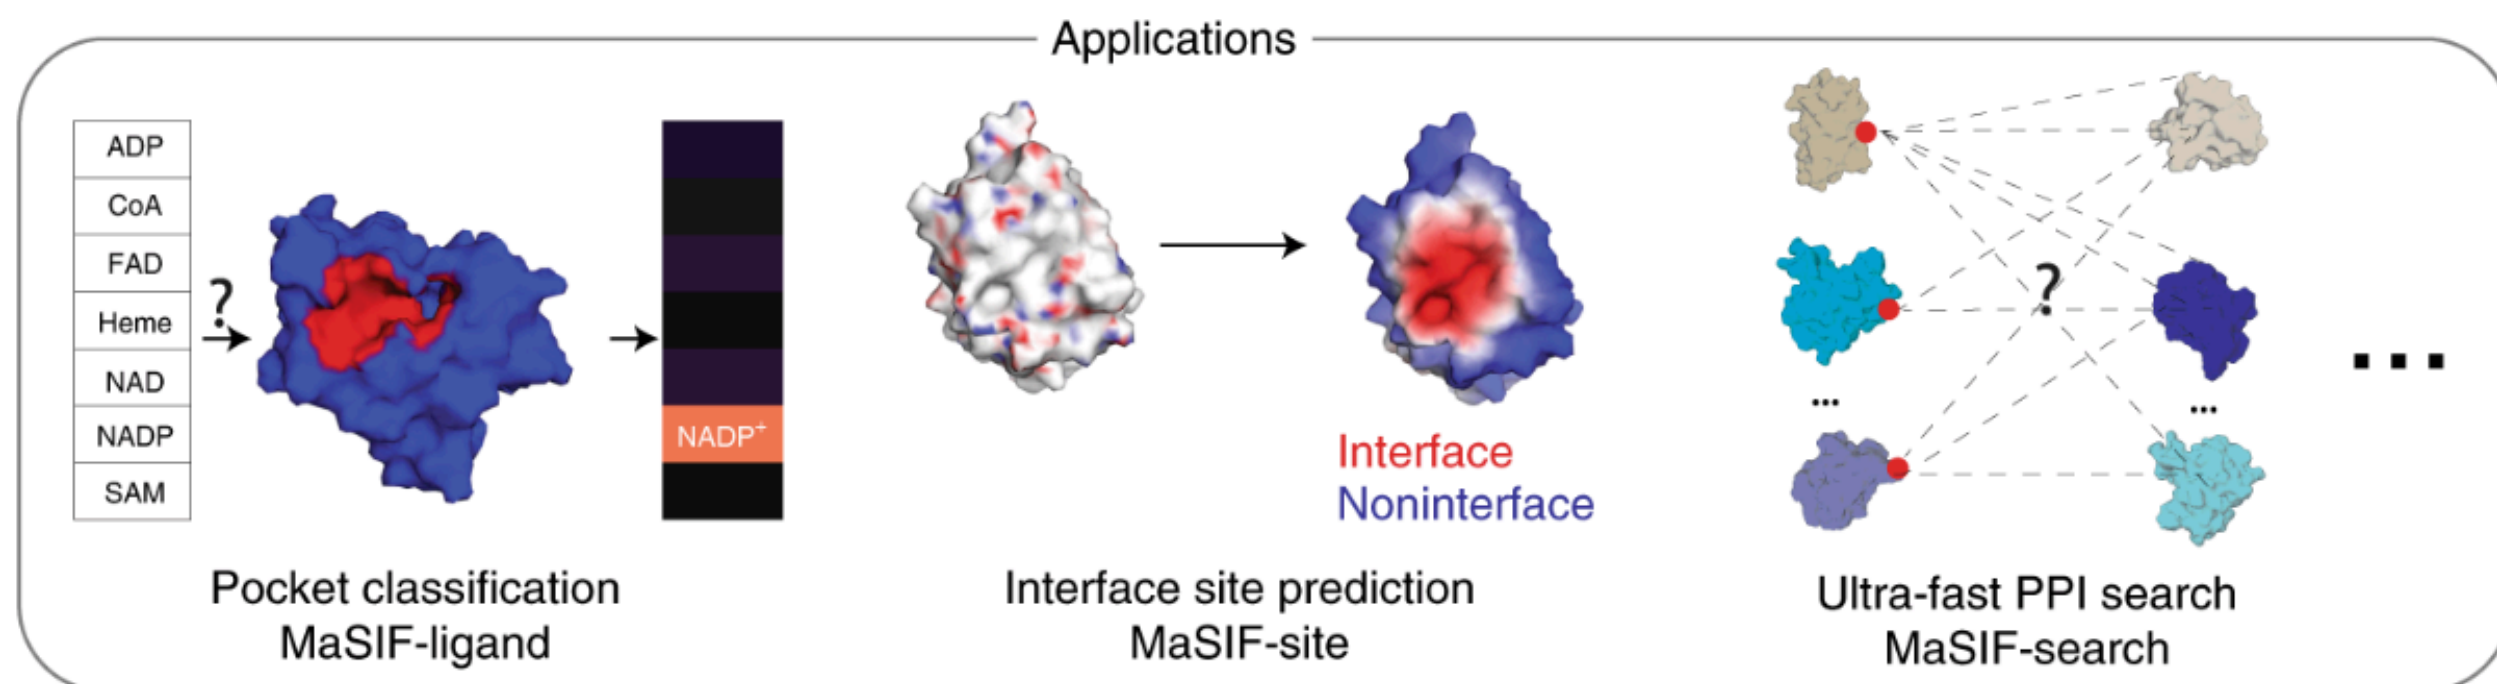
- Exploit GNNs to learn interaction fingerprints in protein molecular surfaces



Gainza et al, Nature methods, 2019

Protein-protein interactions

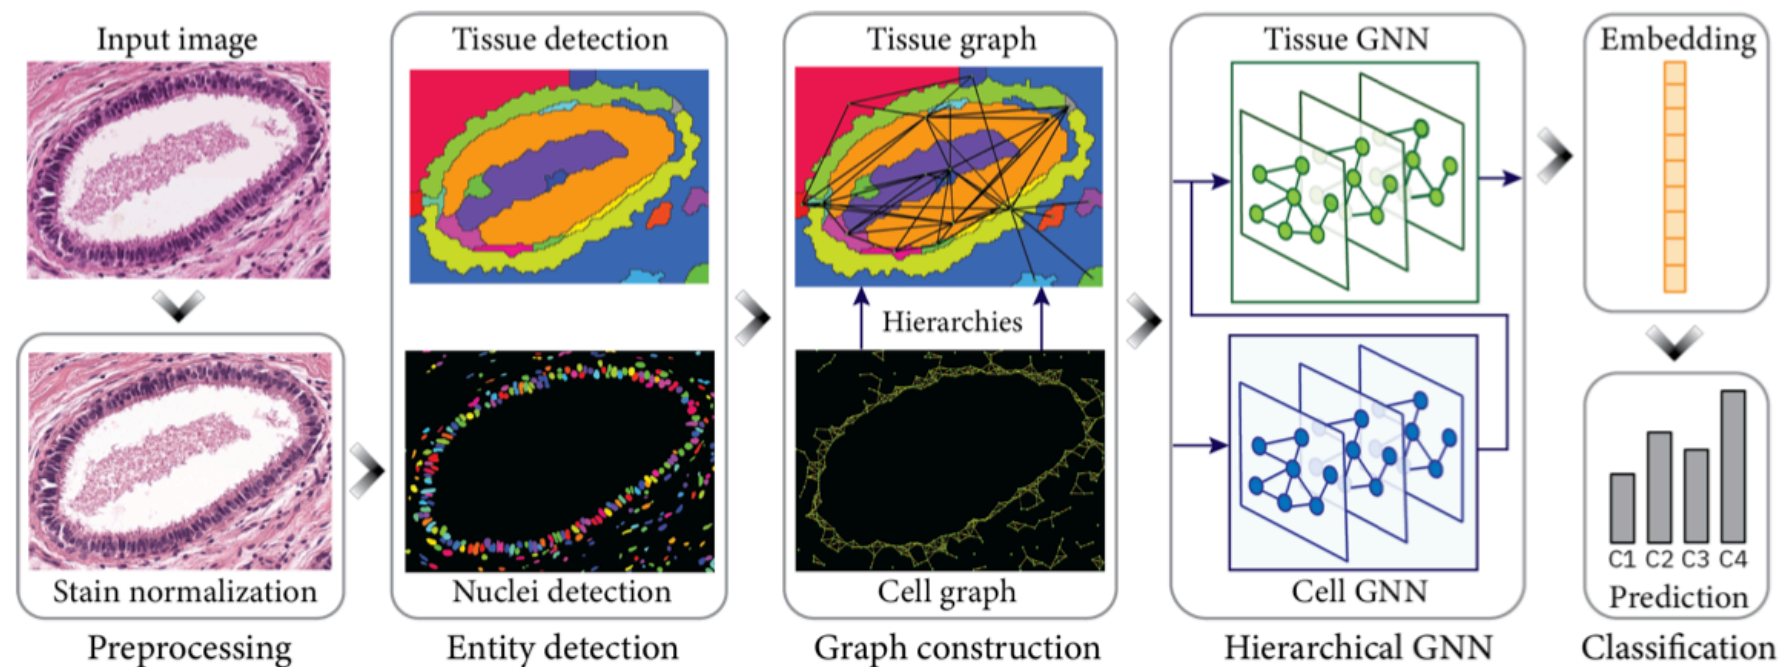
- Exploit GNNs to learn interaction fingerprints in protein molecular surfaces



