

Backpropagation-free Graph Neural Networks

Luca Pasa

Department of Mathematics
University of Padova
Padova, Italy
luca.pasa@unipd.it

Nicolò Navarin

Department of Mathematics
University of Padova
Padova, Italy
nicolo.navarin@unipd.it

Wolfgang Erb

Department of Mathematics
University of Padova
Padova, Italy
wolfgang.erb@unipd.it

Alessandro Sperduti

Department of Mathematics
University of Padova
Padova, Italy
DISI, University of Trento
Trento, Italy
alessandro.sperduti@unipd.it

Abstract—We propose a class of neural models for graphs that do not rely on backpropagation for training, thus making learning more biologically plausible and amenable to parallel implementation in hardware. The base component of our architecture is a generalization of Gated Linear Networks which allows the adoption of multiple graph convolutions. Specifically, each neuron is defined as a set of graph convolution filters (weight vectors) and a gating mechanism that, given a node and its topological context, selects the weight vector to use for processing the node’s attributes. Two different graph processing schemes are studied, i.e., a message-passing aggregation scheme where the gating mechanism is embedded directly into the graph convolution, and a multi-resolution one where neighbouring nodes at different topological distances are jointly processed by a single graph convolution layer. We also compare the effectiveness of different alternatives for defining the context function of a node, i.e., based on hyper-planes or on prototypes. A theoretical result on the expressiveness of the proposed models is also reported. We experimented our backpropagation-free graph convolutional neural architectures on commonly adopted node classification datasets, and show competitive performances compared to the backpropagation-based counterparts.

Index Terms—Graph Convolutional Networks, Graph Neural Network, Deep Learning, Structured Data, Machine Learning on Graphs

I. INTRODUCTION

In the last years, several definitions of neural architectures capable to deal with data in structured form, such as graphs, have been presented [1], [2]. The vast majority of graph neural networks in literature are based on the idea of message passing [3], in which the representation of a node at layer l (or time t if the network is recurrent) is defined as a transformation of the label of the same node and of its neighbors at layer $l - 1$ ($t - 1$).

While many works focused on defining alternative architectures, at the best of author’s knowledge all of them rely on backpropagation to learn the networks’ weights. Backpropagation is a powerful and effective method to train deep neural networks (NNs), that has been successfully applied almost ubiquitously in recent years. When the amount of available data is not huge, however, the standard approach of training a non-linear NN with backpropagation may quickly lead to overfit the training

data. This is in clear contrast to how humans learn, since we do not require nearly the amount of training data modern NNs do to learn how to generalize [4]. Moreover, the backpropagation mechanism is not biologically plausible [5], [6], suggesting that the brain may use different learning algorithms.

Recently, some alternative definitions of multilayer neural networks that do not rely on backpropagation for their training [4], [6] have been proposed. They define local learning rules where each neuron, given its inputs, is trained independently from the rest of the network exploiting a global error signal. This approach allows these networks to: (i) be more biologically plausible (i.e. from the current knowledge about the functioning of animal neurons, it seems implausible for a neuron to have access to the connections in a brain area responsible for a subsequent processing step); (ii) be more sample efficient/simplify the overall training procedure, since each neuron solves an independent (possibly convex) problem;

The aim of this paper is to explore a contamination between these two cutting-edge research fields, studying how to define neural networks for graph processing that do not rely on backpropagation for their training.

Our exploration is based on the recently proposed Gated Linear Networks (GLN) [4], a family of backpropagation-free neural networks that have been developed for online learning and that have shown promising results. The main characteristic of such networks is that, contrarily to the mainstream approach, the non-linearity is achieved via a gating mechanism instead of element-wise non-linear functions. More specifically, each neuron receives a context vector as additional input, that is used to select one weight vector in a pre-defined set. The only non-linearity lies in such gating mechanism. In fact, once the weight vector is selected, each neuron behaves linearly. The resulting network is a piece-wise linear model (similarly to ReLU networks). While the gating mechanism is not trained, each neuron learns to predict (by modifying the weights) a binary output, and can be trained independently from the rest of the network. In the case of multi-class classification, a one-vs-all approach is exploited.

In order to define neural networks capable of processing graph data, we explore two possible generalizations of the above mechanism. The first one is based on the approach adopted by many graph convolutional networks, in which the network architecture reflects the structure of the input graph,

This work was partly funded by the SID/BIRD project *Deep Graph Memory Networks*, Department of Mathematics, University of Padova. The authors acknowledge the HPC resources of the Department of Mathematics, University of Padova, made available for conducting the research reported in this paper.

and node representations are refined at each layer according to the local graph topology via an aggregation operation over neighboring nodes. For this generalization, we also provide a theoretical result on the incremental expressiveness of our models. The second one is inspired to the so-call multi-resolution architecture on which the graph convolutional layer is defined by exploiting the power series of the diffusion operator (also known as Graph-Augmented Multi-Layer Perceptrons (GA-MLPs) [7], or Polynomial-Based Graph Convolutional (PGC) [8]). These models are able to simultaneously and directly consider all topological receptive fields up to k -hops. Moreover, the application of the gating mechanism turns out to be decoupled from the graph topology.

The properties inherited from GLNs ensure that our models are (at least in principle) as expressive as their backpropagation-based counterparts, with a significantly easier training phase. Nonetheless, several choices have to be made, such as which neighbor aggregation mechanism to adopt, how to define the contexts on graphs, and how to define the gating mechanism in an efficient way on graphs for limiting the number of parameters of the network while obtaining good predictive performances.

We experimentally evaluate our backpropagation-free graph convolutional neural architectures on commonly adopted node classification benchmarks, and verify their competitive performance. This work paves the way to novel neural approaches for graph learning.

II. BACKGROUND

In this section, we introduce the background notation and knowledge on which our model hinges.

A. Learning on graph nodes

A learning problem on a graph can be formulated as learning a function that maps nodes to labels. The underlying graph structure is given as $G = (V, E, \mathcal{L})$, where $V = \{v_1, \dots, v_n\}$ is the set of nodes, $E \subseteq V \times V$ is the set of edges connecting the nodes, and $\mathcal{L} : \mathcal{V} \rightarrow \mathbb{R}^s$ is a function associating a vector of attributes to each node. With $\mathcal{N}(v)$ we denote the set of nodes adjacent to v , i.e. $\mathcal{N}(v) = \{u \mid (v, u) \in E\}$. To simplify the notation, we define for a fixed graph G the matrix $\mathbf{X} = [\mathcal{L}(v_1), \dots, \mathcal{L}(v_n)]^\top$

Given a graph G , our training set is composed by the target information associated to some of the graph nodes, i.e., $\{(v, y) \mid v \in W, y \in \mathcal{Y}\}$ with $W \subset V$. For the sake of simplicity, in our presentation we will only consider binary values $\mathcal{Y} \in \{0, 1\}$.

B. Graph Neural Networks

Despite learning from graph data is not a new research field [9]–[11], in the last years graph neural networks have emerged as the machine learning model of choice when dealing with graph problems.

A Graph Neural Network (GNN) is a neural model that exploits the structure of the graph and the information embedded in feature vectors of each node in order to learn a representation $\mathbf{h}_v \in \mathbb{R}^m$ for each vertex $v \in V$. In many GNN

models, the computation of \mathbf{h}_v can be divided in two main steps: *aggregate* and *combine*. We can define aggregation and combination by using two functions, \mathcal{A} and \mathcal{C} , respectively: $\mathbf{h}_v = \mathcal{C}(\mathcal{L}(v), \mathcal{A}(\{\mathcal{L}(u) : u \in \mathcal{N}(v)\}))$.

