Pseudoinverse graph convolutional networks

Fast filters tailored for large Eigengaps of dense graphs and hypergraphs

Dominik Alfke¹ · Martin Stoll¹

Received: 31 July 2020 / Accepted: 12 March 2021 / Published online: 15 April 2021
© The Author(s) 2021

Abstract

Graph Convolutional Networks (GCNs) have proven to be successful tools for semi-supervised classification on graph-based datasets. We propose a new GCN variant whose three-part filter space is targeted at dense graphs. Our examples include graphs generated from 3D point clouds with an increased focus on non-local information, as well as hypergraphs based on categorical data of real-world problems. These graphs differ from the common sparse benchmark graphs in terms of the spectral properties of their graph Laplacian. Most notably we observe large eigengaps, which are unfavorable for popular existing GCN architectures. Our method overcomes these issues by utilizing the pseudoinverse of the Laplacian. Another key ingredient is a low-rank approximation of the convolutional matrix, ensuring computational efficiency and increasing accuracy at the same time. We outline how the necessary eigeninformation can be computed efficiently in each applications and discuss the appropriate choice of the only metaparameter, the approximation rank. We finally showcase our method's performance regarding runtime and accuracy in various experiments with real-world datasets.

Keywords Machine learning · Graph neural networks · Semi-supervised learning · Signal processing on graphs

Responsible editor: Shuiwang Ji.

✉ Dominik Alfke
alfke@math.tu-chemnitz.de

Martin Stoll
stoll@math.tu-chemnitz.de

¹ Chemnitz University of Technologies, Chemnitz, Germany

1 Introduction

One of the central tasks in data science applications is extracting information from relational data encoded in graphs. The umbrella term Graph Neural Networks (GNNs) comprises neural models that seek to combine the theoretical understanding of structured data with the flexibility of machine learning. An outstandingly successful class of GNNs relies on spectral convolution of features along the graph edges (Bruna et al. 2014; Defferrard et al. 2016). These Graph Convolutional Networks (GCNs) have become particularly popular through their success in semi-supervised node classification tasks (Kipf and Welling 2017), where methods aim to benefit from both a small set of training data and clustering information extracted from a large amount of unlabeled data.

In this work, we consider *dense* graphs where any node is connected to most other nodes. This structure occurs in several applications where each node of the graph describes a real-world entity but the edges are constructed artificially based on specific node features. While it is possible to represent such examples using sparse networks this requires additional model assumptions that can be easily avoided using our approach. We here discuss in detail two relevant examples of such constructed, dense graphs. First, a Gaussian kernel function can be used to generate edge weights for fully connected graphs based on spatial node features, e.g., for three-dimensional point clouds as created by LiDAR scans (Nguyen and Le 2013). A localization parameter determines how fast the weights decay with the spatial distance, which can be understood to control the density of the graph. Up to now, only approaches with sparse k -nearest neighbor graphs have been proposed, but these are always associated with the loss of non-local information encoded within the large number of edges with lesser weights. We will show that increasing the density of the graph improves the prediction performance. This phenomenon is also well-known in other related fields like computer vision (Coll and Morel 2005; Tao et al. 2018) and image processing (Gilboa and Osher 2008). For a second example, *hypergraphs* provide a natural extension of graph learning that can be used easily for categorical data (Bretto 2013). For the purpose of many successful methods, these hypergraphs are equivalent to specific dense graphs with exploitable structure. Both these types of constructed graphs consist of edges that do not directly represent real-world connections but still encode valuable information about the dataset.

The density of constructed graphs has far-reaching consequences on the performance of GNNs. Graph learning has traditionally been targeted at purely data-inherent graphs where not only the nodes describe real-world entities, but also every single edge represents a real-world connection. Since the number of edges per node is often limited by real-world factors independent of the network size, the average node degree in these graphs is asymptotically constant, the total number of edges grows linearly with the network size, and the adjacency matrix is sparse. When moving to dense graphs, the increased computational cost and storage requirements might be expected to be a decisive issue. That is however often not the case, as the special structure of the constructed adjacency matrices can be exploited to speed up the relevant algorithms. The true problem is posed by the intrinsically different spectral properties of the dense graph Laplacian. It is well known in graph theory that the smallest eigenvalue λ_0

is always zero and the second smallest eigenvalue λ_1 gives a measure of how close the graph is to not being connected (Bauer and Jost 2013). We call λ_1 the graph's *eigengap* as the difference between the first and second eigenvalue. The eigenvectors corresponding to the first nonzero eigenvalues are known to contain clustering information. Most common GNNs reinforce this clustering information through their feature maps. In *spectral* approaches, this is achieved explicitly via convolution with a spectral filter designed specifically for that purpose. However, denser graphs empirically have much larger eigengaps and almost all eigenvalues are clustered close to 1 (Bauer and Jost 2013). For that reason, the informative eigenvectors are considerably harder to extract by existing filters as well as spatial feature maps. Common GNN architectures hence tend to underperform on dense graphs while our method embraces density and its spectral properties.