Gainza et al, Nature methods, 2019

GNNs for medical imaging

- Digital pathology [1]



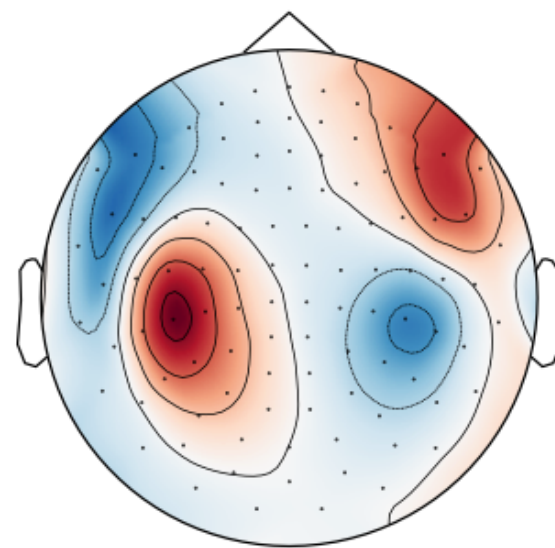
- Graph based representations provide a flexible tool for modelling complex dependencies at different levels of hierarchy (e.g., cells, tissues)

[1] Pati et al, "Hierarchical graph representation in digital pathology," arXiv, 2021

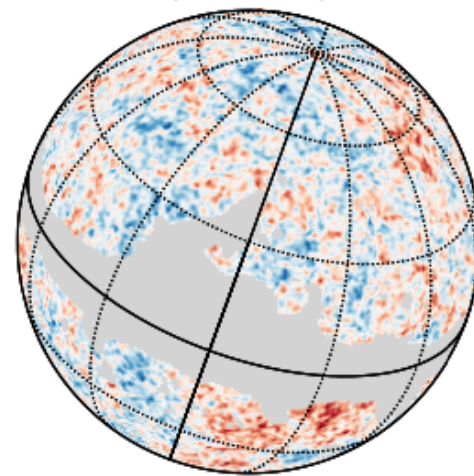
[2] Li et al, Representation learning for networks in biology and medicine: Advancements, challenges, and opportunities, arXiv, 2021

GNNs for spherical imaging

- Spherical data has specific spatial and statistical properties that cannot be captured by regular CNN models

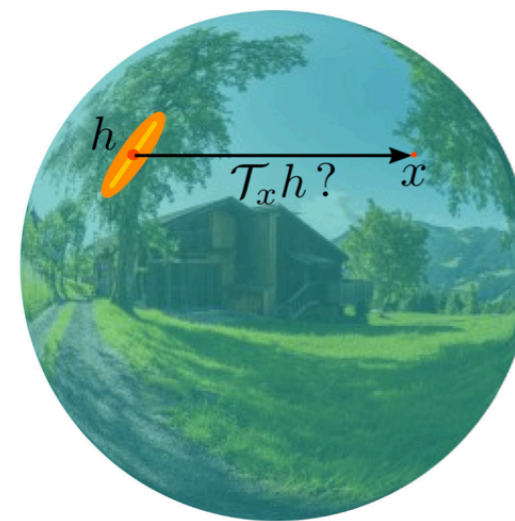


Brain activity (MEG)