The choice of aggregation function \mathcal{A} and combination function \mathcal{C} defines the type of *Graph Convolution (GC)* adopted by the GNN. In [12], the first model that uses graph convolutions is introduced. In the last few years, several different GCs have been proposed [13]–[19].

In this work, we build on top of two widely adopted graph convolutions. The first one is the GCN [13]

$$\mathbf{H}^{(i)} = \mathcal{F}\left(\tilde{\mathbf{D}}^{-\frac{1}{2}}(\mathbf{I} + \mathbf{A})\tilde{\mathbf{D}}^{-\frac{1}{2}}\mathbf{H}^{(i-1)}\mathbf{W}^{(i)}\right), i > 1 \quad (1)$$

where \mathbf{A} denotes the standard adjacency matrix of the graph G and $\tilde{\mathbf{D}}$ a diagonal degree matrix with the diagonal elements defined as $\tilde{d}_{ii} = 1 + \sum_j a_{ij}$. Further, $\mathbf{H}^{(i)} \in \mathbb{R}^{n \times m_i}$ is a matrix containing the representation $\mathbf{h}_v^{(i)}$ of all nodes in the graph (one per row) at layer i , $\mathbf{W}^{(i)} \in \mathbb{R}^{m_{i-1} \times m_i}$ denotes the matrix of the layer's parameters, and \mathcal{F} is the element-wise (usually, nonlinear) activation function.

The second graph convolution we consider is a slight variation of the first model and commonly referred to as GraphConv [2]:

$$\mathbf{h}_v^{(i)} = \mathcal{F}(\mathbf{h}_v^{(i-1)}\mathbf{W}_1^{(i)} + \sum_{u \in \mathcal{N}(v)} \mathbf{h}_u^{(i-1)}\mathbf{W}_2^{(i)}),$$

where $\mathbf{W}_1^{(i)}, \mathbf{W}_2^{(i)} \in \mathbb{R}^{m_{i-1} \times m_i}$ (with $m_0 = s$, the input dimensionality) are the network parameters.

C. Multi-Resolution GNN

Some recent works in literature exploit the idea of extending graph convolution layers to increase the receptive field size, without increasing the depth of the model. The basic idea underpinning these methods is to consider the case in which the graph convolution can be expressed as a polynomial of the powers of a transformation $\mathcal{T}(\cdot)$ of the adjacency matrix. The models based on this idea are able to simultaneously and directly consider all topological receptive fields up to l -hops, just like the ones that are obtained by a stack of graph convolutional layers of depth l , without incurring in the typical limitations related to the complex interactions among the parameters of the GC layers. Formally, the idea is to define a representation as built from the contribution of all topological receptive fields up to l -hops as:

$$\mathbf{H} = f(\mathcal{T}(\mathbf{A})^0\mathbf{X}, \mathcal{T}(\mathbf{A})^1\mathbf{X}, \dots, \mathcal{T}(\mathbf{A})^l\mathbf{X}), \quad (2)$$

where $\mathcal{T}(\cdot)$ is a transformation of the adjacency matrix, (e.g., the Laplacian matrix), and f is a function that aggregates and transforms the various components obtained from the powers of $\mathcal{T}(\mathbf{A})$, for instance the concatenation, the summation, or even something more complex, such as a multi-layer perceptron. The f function can be defined as a parametric function, depending on a set of parameters θ whose values can be estimated from data (e.g., when f involves an MLP). Various models that rely on this idea have been proposed in the last few years [7], [8], [20]–[24].

D. Backpropagation-free neural networks

In this paper we exploit recently defined neurons that can be trained locally and independently instead of exploiting backpropagation. We consider the recently proposed Gated Linear Networks [25] (GLNs) where the local optimization problem obtained for each neuron, adopting an appropriate loss, is convex. Moreover, it has been shown that GLNs can represent any function that represent a probability arbitrarily well [26].

The main differences of GLNs compared to MLPs are the following: first, each neuron in a GLN is a Gated Geometric Mixer. Geometric mixing [27] is an ensemble technique that assigns a weight to each weak predictor in input. In GLNs, every unit produces in output its prediction for the target. Given an input vector of probabilities $\mathbf{p} = [p_1, \dots, p_n]^\top$, geometric mixing is defined as:

$$\sigma(\mathbf{w}^\top \sigma^{-1}(\mathbf{p}))$$

where $\sigma(x) = \frac{1}{1+e^{-x}}$ is the sigmoid function, $\sigma^{-1}(x) = \text{logit}(x) = \log(x) - \log(1-x)$ is the logit function (that is the inverse of the sigmoid function), and both of them are applied element-wise.

To achieve non-linearity, specifically piecewise-linearity, GLNs employ a gating mechanism in each neuron. Each neuron divides its input space in *regions*. A geometric mixing (that is a linear model) is associated to each region. The association from examples to regions is carried by a *region assignment* function c . GLNs assume that for each example we have a vectorial representation available, $\mathbf{x} \in \mathbb{R}^{d(x)}$, and a vector representing side-information (or context), i.e. $\mathbf{z} \in \mathbb{R}^{d(z)}$. The c function is defined depending on side-information associated to each input (in case no side-information is available, it is possible to set $\mathbf{z} = \mathbf{x}$). Each neuron in a GLN solves a convex problem, and is trained independently to predict the target.

For the sake of simplicity, we omit bias terms in the following formulations. Given a neuron j at the i -th layer, its output is defined as:

$$\mathbf{h}_{j,(\mathbf{x},\mathbf{z})}^{(i)} = \sigma\left(\sigma^{-1}\left(\mathbf{h}_{(\mathbf{x},\mathbf{z})}^{(i-1)}\right)^\top \mathbf{w}_{j,(\mathbf{z})}^{(i)}\right), i > 1 \quad (3)$$

with $\mathbf{h}_{(\mathbf{x},\mathbf{z})}^{(0)} = \sigma(\mathbf{x})$. The vector $\mathbf{w}_{j,(\mathbf{z})}^{(i)} \in \mathbb{R}^{m_{i-1}}$ stores the weights associated to the region activated by the context \mathbf{z} for the corresponding neuron. Let us discuss this weight vector in more detail and the gating mechanism that defines how a specific set of weights is selected.

Given an example (\mathbf{x}, \mathbf{z}) , we can select the weights of a single neuron j at the i -th layer as:

$$\mathbf{w}_{j,(\mathbf{z})}^{(i)} = \left(\Theta_j^{(i)} \mathbf{c}_{j,(\mathbf{z})}^{(i)}\right) \quad (4)$$

where $\Theta_j^{(i)} \in \mathbb{R}^{m_{i-1} \times k}$, k is the number of regions (we assume for simplicity that each neuron in the network considers the same number of regions), and $\mathbf{c}_{j,(\mathbf{z})}^{(i)} \in \mathbb{R}^k$. Notice that the main characteristic of a Gated Linear Neuron is that, instead of having a single weight vector, each GL neuron depends on a *matrix* of parameters $\Theta_j^{(i)}$.