In this setting, the present work offers the following main contributions:

- We give motivation for spectral filters that combine a zero-impulse part and an inverse part to overcome the issues related with large eigengaps. In order to make our approach computationally efficient, we add a high-pass part and employ low-rank approximations to the pseudoinverse by computing a small number of eigenvalues of the graph Laplacian.
- We propose a Graph Convolutional Network architecture with a three-dimensional filter function space that represents learning our filter parameters in training. We discuss computational aspects such as asymptotic cost and parameter influence.
- We consider two examples for applications where beneficial dense graphs can be constructed. We discuss how the intrinsic structure of these graphs can be exploited to speed up our setup.
- We showcase the performance of our method in various experiments, comparing it to recent popular GNNs.

1.1 Related work

1.1.1 Graph neural networks

Neural networks have been used for learning on graphs for many years and the vast variety of recent methods and extensions can best be studied from the dedicated review articles (Bronstein et al. 2017; Wu et al. 2019; Zhang et al. 2019). As for methods with a particular connection to our work, we would like to single out Graph Diffusion Convolution (GDC; Klicpera et al. 2019) and ARMA filters (Bianchi et al. 2019). These methods also perform convolution with non-linear spectral filters, in GDC even non-locally. However, their filters are not specifically designed towards boosting clustering information in dense graphs.

1.1.2 GNN techniques for dense graphs

To the best of our knowledge, the challenges posed by dense graphs have rarely been addressed in the context of GNNs. In GDC (Klicpera et al. 2019), however, a graph is transformed into a dense graph whose adjacency matrix incorporates diffusion infor-

mation. This dense graph is then sparsified by a k -Nearest Neighbor approach, which is argued to both address runtime issues and improve prediction accuracy empirically. This sparsification process can equally be applied to general dense graphs, though it has to be noted that non-local information will be lost. Other techniques for increased efficiency like, e.g., node sampling (Hamilton et al. 2017; Chen et al. 2018) are explicitly targeted at large, sparse graphs.

1.1.3 Semi-supervised classification on hypergraphs

Hypergraph generalizations of the graph Laplacian operator have been introduced in competing ways. A definition based on the clique expansion graph (Zhou et al. 2006) has been used for various approaches based on energy minimization (Hein et al. 2013; Bosch et al. 2018) as well as Hypergraph Neural Networks (Feng et al. 2019), a natural generalization of GCNs. A similar work targets hypergraphs constructed from graphs (Bai et al. 2019). Yadati et al. (2019) quote density issues as motivation to avoid the clique expansion graph and instead propose HyperGCN using the Laplacian definition as a nonlinear diffusion operator introduced by Chan et al. (2018). None of these works address the Laplacian structure resulting from categorical data.

1.1.4 Inverse Laplacians in data science

Herbster et al. (2005) use the pseudoinverse of the graph Laplacian for online learning on graphs. Otherwise, mainly inverses of shifted Laplacians can be found in the literature. On multi-layer graphs, higher negative powers of shifted Laplacians can be used to form a Power Mean Laplacian (Mercado et al. 2018). In the context of GNNs, the approximated inverse of a shifted Laplacian is the heart of the diffusion operator for Personalized Page-Rank in GDC (Klicpera et al. 2019). In a broader sense, it is also an example of rational filtering, which is the basis of ARMA networks (Bianchi et al. 2019).

1.2 Problem setting and terminology

Throughout this paper, we assume that we are given an undirected weighted graph whose n nodes form the samples of the dataset, and each sample is associated with a d -dimensional feature vector. We assume that the graph is connected and non-bipartite, which is trivially fulfilled in the setting of dense graphs. On this data we aim to perform semi-supervised node classification. The goal is to assign one of m classes to each sample in such a way that nodes with a strong connection in the graph are likely to belong to the same class. For a small subset of samples called the training set, the true class is known a priori.