-0.00025 0.00025

Cosmic microwave background temperature



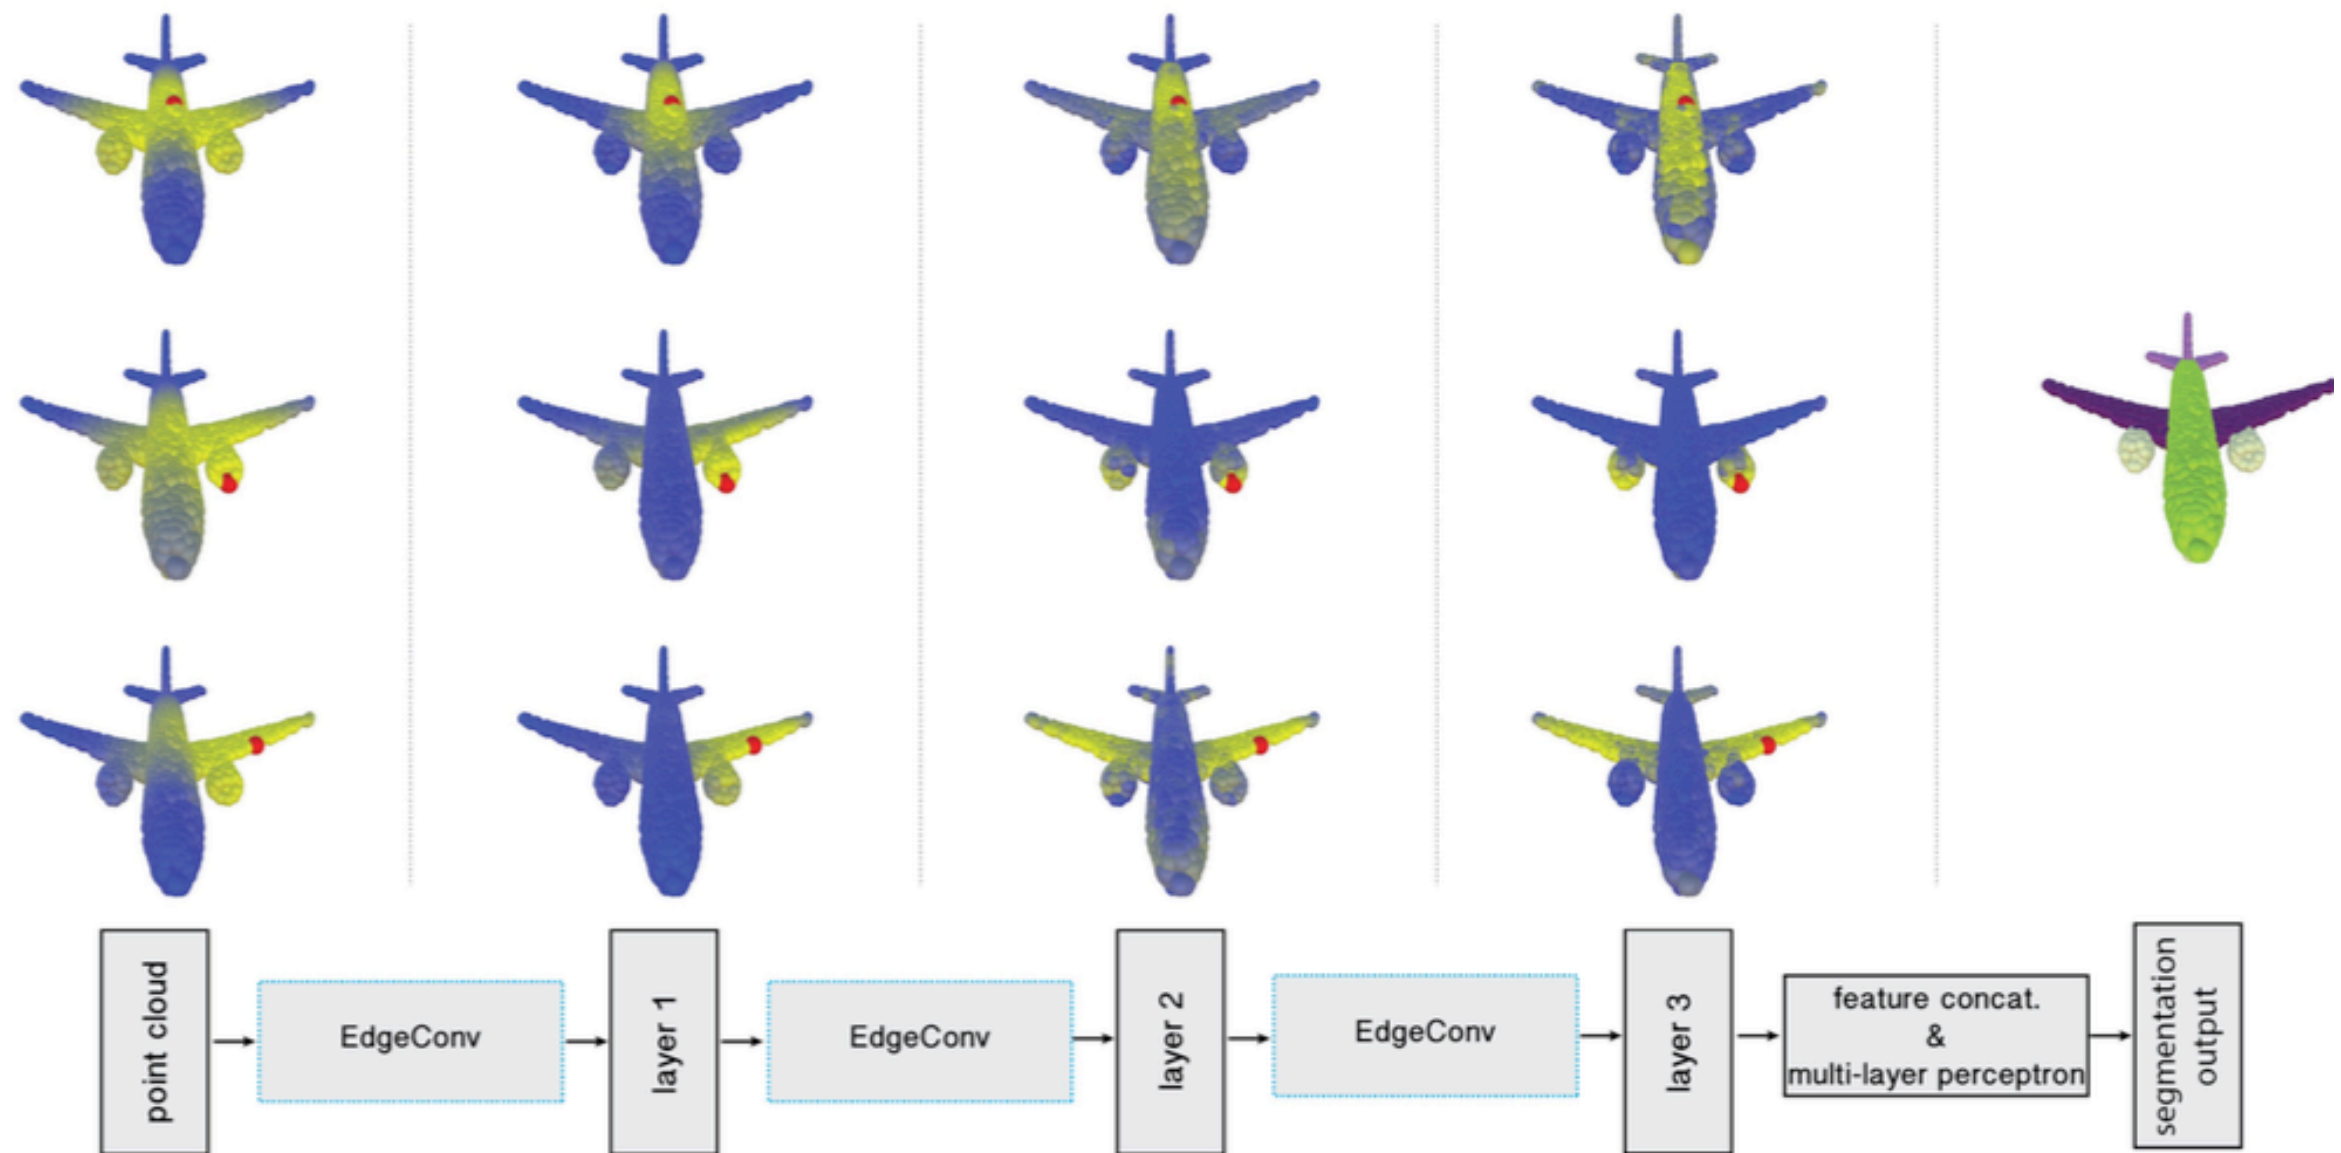
Omnidirectional images

- Sphere is modelled as a graph and classical operation (convolution, translation, pooling...) are performed on the graph