The original paper [25] proposes to implement the gating in the c functions with an halfspace-gating mechanism. Given a vector $\mathbf{z} \in \mathbb{R}^{d(z)}$, and a hyperplane with parameters $\mathbf{a}_i \in \mathbb{R}^{d(z)}$ and $b_i \in \mathbb{R}$, let us define a context function $\tilde{c}_i : \mathbb{R}^{d(z)} \rightarrow \{0, 1\}$ as:

$$\tilde{c}_i(\mathbf{z}) = \begin{cases} 1 & \text{if } \mathbf{a}_i^\top \mathbf{z} > b_i \\ 0 & \text{otherwise} \end{cases}$$

that divides $\mathbb{R}^{d(z)}$ in two half-spaces, according to the hyperplane $\mathbf{a}_i^\top \mathbf{z} = b_i$. We can compose $\log_2(k)$ (assuming k to be a power of 2) context functions of the same kind, obtaining an higher-order context function $\tilde{c} : \mathbb{R}^{d(z)} \rightarrow \{0, 1\}^{\log_2(k)}$, $\tilde{c} = [\tilde{c}_1, \dots, \tilde{c}_k]^\top$. We can then easily define a function f mapping from $\{0, 1\}^{\log_2(k)}$ to $\{0, \dots, k-1\} \subset \mathbb{N}$, obtaining the function $\hat{c} : \mathbb{R}^{d(z)} \rightarrow \{0, \dots, k-1\}$, $\hat{c} = f \circ \tilde{c} = f(\tilde{c}(\mathbf{z}))$. We can exploit the one-hot encoding of the output of such function and re-define it as $c : \mathbb{R}^{d(z)} \rightarrow \{0, 1\}^k$, $c = \text{one_hot}(\hat{c})$.

Given a layer i , each neuron j computes a different function $c_j^{(i)} : \mathbb{R}^{d(z)} \rightarrow \{0, 1\}^k$. For the j -th neuron at the i -th layer, the output of the context function applied to \mathbf{z} is thus the (one-hot) vector $\mathbf{c}_{j,(\mathbf{z})}^{(i)}$.

E. Layer-wise formulation

Exploiting the definition of a single Gated Linear Neuron in the previous section, we can define a whole GLN layer. This formulation will be exploited in the remainder of the paper. The output for the i -th layer in a GLN (with m_i neurons) for a sample (\mathbf{x}, \mathbf{z}) is defined as:

$$\mathbf{h}_{(\mathbf{x},\mathbf{z})}^{(i)} = \sigma\left(\sigma^{-1}\left(\mathbf{h}_{(\mathbf{x},\mathbf{z})}^{(i-1)}\right)^\top \mathbf{W}_{(\mathbf{z})}^{(i)}\right), i > 1 \quad (5)$$

where

$$\mathbf{W}_{(\mathbf{z})}^{(i)} = [\mathbf{w}_{1,(\mathbf{z})}^{(i)}, \dots, \mathbf{w}_{m_i,(\mathbf{z})}^{(i)}], \quad (6)$$

and $\mathbf{W}_{(\mathbf{z})}^{(i)} \in \mathbb{R}^{m_{i-1} \times m_i}$.

Several layers can then be stacked. For a binary classification problem, the last layer will comprise a single neuron, i.e. for a network with l layers we have $\mathbf{W}_{(\mathbf{z})}^{(l)} = \mathbf{w}_{1,(\mathbf{z})}^{(l)}$, $\mathbf{W}_{(\mathbf{z})}^{(l)} \in \mathbb{R}^{m_{l-1} \times 1}$. The resulting model is, by construction, piecewise-linear. Specifically, given a context \mathbf{z} , the model is (up to a final activation function) linear and can be written as:

$$y_{(\mathbf{x},\mathbf{z})} = \sigma\left(\mathbf{x}^\top \mathbf{W}_{(\mathbf{z})}^{(1)} \dots \mathbf{W}_{(\mathbf{z})}^{(l-1)} \mathbf{W}_{(\mathbf{z})}^{(l)}\right) = \sigma\left(\mathbf{x}^\top \mathbf{w}_{(\mathbf{z})}\right) \quad (7)$$

with a weight vector $\mathbf{w}_{(\mathbf{z})} \in \mathbb{R}^{d(x)}$.

III. BACK-PROPAGATION FREE GRAPH NEURAL NETWORKS

In this section, we define our proposed models, which generalize GLNs to graph-structured data. Firstly, we show how to embed the GLN idea into graph convolutions to build models based on the message-passing paradigm. Then, we propose a Multi-Resolution approach that allows to keep the features' propagation through the graph structure separate from the processing of the resulting node representations.

A. Message-Passing GLN

The core concept behind several definitions of Graph Neural Networks is the aggregation function used to obtain information about the local graph structure surrounding a graph node. The simplest aggregation mechanism involves just the summation over the representations of neighbouring nodes. For this simple mechanism, we obtain the following definition for a single layer in a Gated Linear Graph Neural Network:

$$\mathbf{h}_{(v,\mathbf{z})}^{(i)} = \sigma \left(\sigma^{-1} \left(\mathbf{h}_{(v,\mathbf{z})}^{(i-1)} \right)^\top \mathbf{W}_{(\mathbf{z})}^{(i,1)} + \left(\sum_{(u,\mathbf{z}') \in \mathcal{N}_v} \sigma^{-1} \left(\mathbf{h}_{(u,\mathbf{z}')}^{(i-1)} \right)^\top \mathbf{W}_{(\mathbf{z})}^{(i,2)} \right) \right), \quad (8)$$

for $i \geq 1$, and $\mathbf{h}_{v,\mathbf{z}}^{(0)} = \sigma(\mathcal{L}(v))$. The weights $\mathbf{W}_{(\mathbf{z})}^{(i,2)}$, $\mathbf{W}_{(\mathbf{z})}^{(i,1)}$, and $\mathbf{W}_{(\mathbf{z})}^{(i)}$ are defined as per eq. (6) and can be obtained by a backpropagation-free training. This model can be considered as a modification of GraphConv proposed in [2] in which gated geometric mixing has been applied. For this reason we refer to this model also as Backpropagation Free - GraphConv (BF-GraphConv). Similarly to common formulations of graph neural networks, we can express the hidden representation for all the nodes in the graph as a single matrix. We obtain the following form of BF-GraphConv:

$$\mathbf{H}_{(\mathbf{z})}^{(i)} = \sigma^{-1} \left(\mathbf{H}_{(\mathbf{z})}^{(i-1)} \right) \mathbf{W}_{(\mathbf{z})}^{(i,1)} + \mathbf{A} \sigma^{-1} \left(\mathbf{H}_{(\mathbf{z})}^{(i-1)} \right) \mathbf{W}_{(\mathbf{z})}^{(i,2)} \quad (9)$$

where $\mathbf{H}_{(\mathbf{z})}^{(i)} \in \mathbb{R}^{n \times m_i}$ and $\mathbf{H}_{(\mathbf{z})}^{(0)} = \sigma(\mathbf{X})$. BF-GraphConv can therefore be regarded as a piecewise linear GNN depending on the context information \mathbf{z} of the neurons.