The graph can be described mathematically by its weighted adjacency matrix $A \in \mathbb{R}^{n \times n}$, where A_{ij} holds the weight of the edge between nodes i and j . This is set to 0 if the nodes are not connected. The degree matrix $D \in \mathbb{R}^{n \times n}$ is defined as the diagonal matrix holding the node degrees $D_{ii} = \sum_{j=1}^n A_{ij}$. Together they can be used to create

the symmetrically normalized graph Laplacian matrix

$$\mathcal{L}_{\text{sym}} = I - D^{-1/2} A D^{-1/2},$$

where I is the identity matrix. $\mathcal{L}_{\text{sym}}$ is known to encode many useful clustering properties of the graph (von Luxburg 2007).

Spectral graph theory now utilizes the eigenvalue decomposition of the graph Laplacian. Let (λ_i, u_i) be the eigenpairs of $\mathcal{L}_{\text{sym}}$ for $i = 0, \dots, n-1$ such that $\mathcal{L}_{\text{sym}} u_i = \lambda_i u_i$ and $\|u_i\|_2 = 1$. Since $\mathcal{L}_{\text{sym}}$ is symmetric, we have

$$\mathcal{L}_{\text{sym}} = U \Lambda U^T \quad \text{with} \quad \Lambda = \begin{bmatrix} \lambda_0 & & \\ & \ddots & \\ & & \lambda_{n-1} \end{bmatrix} \quad \text{and} \quad U = \begin{bmatrix} | & & | \\ u_0 & \cdots & u_{n-1} \\ | & & | \end{bmatrix}.$$

It is known that the smallest eigenvalue is always 0 with multiplicity one since there is only one connected component, and the largest eigenvalue is less than 2 since the graph is non-bipartite (von Luxburg 2007). We assume the eigenvalues to be sorted increasingly, $0 = \lambda_0 < \lambda_1 \leq \dots \leq \lambda_{n-1} < 2$. The corresponding eigenvector to $\lambda_0 = 0$ is u_0 with entries equal to the square roots of the node degrees,

$$u_0 = \left[D_{11}^{1/2}, \dots, D_{nn}^{1/2} \right]^T. \quad (1)$$

In case of multiple connected components, all eigenvectors for λ_0 can be computed in a similar fashion.

1.2.1 Convolution on graphs

The generalization of Fourier transforms and convolution from regular grids to irregular graphs is a central result of spectral graph theory (Shuman et al. 2013). Let $x \in \mathbb{R}^n$ be a spatial graph signal, i.e., x_i describes the magnitude of some quantity on node i . Then the function $\hat{x}(\lambda_j) = u_j^T x$, whose domain is the spectrum of the graph Laplacian, is the result of the graph Fourier transform. Its inverse is $x = \sum_{j=0}^{n-1} \hat{x}(\lambda_j) u_j$. In the spectral space, convolution with another spectral element φ is simply point-wise multiplication, $\hat{y}(\lambda_j) = \varphi(\lambda_j) \hat{x}(\lambda_j)$. Note that φ must be a real-valued function defined on the eigenvalues λ_j , but for simplicity, it is commonly defined on the whole interval $[0, 2]$. Put together, the spatial representation of the convolution of a spatial signal x with a spectral filter φ turns out to be $y = U \varphi(\Lambda) U^T x$, where $\varphi(\Lambda)$ denotes the diagonal matrix holding the function values of φ in the diagonal elements of Λ . The operator $\mathcal{K} = U \varphi(\Lambda) U^T$ is sometimes called the convolutional matrix associated with φ .

1.2.2 Absence of loops

Some popular GNN methods preprocess the graph by adding loops (edges with the same start and end node) with a uniform weight. For GCN, this has been interpreted as

a *re-normalization trick* (Kipf and Welling 2017). Besides that interpretation, the self loops empirically lead to slightly reduced eigengaps and hence better performance of the traditional methods that benefit from small eigengaps. This is because the normalized weight of non-loop edges is decreased, making the graph slightly sparser from a spectral point of view. Since this behaviour is not required in our method, we generally do not use loops.