Perraudin et al., "DeepSphere", Astronomy and Computing, 2019

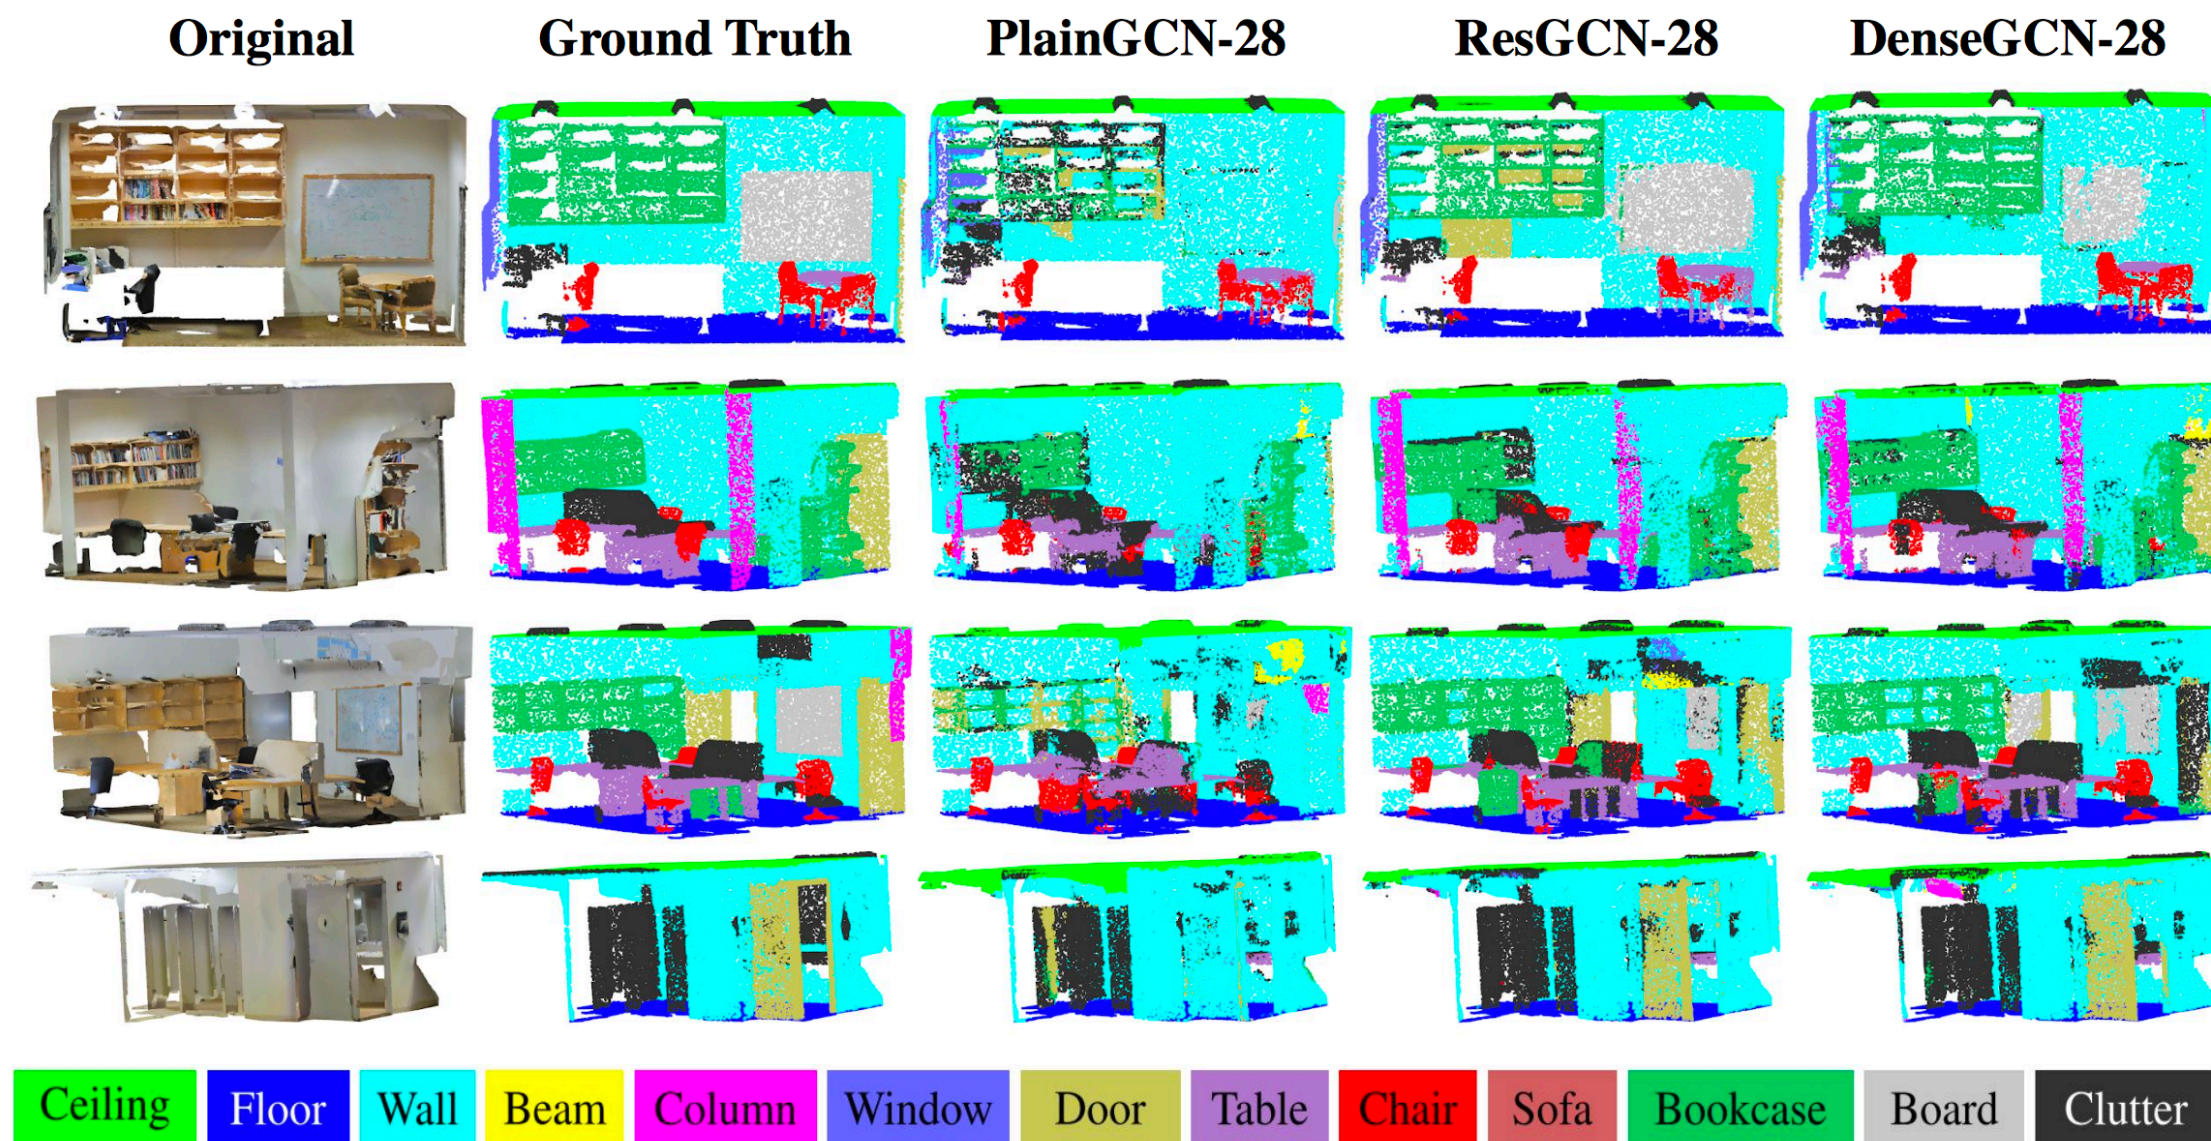
Bidgoli et al, OSLO: On-the-Sphere Learning for Omnidirectional images and its application to 360-degree image compression, arXiv, 2021