Following common definitions of Graph Neural Networks, we can resort to any message passing mechanism and define the Gated Linear counterpart. For instance, we can also consider the GCN presented in eq. (1) which leads to the following BF-GCN:

$$\mathbf{H}_{(\mathbf{z})}^{(i)} = \sigma \left(\tilde{\mathbf{D}}^{-\frac{1}{2}} (\mathbf{I} + \mathbf{A}) \tilde{\mathbf{D}}^{-\frac{1}{2}} \sigma^{-1} \left(\mathbf{H}_{(\mathbf{z})}^{(i-1)} \right) \mathbf{W}_{(\mathbf{z})}^{(i)} \right). \quad (10)$$

The resulting model can be regarded as a piecewise linear GCN. In particular, after l layers, the output $\mathbf{H}_{(\mathbf{z})}^{(l)}$ for the context \mathbf{z} can be written as

$$\mathbf{H}_{(\mathbf{z})}^{(l)} = \sigma \left(\left(\tilde{\mathbf{D}}^{-\frac{1}{2}} (\mathbf{I} + \mathbf{A}) \tilde{\mathbf{D}}^{-\frac{1}{2}} \right)^l \mathbf{X} \mathbf{w}_{(\mathbf{z})} \right), \quad (11)$$

i.e., the BF-GCN model with l layers is a generalization of the simple graph convolutional network (SGC) introduced in [20] and further investigated in [28], where the vector of weights $\mathbf{w}_{(\mathbf{z})}$ changes based on the input context. Notice that the main differences between the Gated Linear Graph Neural Networks and commonly adopted GNN formulations are the local training and the gating mechanism. A graphical layout of the proposed architecture is reported in Fig.1.




Fig. 1: A graphical layout of the proposed Message-Passing GLN, with an expanded view of the GLN neuron.

B. Incremental expressivity of GLNs

The usage of a region assignment function c for the side information $\mathbf{z} \in \mathbb{R}^{d_z}$ is equivalent to a partitioning $\mathcal{P}(c)$ of the space \mathbb{R}^{d_z} into k disjoint regions. We can therefore compare different GLNs based on the corresponding partitioning of \mathbb{R}^{d_z} . If a partitioning $\mathcal{P}(c_1)$ is a refinement of a second partitioning $\mathcal{P}(c_2)$, we intuitively expect that the GLN corresponding to the refined partitioning $\mathcal{P}(c_1)$ is more expressive. In the following, we will show this intuition for the two gated linear networks BF-GCN and BF-GraphConv.

Theorem 1. Consider two Gated Linear GNNs (either BF-GCN or BF-GraphConv) with gated geometric mixing based on two region assignment functions c_1 and c_2 . We assume that the partitioning $\mathcal{P}(c_1)$ is a refinement of the partitioning $\mathcal{P}(c_2)$. Then, the Gated Linear GNN based on c_1 is more expressive than the GNN based on c_2 .

Proof. We will only consider the BF-GCN model. For simplicity, we will also assume that $\mathbf{X} = \mathbf{x} \in \mathbb{R}^{n \times 1}$, i.e., that the dimension of the input variable is $s = 1$ and, thus, that we have real-valued weights $w_{(\mathbf{z})} \in \mathbb{R}$. Now, let $\mathcal{H}(c)$ denote the set of all possible functions

$$\mathcal{H}(c) = \left\{ \mathbf{y}_{(\mathbf{z})} = \sigma \left(\left(\tilde{\mathbf{D}}^{-\frac{1}{2}} (\mathbf{I} + \mathbf{A}) \tilde{\mathbf{D}}^{-\frac{1}{2}} \right)^l \mathbf{x} w_{(\mathbf{z})} \right) \mid w_{(\mathbf{z})} \in \mathbb{R} \right\}$$

generated by a BF-GCN network with a region assignment function c . Let $R \in \mathcal{P}(c_2)$ be one of the regions in the partitioning given by the function c_2 . As $\mathcal{P}(c_1)$ is a refinement of $\mathcal{P}(c_2)$ we can decompose R as $R = R_1 \cup \dots \cup R_r$ with $R_i \in \mathcal{P}(c_1)$ for $i \in \{1, \dots, r\}$. Now, if in $\mathcal{H}(c_1)$ we impose that $w_{\mathbf{z}_i} = w_{\mathbf{z}_j}$ if $\mathbf{z}_i \in R_i$ and $\mathbf{z}_j \in R_j$ for $i, j \in \{1, \dots, r\}$, then every $\mathbf{z} \in R$ obtains the same weight $w_{\mathbf{z}}$. Thus, we see that every function in $\mathcal{H}(c_2)$ can be formulated as a function in $\mathcal{H}(c_1)$. On the other hand, if we can choose $w_{\mathbf{z}_i} \neq w_{\mathbf{z}_j}$ freely for $\mathbf{z}_i \in R_i$ and $\mathbf{z}_j \in R_j$ we can enforce $\mathbf{y}_{(\mathbf{z}_1)} \neq \mathbf{y}_{(\mathbf{z}_2)}$ for $\mathbf{z}_i, \mathbf{z}_j \in R$ (excluding exceptional cases in which $\left(\tilde{\mathbf{D}}^{-\frac{1}{2}} (\mathbf{I} + \mathbf{A}) \tilde{\mathbf{D}}^{-\frac{1}{2}} \right)^l \mathbf{x}$ vanishes on domains linked to R_i and R_j). Therefore, the set $\mathcal{H}(c_1)$ is, in general, strictly larger than $\mathcal{H}(c_2)$. \square

From Theorem 1, it is easy to show the following corollary.

Corollary 1. *Every BF-GCN with more than one region assignment is more expressive than a Simplified Graph Convolutional Networks (SGC) [20]*

$$\mathbf{H} = \sigma((\tilde{\mathbf{D}}^{-\frac{1}{2}} \tilde{\mathbf{A}} \tilde{\mathbf{D}}^{-\frac{1}{2}})^l \mathbf{X} \mathbf{W}) \quad (12)$$

C. Multi-Resolution-GLN

A different definition of GNN is based on the idea of exploiting the power series of the diffusion operator in order to obtain a multi-scale representation of the graph features. The obtained representation is usually fed to an MLP that projects it into the output space. Our proposal is to substitute the MLP with a GLNs architecture.

Considering the explorative purpose of this work we decide to adopt the most general multi-scale representation definition proposed in [8]:

$$\mathbf{R}_{i,\mathcal{T}} = [\mathbf{X}, \mathcal{T}(\mathbf{A})\mathbf{X}, \mathcal{T}(\mathbf{A})^2\mathbf{X}, \dots, \mathcal{T}(\mathbf{A})^l\mathbf{X}],$$

where $\mathcal{T} : \bigcup_{j=1}^{\infty} (\mathbb{R}^{j \times j} \rightarrow \mathbb{R}^{j \times j})$ is a generic transformation of the adjacency matrix that preserves its shape, i.e. $\mathcal{T}(\mathbf{A}) \in \mathbb{R}^{n \times n}$. For instance, \mathcal{T} can be defined as the function returning the Laplacian matrix starting from the adjacency matrix. Then we can apply the GLN-based classifier. In particular, for each degree of diffusion operator (up to l) we consider c neurons, where c is the number of the classes considered in the classification problem. We recall that each neurons performs a binary classification problem. Therefore, similar to the message-passing GLN case, one-vs-rest approaches is exploited.

$$\begin{aligned} \mathbf{h}_i &= [\mathbf{h}_{j,(\mathbf{R}_{i,\mathcal{T}}, \mathbf{X})}]_{\forall j \in [0 \dots c-1]} \\ \mathbf{h} &= [\text{GeomMean}(\mathbf{h}_0[j], \dots, \mathbf{h}_l[j])]_{\forall j \in [0 \dots c-1]}. \end{aligned} \quad (13)$$

Note that \mathbf{h}_i is computed considering only the multi-resolution representation up to degree i of the diffusion operator. This allows us to obtain the same effect that the authors have in [8] where l Multi-Resolution convolutions of degree ranging from 0 up to l are concatenated. Finally, in order to combine the results of all the neurons of each class, we compute the geometric mean.