2 Proposed architecture

2.1 Pseudoinverse spectral filter functions

Many popular machine learning methods involve repeated multiplication of a feature matrix with some kind of adjacency matrix. The most common form of GCN uses the normalized adjacency matrix, $\hat{A} = D^{-1/2}AD^{-1/2}$, as the convolutional matrix with the spectral filter $\varphi(\lambda) = 1 - \lambda$ (Kipf and Welling 2017). If all non-zero eigenvalues are close to 1, then $\varphi(\lambda)$ is close to zero for all eigenvalues except the first one. Hence $\hat{A}x = U\varphi(\Lambda)U^Tx$ will be dominated by the first eigenvector u_0 for most x , i.e., $\hat{A}x$ will almost be a multiple of u_0 and close to orthogonal to all other u_i ($i > 0$). This can be seen via the inner product

$$u_i^T \hat{A}x = \varphi(\lambda_i)u_i^T x \begin{cases} = u_0^T x & \text{if } i = 0, \\ \approx 0 & \text{else.} \end{cases}$$

Since all entries of u_0 have the same sign, this makes it hard for the network output to contain meaningful clustering information.

In order to overcome the issues of multiplying with the adjacency matrix, we would like to introduce the *pseudoinverse* of the graph Laplacian, denoted by $\mathcal{L}_{\text{sym}}^\dagger$. This has been used, e.g., for online learning (Herbster et al. 2005). At the same time, $\mathcal{L}_{\text{sym}}^\dagger$ is also the convolutional matrix of the spectral filter function (cf. Golub and Van Loan 1996)

$$\varphi^\dagger(\lambda) = \begin{cases} 0 & \text{if } \lambda = 0, \\ \frac{1}{\lambda} & \text{if } \lambda > 0. \end{cases}$$

This function is decreasing on the nonzero eigenvalues, so low-frequency signals are reinforced and high-frequency signals are damped. Because of $\varphi(0) = 0$, the first eigenvector u_0 is completely removed from the output of the convolution $\mathcal{L}_{\text{sym}}^\dagger x$, i.e., that output will always be orthogonal to u_0 . The problematic behaviour described above is hence completely avoided by the pseudoinverse filter. Note that the decay on the positive eigenvalues can also be achieved by different filters that might warrant investigation in the future.

However, this is a double-edged sword, since often some limited presence of u_0 may be beneficial. Hence we propose custom filters that allow for the combination of the pseudoinverse approach with added u_0 . This can be achieved by filter functions of

the form

$$\varphi_{\alpha,\beta}(\lambda) = \begin{cases} \alpha & \text{if } \lambda = 0, \\ \frac{\lambda_1\beta}{\lambda} & \text{if } \lambda > 0, \end{cases} \quad (2)$$

where the parameters α, β can either be assigned manually or learned. The eigengap λ_1 appears in the pseudoinverse part as a normalization factor. The corresponding convolutional matrix is

$$\mathcal{K}_{\alpha,\beta} = U\varphi_{\alpha,\beta}(\Lambda)U^T = \alpha u_0 u_0^T + \lambda_1 \beta \mathcal{L}_{\text{sym}}^\dagger.$$

2.2 Low-rank approach

In most scenarios we cannot compute the full eigenvalue decomposition of $\mathcal{L}_{\text{sym}}$. Instead, we compute only the $r + 1$ smallest eigenvalues and replace the term β/λ in (2) with a third parameter γ for all $\lambda > \lambda_r$. This means switching to a high-pass filter for higher frequencies, leading to the filter function

$$\varphi_{\alpha,\beta,\gamma,r}(\lambda) = \begin{cases} \alpha & \text{if } \lambda = 0, \\ \frac{\lambda_1\beta}{\lambda} & \text{if } 0 < \lambda \leq \lambda_r, \\ \lambda_1\gamma & \text{if } \lambda > \lambda_r \end{cases} \quad (3)$$

and the corresponding convolutional matrix

$$\mathcal{K}_{\alpha,\beta,\gamma,r} = (\alpha - \lambda_1\gamma)u_0 u_0^T + \lambda_1 U_r(\beta\Lambda_r^{-1} - \gamma I)U_r^T + \lambda_1\gamma I,$$

where

$$\Lambda_r = \begin{bmatrix} \lambda_1 & & \\ & \ddots & \\ & & \lambda_r \end{bmatrix} \in \mathbb{R}^{r \times r} \quad \text{and} \quad U_r = \begin{bmatrix} | & & | \\ u_1 & \cdots & u_r \\ | & & | \end{bmatrix} \in \mathbb{R}^{n \times r}$$

denote the matrices holding the second through $(r + 1)$ -st eigenvalues and corresponding eigenvectors in an ascending order. Note that $U_r\Lambda_r^{-1}U_r^T$ is the best rank- r approximation to $\mathcal{L}_{\text{sym}}^\dagger$, i.e.,

$$U_r\Lambda_r^{-1}U_r = \arg \min_{A \in \mathbb{R}^{n \times n}, \text{rank}(A) \leq r} \|\mathcal{L}_{\text{sym}}^\dagger - A\|_2,$$

which can be proven following the argumentation by, e.g., (Horn and Johnson 1985, Section 7.4.2).