Point cloud semantic



Wang et al., 2019

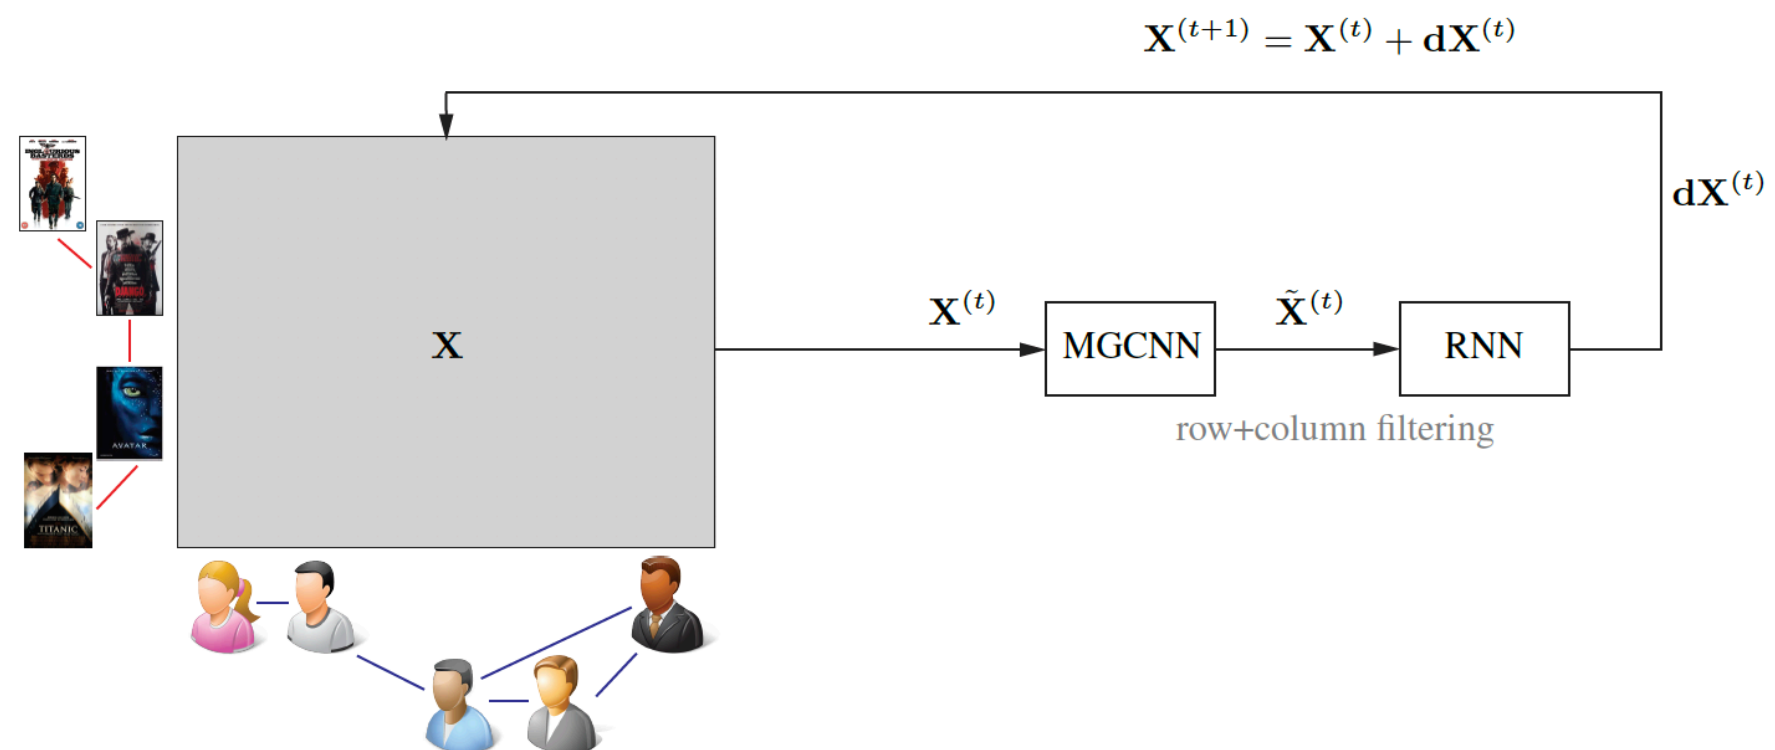
3D point clouds semantic



Li et al. 2019

Recommender systems

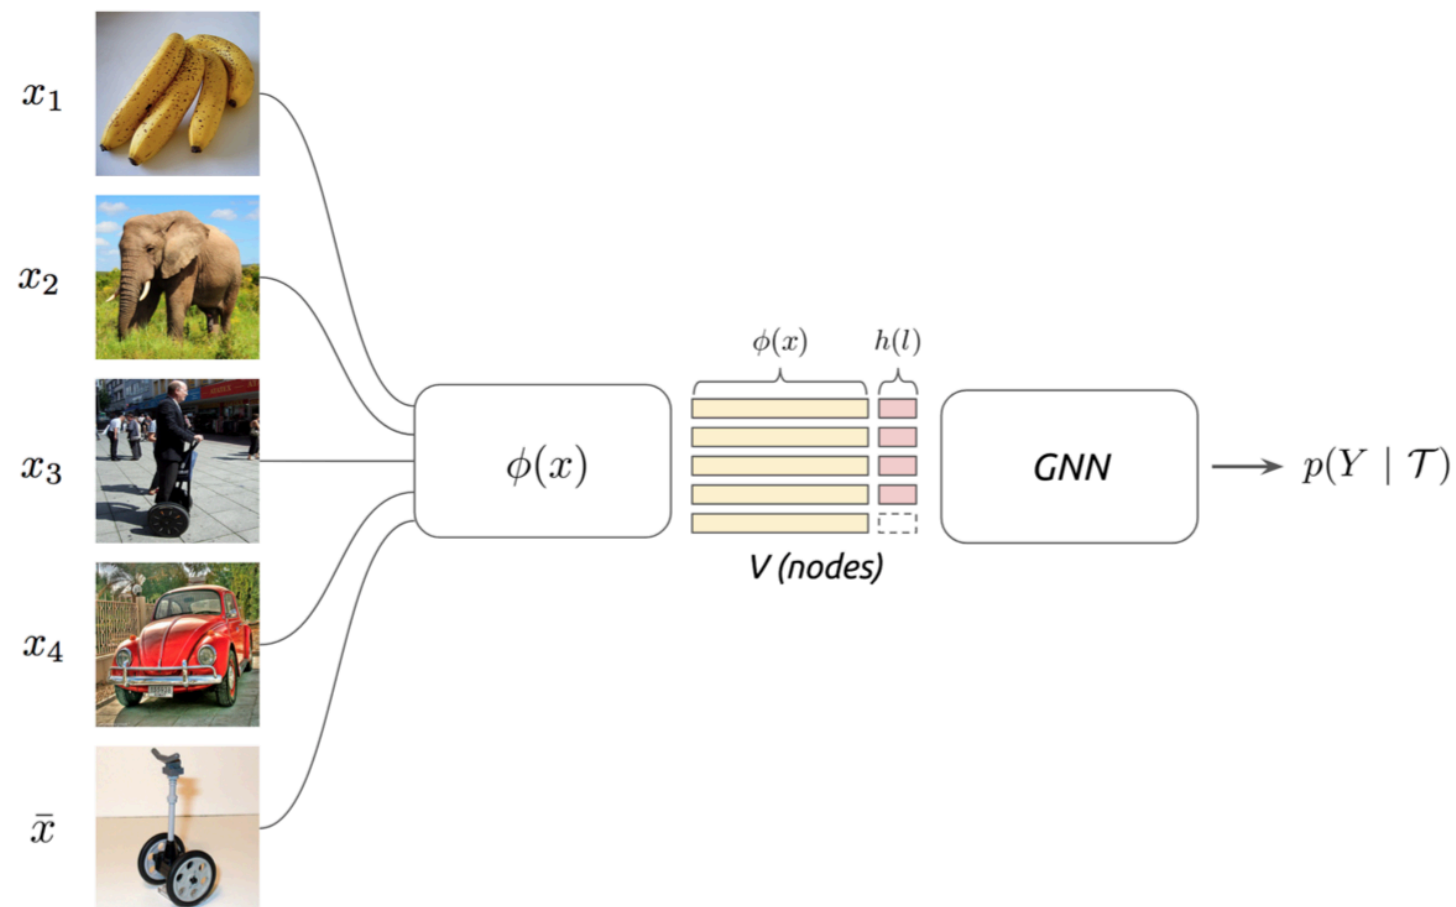
- Matrix completion: deep learning on user and item graphs
- Multi-graph convolution (spatial features), followed by LSTM (diffusion process)