D. Context functions

The context function presented in Section III and exploited in (4) is based on random half-space gating. That definition is suited for online learning, where the training data distribution is not known beforehand. However, it is not data-driven and may result in the necessity of defining a high number of context regions to obtain a sufficiently non-linear model. Notice that the halfspace gating mechanism depends on some hyperparameters: in addition to the number of regions k , one has to choose the parameters of the distribution from which to sample the weights corresponding to each hyperplane (e.g. mean and variance assuming they are sampled from a normal distribution). Setting these hyperparameters may be challenging, since results can be strongly affected by their choice.

In this section, we propose an alternative approach that can be exploited in the batch learning scenario and that does not depend on any parameter but the number of regions to consider.

Dataset	#Classes	#Edges	#Train	#Val	#Test
Citeseer	6	9228	1995	666	666
Cora	7	10556	1624	542	542
Pubmed	3	88651	11829	3944	3944
WikiCS	10	216123	7021	2340	2340

TABLE I: Datasets statistics. The columns #Train, #Val, and #Test report the number of nodes in the training, validation and test sets, respectively.

In particular, we propose to define a partition of the space based on a set of prototypes [29]. Each point in the space is assigned to its closest prototype, obtaining a Voronoi tessellation. Note that half-space gating generates a division of the context space that can be represented as a planar straight-line graph (PSLG) instead. It is possible to show that any PSLG coincides with the Voronoi diagram of some set S of points (i.e. prototypes) [30]. Similarly to the half-space gating mechanism, the prototypes are not learned. However, instead of randomly generating them, we propose to sample at random among the training examples. This ensures that each prototype will lie on the input data manifold. Moreover, as mentioned before, this approach relieves us from many hyper-parameter choices.

Let $\mathbf{P}_j^{(i)} \in \mathbb{R}^{2^m \times m_z}$ be the matrix of prototypes. We can formally define the context vector $\mathbf{c}_{(z)}^{(i)} \in \{0, 1\}^{c_i}$ as:

$$\mathbf{c}_{j,(z)}^{(i)} = \text{one_hot}(\text{argmin}(2_dist(\mathbf{P}_j^{(i)}, \mathbf{z}))$$

where 2_dist computes the 2-norm distance between each row of $\mathbf{P}_j^{(i)}$ and \mathbf{z} is the context vector.

IV. EXPERIMENTAL RESULTS

We empirically validated the proposed backpropagation-free graph neural networks on four widely adopted datasets of node classification: Citeseer, Cora, Pubmed and WikiCS [31]. For what concerns the first three datasets, each of them is a single graph, where nodes represent documents and node features are sparse bag-of-words feature vectors. Specifically, in Citeseer, Cora, and Pubmed the task requires to classify the research topics of papers. Each node represents a scientific publication described by a 0/1-valued word vector indicating the absence/presence of the corresponding word from a dictionary. The WikiCS dataset is a graph where the nodes represent Wikipedia Computer Science articles and the edges represent the hyperlinks among them. Also in this case the node features are derived from the text, but differently from the other considered datasets, here they are calculated as the average of pre-trained *GloVe* word embeddings. For each dataset, we split the nodes in training (60%), validation (20%), and test (20%) sets. Relevant statistics of the dataset are reported in Table I.

We developed all the models involved in the comparison using PyTorch Geometric [32]. As baseline models, we considered the GCN and the GraphConv convolutions. For these models we exploit the implementation provided by PyTorch Geometric. In addition we also evaluate as a baseline a multi-resolution architecture (MRGNN) that exploits the same graph

augmented representation of the proposed backpropagation-free Multi-Resolution-GLN, followed by an MLP that projects the node representations to the output space. All the baseline models are trained in an end-to-end fashion using the backpropagation algorithm.

Furthermore, we do not limit our exploration of backpropagation-free models, and in particular the multi-layer ones (BF-GCN and BF-GraphConv), to the standard methodology that uses the input \mathbf{X} as context, but we also assess the effects of using each node’s inputs, i.e. the hidden representation computed at the previous layer label (\mathbf{H}) as a context for all gated neurons. For all the datasets, we solve the resulting optimization problems with the Adam algorithm (a variant of stochastic gradient descent with momentum and adaptive learning rate). We used early stopping (with the patience set to 15 epochs) and model checkpoint, monitoring the accuracy on the validation set. We set the maximum number of epochs to 250. All the baseline experiments involved softmax activation function applied to the last layer. The results were obtained by performing 5 runs for each model. For our experiments, we adopted a machine equipped with: 2 x Intel(R) Xeon(R) CPU E5-2630L v3, 192GB of RAM and a Nvidia Tesla V100. For more details please check the publicly available code¹.

A. Model selection

Before discussing the results of the proposed backpropagation-free graph neural networks, we would like to point out that for many reasons, the results reported in literature for different graph convolutions are not always comparable one to each other. A key aspect to consider is the procedure adopted to select the hyper-parameters (such as learning rate, regularization, network architecture, etc.). Many papers report, for each dataset, the best performance (on the test set) obtained after testing many hyper-parameter configurations. This procedure favours complex methods that depend on many hyper-parameters, since they have a larger set of trials to select from compared to simpler methods. However, the predictive performances computed in this way are not unbiased estimations of the true error, thus these results are not comparable to other model selection methods [33]. Another methodology used in many papers is to validate only a subset of the hyper parameters of the model, while many others are kept fixed. Unfortunately the values set for the hyper-parameters not considered in the validation process significantly impact the final results.

We recall that the contribution of this paper is not to present yet another graph neural network architecture performing slightly better than other alternatives. Instead, we want to show that it is possible to match the performance of different graph convolutional neural networks (and thus to perform an effective representation learning) even not relying on backpropagation.

For these reasons, we decided to focus on three widely adopted graph convolutions, and run all the experiments using

hyper-parameter	values
m	2, 4, 8, 16, 24, 32, 64
l	1, 2, 3, 4, 5, 6, 7, 8
k	1, 2, 4, 8, 16
learning rate	0.1, 0.2, 0.01, 0.001
weight decay	$0, 5^{-4}, 5^{-3}$
drop out	0, 0.2, 0.5
$\mathcal{T}(\mathbf{A})$	\mathbf{A}, \mathbf{L}

TABLE II: Sets of hyper-parameters values used for model selection via grid search.

the same fair model selection procedure, where we select all the hyper-parameters of each method on the validation set. The hyper-parameters of the model (number of hidden units, number of layers, learning rate, weight decay, and dropout only for backpropagation-based models) were selected by using a limited grid search, where the explored sets of values do change based on the considered dataset. We performed some preliminary tests in order to select the set of values taken into account for each hyper-parameter. In Table II, we report the sets of hyper-parameter values used for the grid search. In order to perform a fair comparison among the proposed models and the baselines, we use the same hyper parameter grid for all the models. To ensure a fair comparison between baseline and the proposed model that adopts a one-vs-rest approach, for the baselines we also consider the model where the number of hidden is the values reported in Table II multiplied by the number of classes of the considered task. As evaluation measure to perform model selection, we used the average accuracy computed on the validation set, while we report in Table III the average accuracy on the test set along with the hyper-parameters selected by the validation process.