On the one hand, this low-rank approach is necessary to avoid computing the full eigenvalue decomposition. On the other hand, it also has spectral benefits. In spectral graph theory for clustering applications, the smallest, but nonzero eigenvalues $\lambda_1, \lambda_2, \dots$ are called the *informative* eigenvalues. Their quantity depends on the graph.

The corresponding eigenvectors contain clustering information in the sense that if two nodes have a strong connection in the graph, the corresponding entries in the eigenvector are likely to be similar, especially in terms of their sign. If r is chosen appropriately, this low-rank filter can be argued to perform pseudoinverse convolution on the informative part of the spectrum, while the non-informative eigenvectors – especially noise – are damped uniformly. Hence, the low-rank approximation may very well have positive effects on the accuracy. Since there is no hard boundary between informative and non-informative eigenvalues, there is a wide range of ranks that can yield these benefits.

2.3 Pseudoinverse filters in graph convolutional networks

The general idea of convolutional feature maps in GNNs is that each column y_j of an output matrix $Y \in \mathbb{R}^{n \times N_1}$ is obtained by convolving each column x_i of an input matrix $X \in \mathbb{R}^{n \times N_0}$ with its own learned filter φ_{ij} and then summing up the convolution results. The learned filters are restricted to a given function space, which is arguably the decisive property of each convolutional GNN variant. Let that filter space be K -dimensional and spanned by the basis functions $\varphi^{(1)}, \dots, \varphi^{(K)}$. Then the coefficients of φ_{ij} in the given basis are the trained layer parameters, denoted by $W_{ij}^{(1)}, \dots, W_{ij}^{(K)}$. The individual filter functions are given as

$$\varphi_{ij} = \sum_{k=1}^K W_{ij}^{(k)} \varphi^{(k)}. \quad (4)$$

By organising the parameters in matrices $W^{(k)} \in \mathbb{R}^{N_0 \times N_1}$, this leads to the formula (Bruna et al. 2014; Defferrard et al. 2016)

$$y_j = \sum_{i=1}^{N_0} U \varphi_{ij}(\Lambda) U^T x_i = \sum_{i=1}^{N_0} \sum_{k=1}^K W_{ij}^{(k)} \underbrace{U \varphi^{(k)}(\Lambda) U^T}_{=\mathcal{K}^{(k)}} x_i,$$

and by putting together the full input feature matrix X with columns x_i and the output feature matrix Y with columns y_j , we obtain the feature map

$$Y = \sum_{k=1}^K \mathcal{K}^{(k)} X W^{(k)}. \quad (5)$$

We now aim to use a small filter space that contains the proposed pseudoinverse filters. One possibility is to choose the parameters in (3) manually and set the resulting $\varphi_{\alpha, \beta, \gamma, r}$ as the only basis function for a one-dimensional filter space. That requires a-priori identification of the parameter impact on desirable behaviour, which differs from dataset to dataset. For that reason, this approach is infeasible in practice.

Instead, we only fix r in (3) and have the other parameters learned in training, which means using the filter basis functions

$$\varphi^{(1)}(\lambda) = \begin{cases} 1 & \text{if } \lambda = 0, \\ 0 & \text{else,} \end{cases} \quad (\text{zero impulse part}) \quad (6)$$

$$\varphi^{(2)}(\lambda) = \begin{cases} \frac{\lambda_1}{\lambda} & \text{if } 0 < \lambda \leq \lambda_r, \\ 0 & \text{else,} \end{cases} \quad (\text{low-rank pseudoinverse part}) \quad (7)$$

$$\varphi^{(3)}(\lambda) = \begin{cases} \lambda_1 & \text{if } \lambda > \lambda_r, \\ 0 & \text{else.} \end{cases} \quad (\text{high-pass part}) \quad (8)$$