Monti et al. NIPS, 2017

Few-shot learning

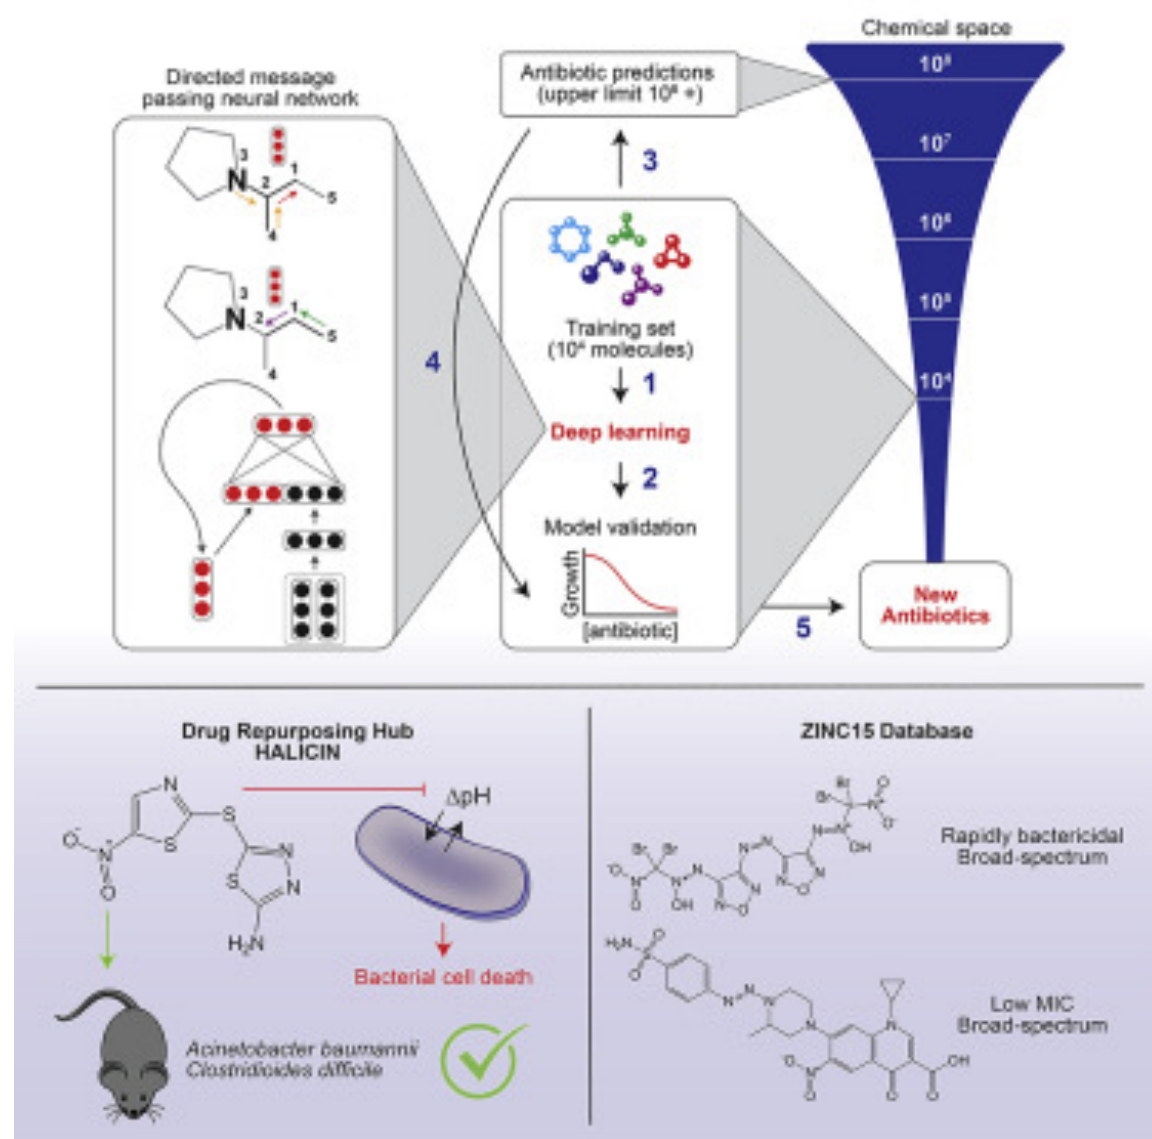
- Represent images as a fully connected graph
- Pose the problem as a supervised task using GNN



Garcia et al., ICLR, 2018

Molecular graph generation

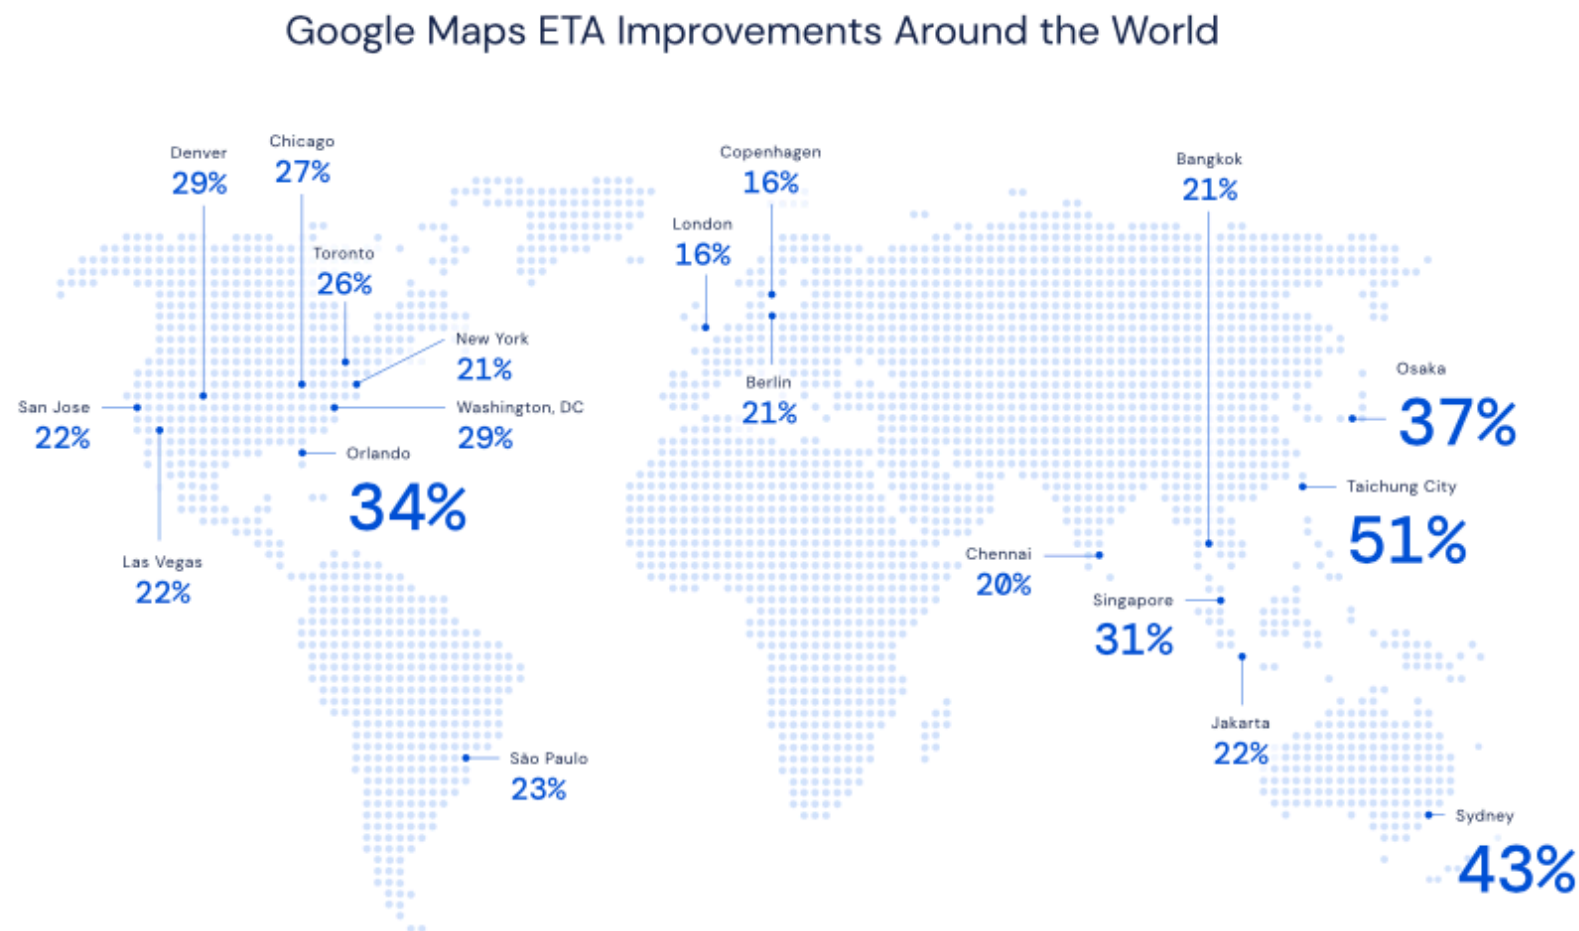
- Recent advances in antibiotic discovery ...



The screenshot shows a BBC News article titled "Scientists discover powerful antibiotic using AI". The article is dated 21 February 2020. The main image shows a scientist in a lab coat and mask looking through a microscope. Below the image, there is a section titled "Support the Guardian" with a "Contribute" button. The article text includes the headline "Powerful antibiotic discovered using machine learning for first time" and a sub-headline "Team at MIT says halicin kills some of the world's most dangerous strains". There is also a small image of petri dishes with bacterial cultures.