B. Discussion

The results obtained by the backpropagation-free graph neural networks, namely BF-GCN, BF-GraphConv, and BF-MRGNN, are in general comparable and sometimes higher than the ones of the baselines trained with backpropagation. We recall that the backpropagation-free alternatives are significantly easier to train since they optimize convex optimization problems (one for each neuron, that can be parallelized), compared to the single but significantly more complex nonlinear problem optimized by the baselines. In Table III we report the results obtained validating all the hyper-parameters on the validation set. For each method and dataset we report the average accuracy and the standard deviation over 5 runs. Let us start discussing the GCN convolution mechanism, where the results of backpropagation-based (GCN) and the proposed backpropagation-free (BF-GCN) models with the two alternative context definitions, i.e. \mathbf{H} and \mathbf{X} are reported. Moreover, for BF-GCN (as for the all backpropagation-free models) we experimented with both the context function based on random half-space gating and our proposal of defining the partition of the space based on a set of prototypes (*context function* column in Table III). In this case, the backpropagation-

¹<https://github.com/lpasa/BF-GNN>

free models perform slightly better than the backpropagation counterpart in 2 datasets out of four. In all the cases, the methods are less than one standard deviation apart. We can conclude that, in this case, all the backpropagation-free models based on GCN are capable of learning a representation that is comparably expressive to the one of the backpropagation-based GCN. With this convolution mechanism there is not a clear advantage of using either \mathbf{H} or \mathbf{X} as contexts. These results show that the proposed backpropagation-free methods are pretty resilient and show consistent performance even with significantly different choices of the context space. Moreover, the results suggest that the prototype-based context function allows to reach slightly better performance in terms of accuracy compared to half-space gating.

Let us now consider the GraphConv convolution mechanism. In this case, backpropagation-free models show slight but consistent accuracy improvements on the Citeseer (up to 1.1%) and the Cora (up to 1.4%) datasets. On Pubmed, their results are at most 2.8% lower than the baseline GraphConv, while on WikiCS all the backpropagation-free variants improve over GraphConv, showing an accuracy improvement up to 5.5%. Notice that the differences in these last two cases are greater than one standard deviation. We can notice that with GraphConv, using \mathbf{H} as context tends to provide slightly higher performances compared to using \mathbf{X} . Analyzing the accuracy, no clear advantages can be noticed in using the prototypes-based context function instead of a half-space gating mechanism.

For what concerns BF-MRGNN the results obtained on Citeseer show a significant improvement with respect to its counterpart trained via backpropagation, i.e. MRGNN. Specifically, BF-MRGNN with prototypes as contexts achieved the overall best results in the Citeseer dataset, improving 3.4% over MRGNN. On the other datasets, the accuracy difference between MRGNN and its backpropagation-free versions is within the standard deviation interval. Note that in the BF-MRGNN, due to its definition (eq. (13)), only \mathbf{X} can be used as context.

We can conclude that the proposed backpropagation-free graph convolutions are competitive with their backpropagation-based counterparts, while inheriting all the advantages of backpropagation-free methods.

For what concerns the comparison between the two considered gating mechanisms (halfspace and prototype), we obtained no strong evidence in term of accuracy in favour of using one approach over the other. However, the prototype approach does have an advantage in reducing the number of hyperparameters. In fact, it is not straightforward to define the half-space gating hyperparameters, as random initialization of hyperplanes introduces a strong assumption on the data distribution. In our experiments, we decided to keep the same parameters used in [4] for the distribution from which the weights corresponding to each hyperplane are sampled, since modifying them even slightly seemed to impact the predictive performance. On the other hand, the proposed prototype-based context function allows to initialize the gating mechanism in a data-driven way, which turns out to be very simple since we can just

uniformly sample them from the training set. Finally, in terms of time complexity the BackPropagation-free models present a huge advantage with respect to the standard GNN. Indeed, the computation (both forward and backward step) of each neuron is independent from all the others, thus it is possible to perform the computation of each unit in parallel. Considering the Message-Passing GLNs, the layerwise construction of the model allows all neurons in the same layer to be computed in parallel. This constraint is overcome by the BL-MRGNN where all the nodes are completely independent, and thus it is possible to parallelize the computation on all neurons.

C. Computed representation

In Figure 2 we plot the representation of each node of the four considered datasets, and its representations obtained by exploiting a layer of BF-GCN using context \mathbf{H} and context \mathbf{X} and adopting the context functions based on random half-space gating and the prototypes-based mechanism. The configurations considered to generate these figures are selected through the validation phase (i.e., the ones reported in Table III.). All the figures are generated using UMAP [34]. The color of each node represents the class to which it belongs to. In Figure 2, the first plot of each row represents the manifold of the input space. In Cora and Citeseer datasets the positioning of the nodes is chaotic, while in Pubmed and WikiCs it is possible to note some areas where the nodes of a certain class are grouped.

The second and third plots of each row show the spatial representation obtained by the BF-GCN after the second hidden layer, using \mathbf{X} or \mathbf{H} as a context and the random half-space gating mechanism. Similarly, the fourth and the fifth plots of each row concern the spatial representation obtained by the BF-GCN adopting the prototypes-based mechanism context function, considering \mathbf{X} or \mathbf{H} as contexts, respectively. We do not consider the first layer, since in both cases the context is the same (the input embedding).

All the models that we examined have only two hidden layers, and thus the second one has as many hidden units as the number of classes. In Cora and Citeseer, the obtained embeddings are in different positions but the model tends to separate the classes similarly in both cases. Moreover, the type of context function does not significantly affect the representation. In both cases, using \mathbf{X} seems to be preferable since the division between the classes seems sharper than using \mathbf{H} . However, these differences are very slight, given the very slight difference in the predictive performances of the two models. Differently, in WikiCS and Pubmed, the use of the different context functions significantly influences the obtained representation. On Pubmed, the separation between the classes seems not sharp, but there are three clusters where most of the samples of each class are grouped. On these datasets, the prototype context function approach seems to enhance the spatial separation of classes. Differently, In WikiCS we see a more chaotic division. This behavior can be due to the high number of edges of the Graph of WikiCS. Indeed, having a stronger connection between nodes influence significantly the result of a convolution operator. Note that, using the

Model	context function	context	Citeseer	Cora	Pubmed	WikiCS
GCN	-	-	76.6±1.0	87.5±0.9	88.5±0.3	81.6±0.7
GraphConv	-	-	75.1±1.6	87.1±0.5	88.9±0.4	76.8±1.3
MRGNN	-	-	74.15±1.6	86.5±0.8	87.15±0.9	80.6±0.7
BF-GCN	Halfspace	H	76.0±1.4	87.7±0.5	88.1±0.3	79.8±0.8
			(2, 16, 2)	(1, 16, 2)	(1, 8, 2)	(2, 6, 2)
		X	76.3±1.8	88.0±0.5	88.1±0.4	80.7±0.7
		(2, 24, 2)	(2, 16, 2)	(2, 4, 2)	(2, 6, 8)	
	Proto	H	77.0±0.4	87.8±0.3	88.1±0.7	81.0±0.9
			(2, 16, 4)	(2, 8, 2)	(2, 32, 4)	(2, 4, 16)
X		76.9±2.0	88.0±0.4	88.0±0.6	80.4±0.8	
	(2, 8, 1)	(2, 32, 2)	(2, 8, 2)	(2, 6, 16)		
BF-GraphConv	Halfspace	H	76.5±0.9	88.0±1.0	86.8±0.9	80.6±0.4
			(1, 32, 2)	(2, 24, 2)	(2, 2, 2)	(4, 8, 2)
		X	76.3±1.8	88.0±1.1	86.5±0.4	80.5±0.5
		(2, 24, 2)	(1, 16, 2)	(2, 8, 4)	(2, 6, 2)	
	Proto	H	76.2±1.2	88.5±1.3	86.1±0.2	81.4±0.8
			(2, 8, 2)	(2, 8, 1)	(2, 32, 4)	(2, 32, 2)
X		74.9±0.6	87.8±1.1	86.5±0.3	82.3±0.7	
	(2, 16, 2)	(2, 8, 4)	(2, 16, 4)	(2, 8, 16)		
BF-MRGNN	Halfspace	-	76.3±1.5	87.6±0.6	87.6±0.5	78.4±0.7
			(3, 2)	(3, 2)	(2, 2)	(3, 4)
	Proto	-	77.5±1.0	87.1±0.9	87.9±0.4	79.4±0.6
	(2, 2)	(3, 2)	(2, 2)	(2, 2)	(2, 2)	

TABLE III: Accuracy comparison between the Back-propagation-free models and standard models. The model selection is preformed considering the results obtained on the validation set. Under each result we report the hyper-parameters selected via validation process: (l, m, k) for the BF-GCN and BF-GraphConv and (l, k) for the BF-MRGNN.

halfspace gating approach, the version using **H** as context, while apparently providing a worse representation compared to the counterpart using **X**, is actually the one performing better from the accuracy point of view. It is worth noticing that both the adopted manifold learning method (UMAP) and the hyper-parameters of the model can influence the obtained plots.