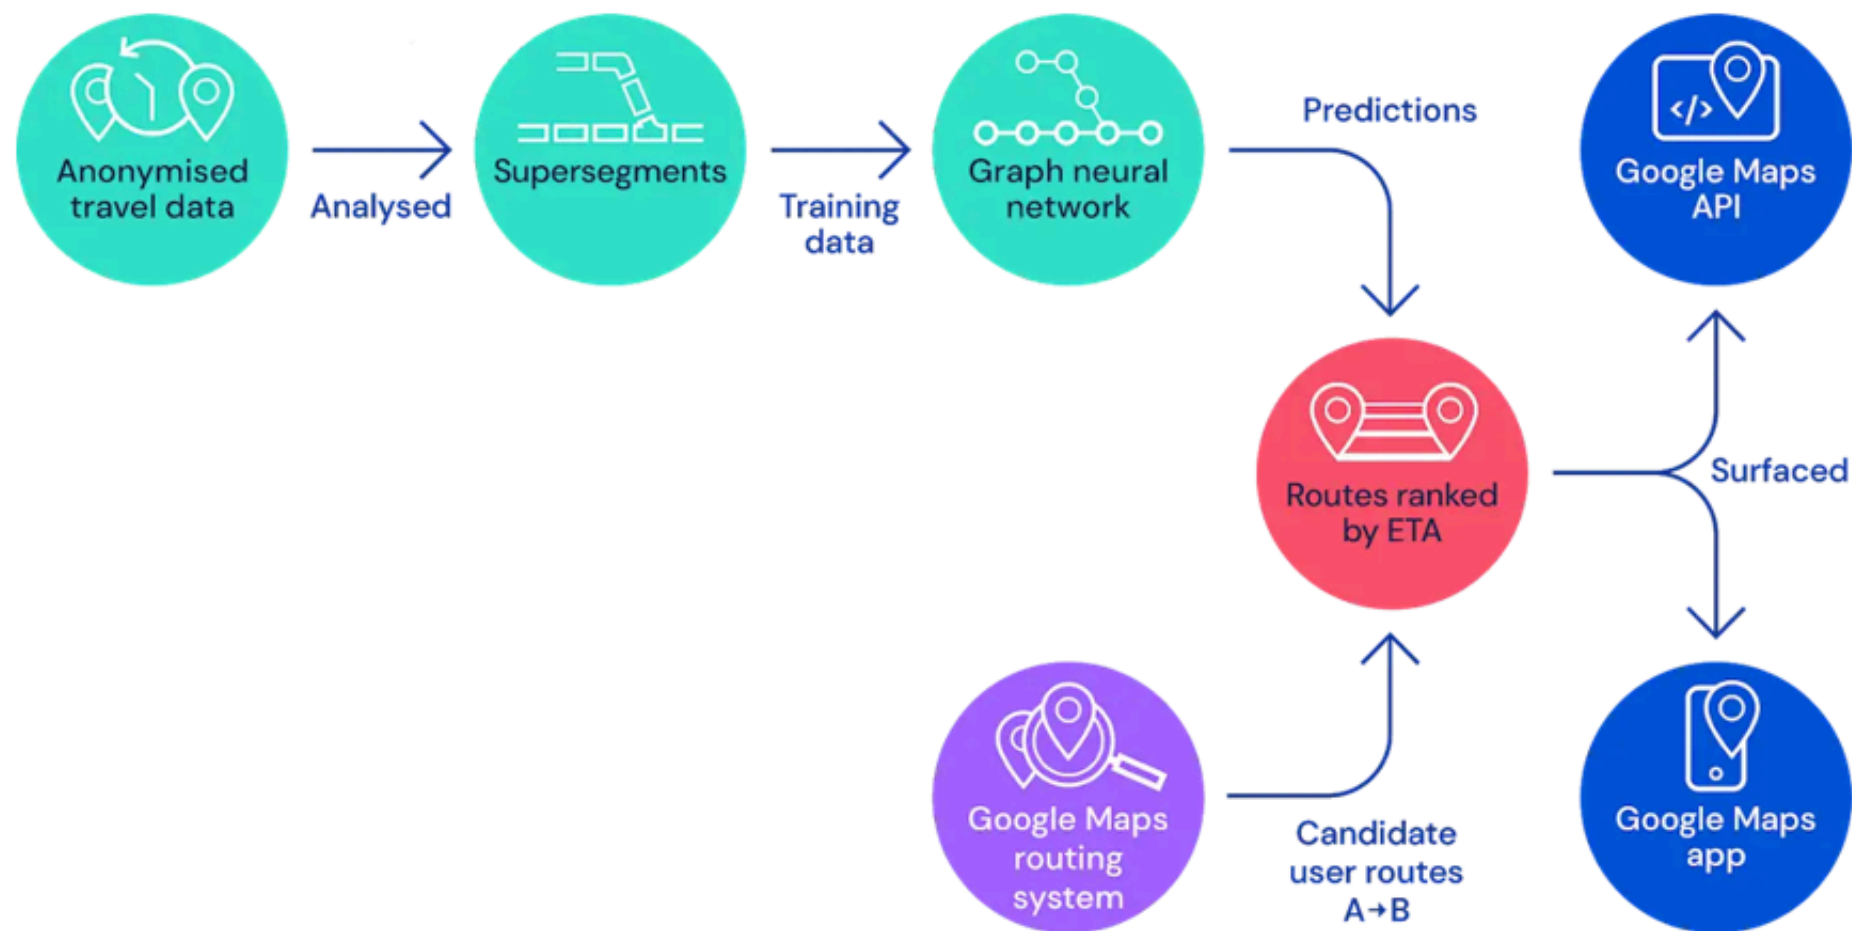
Simonovsky et al, 2017, De Cao et al 2018, Stokes et al 2020

Traffic prediction



[Derrow-Pinion et al., 2021]

Traffic prediction

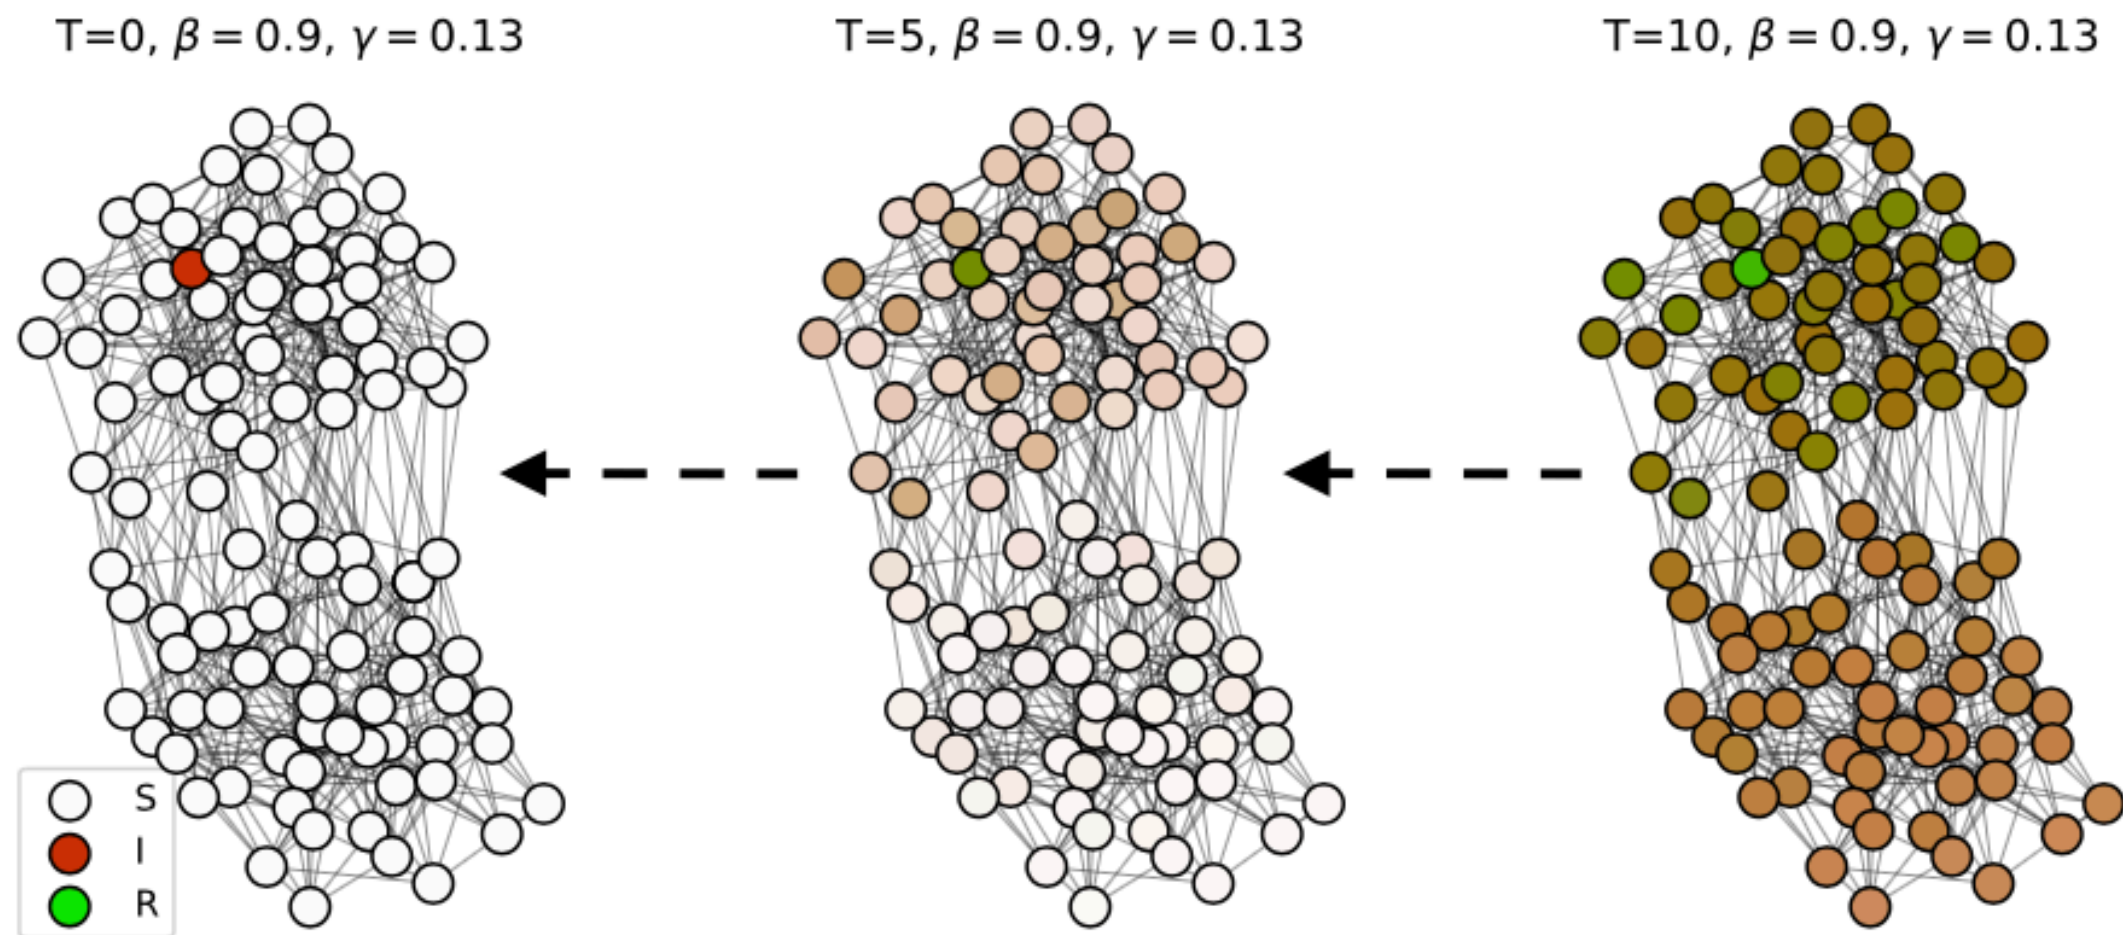


The model architecture for determining optimal routes and their travel time.

[Derrow-Pinion et al., 2021]

GNN for COVID-19

- Use GNN to locate the source of the epidemics [Shah et al. 2020]



Concluding remarks

- Deep learning on graphs
 - Emerging field that extends data analysis to irregular domains
 - Highly disciplinary topic: machine learning, signal processing, graph theory, harmonic analysis, statistics
 - Spectral- and spatial-domain approaches: different frameworks, significant overlaps
 - More and more applications are emerging

Open issues/Future directions

- Scalability issues
- Limited theoretical understanding
- Lack of performance guarantees; vulnerability to adversarial attacks
- Lack of interpretability
- Dealing with dynamic graphs
- Incorporating higher-order structures into GNNs
- Lack of standardized benchmarks
- <https://towardsdatascience.com/graph-deep-learning/home>

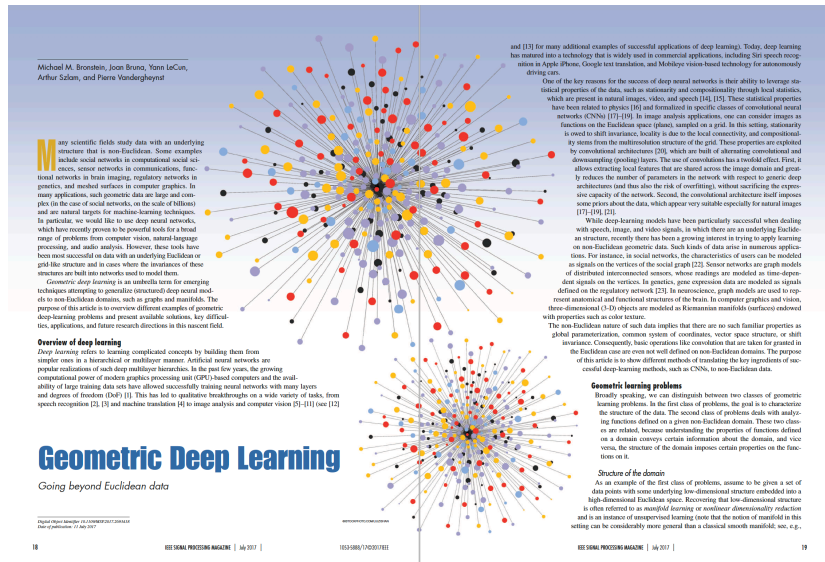
References

A Comprehensive Survey on Graph Neural Networks

Zonghan Wu, Shirui Pan, *Member, IEEE*, Fengwen Chen, Guodong Long, Chengqi Zhang, *Senior Member, IEEE*, Philip S. Yu, *Fellow, IEEE*

Abstract—Deep learning has revolutionized many machine learning tasks in recent years, ranging from image classification and video processing to speech recognition and natural language understanding. The data in these tasks are typically represented in the Euclidean space. However, there is an increasing number of applications where data are generated from non-Euclidean domains and are represented as graphs with complex relationships and interdependency between objects. The complexity of graph data has imposed significant challenges on existing machine learning algorithms. Recently, many studies on extending deep learning approaches for graph data have emerged. In this survey, we provide a comprehensive overview of graph neural networks (GNNs) in data mining and machine learning fields. We propose a new taxonomy to divide the state-of-the-art graph neural networks into different categories. With a focus on graph convolutional networks, we review alternative architectures that have recently been developed; these learning paradigms include graph attention networks, graph autoencoders, graph generative networks, and graph spatial-temporal networks. We further discuss the applications of graph neural networks across various domains and summarize the open source codes and benchmarks of the existing algorithms on different learning tasks. Finally, we propose potential research directions in this fast-growing field.

Index Terms—Deep Learning, graph neural networks, graph convolutional networks, graph representation learning, graph autoencoder, network embedding



Geometric Deep Learning Grids, Groups, Graphs, Geodesics, and Gauges

Michael M. Bronstein¹, Joan Bruna², Taco Cohen³, Petar Veličković⁴

May 4, 2021

https://www.cs.mcgill.ca/~wlh/grl_book/

<https://github.com/DeepGraphLearning/LiteratureDL4Graph>

<https://github.com/thunlp/NRLPapers>

<https://github.com/thunlp/GNNPapers>

Useful resources

- **Toolboxes**

- https://github.com/rusty1s/pytorch_geometric
- <https://github.com/dmlc/dgl>

- **Datasets**

- <https://chrsmrrs.github.io/datasets/>
- <https://ogb.stanford.edu>

Thank you for your attention!