V. CONCLUSIONS AND FUTURE DIRECTIONS

In this paper, we explored a novel locally trainable graph convolutional operator dubbed backpropagation-free graph convolutional network. The proposed GC is inspired by Gated Linear Networks and extends them to be applied to graph-structured data. It relies on a representation space of graph nodes that is shattered into different subspaces according to the node context. Indeed, each neuron that composes the GC operator is defined as a *set* of weight vectors. A gating mechanism within each neuron selects the weight vector to use for processing the input, based on its context. This mechanism allows training each neuron independently, without using back-propagation, resulting in a set of convex problems to solve. We analyzed the strengths and the weaknesses of two variants of our approach exploiting both the common message-passing based convolutions (GCN and GraphConv) and a multi-resolution graph architecture (MRGNN). It is worth noticing that the proposed approach is not limited to the considered graph operators but it can be applied to any graph convolution operator. We empirically assessed the performances of BF-GCN, BF-GraphConv and BF-MRGNN on four commonly adopted node classification benchmarks, and verified their competitive

performances. Moreover, we analyzed the behavior of such models considering different options for shattering the context space associated to graph nodes. In our implementation, we use layer-wise training via Stochastic Gradient Descent (SDG), but many other methods can be exploited to solve the resulting convex problem., making it suitable for online and continual learning scenarios. Indeed, we plan to explore the application of the proposed backpropagation-free graph neural networks to continuous learning tasks in the near future.

REFERENCES

- [1] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *ICLR*, 2017, citation Key: Kipf2016a. [Online]. Available: <http://arxiv.org/abs/1609.02907>
- [2] C. Morris, M. Ritzert, M. Fey, W. L. Hamilton, J. E. Lenssen, G. Rattan, and M. Grohe, "Weisfeiler and leman go neural: Higher-order graph neural networks," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, Oct 2019, p. 4602–4609. [Online]. Available: <http://arxiv.org/abs/1810.02244>
- [3] J. Gilmer, S. S. Schoenholz, P. F. Riley, O. Vinyals, and G. E. Dahl, "Neural Message Passing for Quantum Chemistry," in *Proceedings of the 34th International Conference on Machine Learning*, apr 2017, pp. 1263—1272.
- [4] J. Veness, T. Lattimore, D. Budden, A. Bhoopchand, C. Mattern, A. Grabska-Barwinska, E. Sezener, J. Wang, P. Toth, S. Schmitt, and et al., "Gated linear networks," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 11, p. 10015–10023, May 2021.
- [5] J. C. Whittington and R. Bogacz, "Theories of error back-propagation in the brain," 2019.
- [6] D. G. Clark, L. F. Abbott, and S. Chung, "Credit assignment through broadcasting a global error vector," *arXiv:2106.04089 [cs, q-bio]*, Jun 2021, arXiv: 2106.04089. [Online]. Available: <http://arxiv.org/abs/2106.04089>
- [7] L. Chen, Z. Chen, and J. Bruna, "On graph neural networks versus graph-augmented mlps," *arXiv preprint arXiv:2010.15116*, 2020.

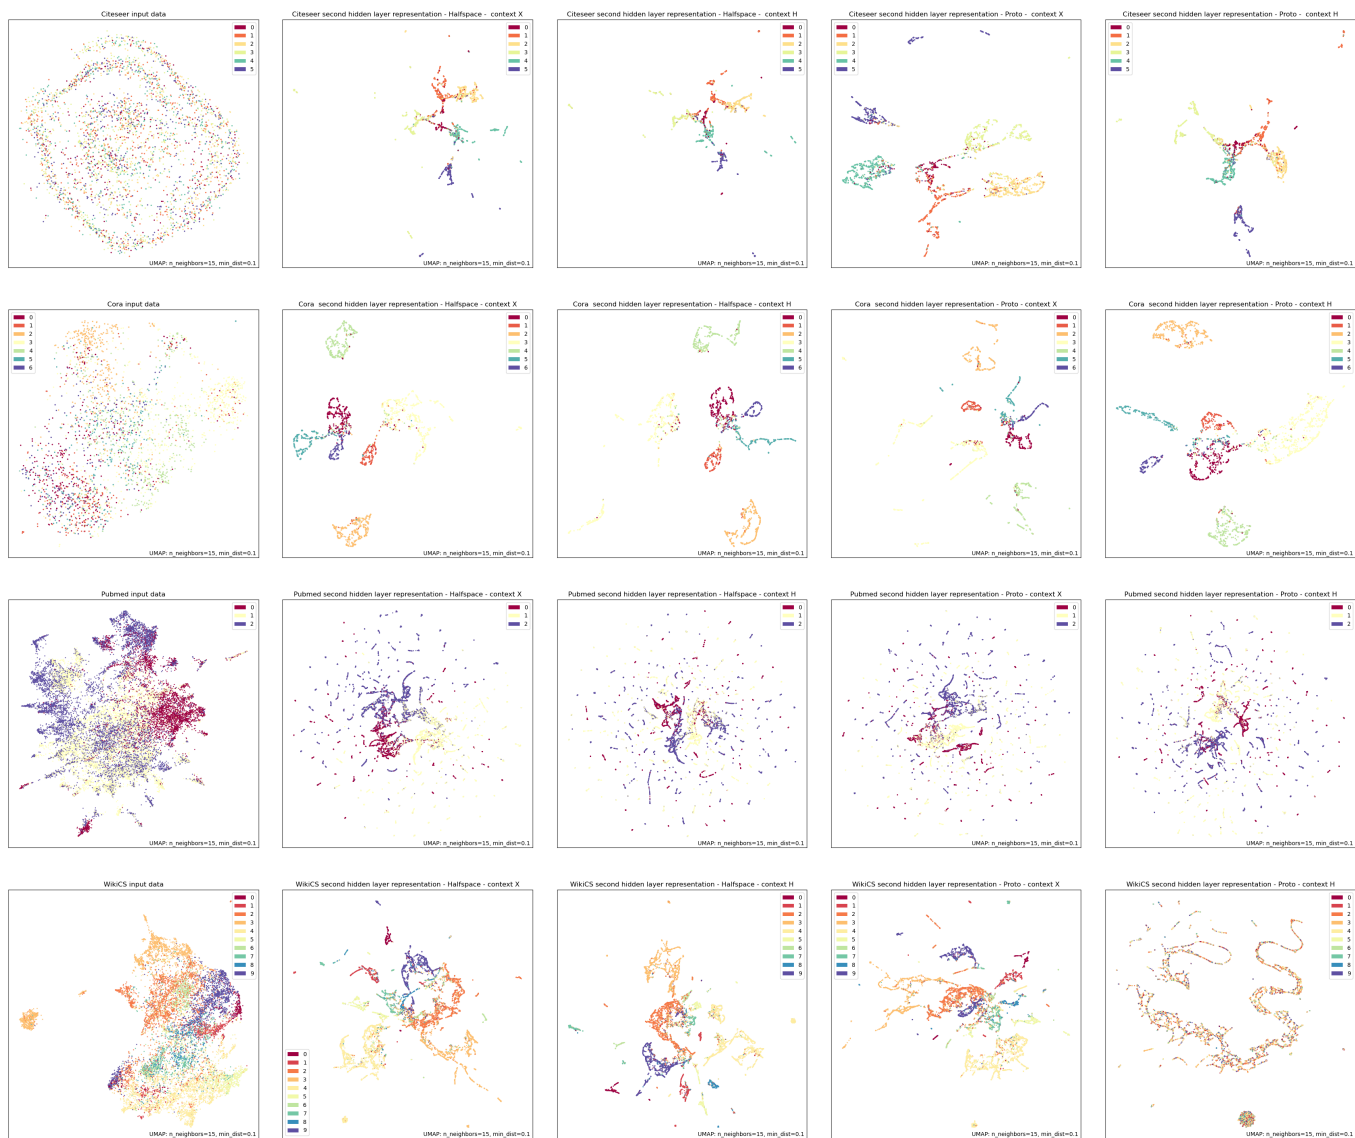


Fig. 2: UMAP 2D graphical representations of the embeddings of nodes by BF-GCN for all datasets. 1st column: original data; 2nd column: second hidden layer using context X and halfspace gating mechanism; 3rd column: second hidden layer using context X and halfspace gating mechanism ; 4nd column: second hidden layer using context X and prototypes-based gating mechanism; 3rd column: second hidden layer using context X and prototypes-based gating mechanism. The nodes are colored based on their target class.

[8] L. Pasa, N. Navarin, and A. Sperduti, “Polynomial-based graph convolutional neural networks for graph classification,” *Machine Learning*, 2021.

[9] A. Sperduti and A. Starita, “Supervised neural networks for the classification of structures,” *IEEE Trans. Neural Networks*, vol. 8, no. 3, pp. 714–735, 1997.

[10] T. Gärtner. “A survey of kernels for structured data,” *ACM SIGKDD Explorations Newsletter*, vol. 5, no. 1, p. 49, Jul 2003, citation Key: Gartner2003 publisher-place: New York, NY, USA.

[11] F. Scarselli, M. Gori, A. C. Ah Chung Tsoi, M. Hagenbuchner, and G. Monfardini, “The Graph Neural Network Model,” *IEEE Transactions on Neural Networks*, vol. 20, no. 1, pp. 61–80, 2009.

[12] A. Micheli, “Neural network for graphs: A contextual constructive approach,” *IEEE Transactions on Neural Networks*, vol. 20, no. 3, pp. 498–511, 2009.

[13] T. N. Kipf and M. Welling, “Semi-Supervised Classification with Graph Convolutional Networks,” in *ICLR*, 2017, pp. 1–14.

[14] M. Defferrard, X. Bresson, and P. Vandergheynst, “Convolutional neural networks on graphs with fast localized spectral filtering,” in *NIPS*, 2016, pp. 3844–3852.

[15] W. Hamilton, Z. Ying, and J. Leskovec, “Inductive representation learning on large graphs,” in *NIPS*, 2017, pp. 1024–1034.

[16] Y. Li, D. Tarlow, M. Brockschmidt, and R. Zemel, “Gated Graph Sequence Neural Networks,” in *ICLR*, 2016. [Online]. Available: <http://arxiv.org/abs/1511.05493>

[17] K. Xu, W. Hu, J. Leskovec, and S. Jegelka, “How Powerful are Graph Neural Networks?” in *International Conference on Learning Representations*, 2019.

[18] D. V. Tran, N. Navarin, and A. Sperduti, “On Filter Size in Graph Convolutional Networks,” in *IEEE SSCI*. Bengaluru, India: IEEE, nov 2018, pp. 1534–1541. [Online]. Available: <https://ieeexplore.ieee.org/document/8628758/>

- [19] Z. Xinyi and L. Chen, “Capsule graph neural network,” in *ICLR*, 2019.
- [20] F. Wu, T. Zhang, A. H. de Souza, C. Fifty, T. Yu, and K. Q. Weinberger, “Simplifying Graph Convolutional Networks,” *ICML*, feb 2019.
- [21] T. Chen, S. Bian, and Y. Sun, “Are powerful graph neural nets necessary? a dissection on graph classification,” *arXiv preprint arXiv:1905.04579*, 2019.
- [22] S. Luan, M. Zhao, X.-W. Chang, and D. Precup, “Break the ceiling: Stronger multi-scale deep graph convolutional networks,” in *Advances in neural information processing systems*, 2019, pp. 10 945–10 955.
- [23] E. Rossi, F. Frasca, B. Chamberlain, D. Eynard, M. Bronstein, and F. Monti, “Sign: Scalable inception graph neural networks,” *arXiv preprint arXiv:2004.11198*, 2020.
- [24] L. Pasa, N. Navarin, and A. Sperduti, “Multiresolution reservoir graph neural network,” *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–12, 2021.
- [25] J. Veness, T. Lattimore, D. Budden, A. Bhoopchand, C. Mattern, A. Grabska-Barwinska, E. Sezener, J. Wang, P. Toth, S. Schmitt, and M. Hutter, “Gated linear networks,” *arXiv*, 9 2019. [Online]. Available: <http://arxiv.org/abs/1910.01526>
- [26] J. Veness, T. Lattimore, A. Bhoopchand, A. Grabska-Barwinska, C. Mattern, and P. Toth, “Online learning with gated linear networks,” *arXiv*, pp. 1–40, 2017.
- [27] C. Mattern, “Linear and geometric mixtures - Analysis,” *Data Compression Conference Proceedings*, pp. 301–310, 2013.
- [28] N. Navarin, W. Erb, L. Pasa, and A. Sperduti, “Linear Graph Convolutional Networks,” in *European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning*, 2020.
- [29] M. Munari, L. Pasa, D. Zambon, C. Alippi, and N. Navarin, “Understanding catastrophic forgetting of gated linear networks in continual learning,” *2022 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, 2022.
- [30] G. Aloupis, H. Pérez-Rosés, G. Pineda-Villavicencio, P. Taslakian, and D. Trinchet, “Fitting voronoi diagrams to planar tessellations,” *arXiv:1308.5550 [cs]*, Aug 2013, arXiv: 1308.5550. [Online]. Available: <http://arxiv.org/abs/1308.5550>
- [31] P. Mernyei and C. Cangea, “Wiki-cs: A wikipedia-based benchmark for graph neural networks,” *arXiv preprint arXiv:2007.02901*, 2020.
- [32] M. Fey and J. E. Lenssen, “Fast graph representation learning with pytorch geometric,” in *ICLR 2019 (RLGM Workshop)*, Apr 2019.
- [33] L. Oneto, *Model selection and error estimation in a nutshell*. Springer, 2020.
- [34] A. Narayan, B. Berger, and H. Cho, “Density-preserving data visualization unveils dynamic patterns of single-cell transcriptomic variability,” *bioRxiv*, 2020. [Online]. Available: <https://www.biorxiv.org/content/early/2020/05/14/2020.05.12.077776>