This means that the individual filter functions (4) are equal to (3) with $\alpha = W_{ij}^{(1)}$, $\beta = W_{ij}^{(2)}$, and $\gamma = W_{ij}^{(3)}$. The corresponding convolutional matrices can be set up as

$$\mathcal{K}^{(1)} = u_0 u_0^T, \quad \mathcal{K}^{(2)} = \lambda_1 U_r \Lambda_r^{-1} U_r^T, \quad \mathcal{K}^{(3)} = \lambda_1 (I - u_0 u_0^T - U_r U_r^T). \quad (9)$$

This feature map can now be embedded in a classical GNN architecture. We follow the traditional GCN setup for semi-supervised classification in that we use two layers each consisting of the convolutional feature map with an added bias and ReLU activation between the layers. The numbers of channels before and after each layer are given by the feature dimension d , the *hidden width* hyperparameter h , and the number of classes m . Let $X^{(0)} \in \mathbb{R}^{n \times d}$ be the input feature matrix. Extending the notation by an index for the layer, we get the propagation scheme

$$X^{(1)} = \sigma \left(\sum_{k=1}^3 \mathcal{K}^{(k)} X^{(0)} W^{(1,k)} + b^{(1)} \right), \quad X^{(2)} = \sum_{k=1}^3 \mathcal{K}^{(k)} X^{(1)} W^{(2,k)} + b^{(2)} \quad (10)$$

Here σ is the ReLU function applied element-wise and $W^{(1,k)} \in \mathbb{R}^{d \times h}$, $b^{(1)} \in \mathbb{R}^h$, $W^{(2,k)} \in \mathbb{R}^{h \times m}$, and $b^{(2)} \in \mathbb{R}^m$ are the trainable network parameters ($k = 1, \dots, K$). The addition of the biases $b^{(l)}$ to the matrices $\sum_{k=1}^K \mathcal{K}^{(k)} X^{(l-1)} W^{(l,k)}$ is understood row-wise. i.e., $b^{(l)}$ is added as a row vector to each row of the former matrix ($l = 1, 2$).

2.4 Computational aspects

2.4.1 Setup

In order to assemble the convolutional matrices, we only need to compute the $r + 1$ smallest eigenvalues of $\mathcal{L}_{\text{sym}}$. This can be done by efficiently computing the largest eigenvalues of the *signless Laplacian* via the state-of-the-art Krylov-Schur method (Stewart 2002). Since the eigenvector to $\lambda_0 = 0$ is known a priori via (1), we can use Wielandt deflation to remove the eigengap and significantly accelerate the method

(Saad 2011, Chapter 4.2). Put together, the system matrix of the eigenvalue computation is

$$2I - \mathcal{L}_{\text{sym}} - 2u_0u_0^T = I + D^{-1/2}AD^{-1/2} - 2u_0u_0^T.$$

If $\mu_0 \geq \dots \geq \mu_{r-1}$ are the r largest eigenvalue of that matrix, then the non-zero eigenvalues of $\mathcal{L}_{\text{sym}}$ can be recovered via $\lambda_i = 2 - \mu_{i-1}$ for $i = 1, \dots, r$. The corresponding eigenvectors are the same. This way, the asymptotic setup cost amounts to the number of Krylov iterations times the cost of one matrix-vector product, which is $\mathcal{O}(n^2)$ in the worst case but may be significantly less if we are able to exploit any special problem-dependent structure. The number of required iterations, on the other hand, depends on the desired rank r .

2.4.2 Asymptotic cost of layer operations

Similar to $\mathcal{L}_{\text{sym}}^\dagger$, the convolutional matrices $\mathcal{K}^{(k)}$ from (9) will in general not be sparse. However, we do not have to store and apply the full matrices, but rather exploit their low rank by keeping them in their factorized form and computing the feature map (5) via

$$\begin{aligned} Y = u_0 \left((u_0^T X) W^{(1)} - \lambda_1 u_0^T (X W^{(3)}) \right) \\ + \lambda_1 U_r \left(A_r^{-1} (U_r^T X) W^{(2)} - U_r^T (X W^{(3)}) \right) + X W^{(3)}. \end{aligned} \quad (11)$$

This way, the asymptotic cost of a single feature map is $\mathcal{O}(N_0 nr)$, where N_0 is the number of input features. Note that in general, multiplications with the dense adjacency matrix are already in $\mathcal{O}(n^2)$.

2.4.3 Choice of rank

As stated in Sect. 2.2, the target rank r should be chosen roughly as the number of informative eigenvalues. However, higher ranks may have additional benefits. Typically, the rank choice for low-rank approximations comes down to a trade-off between accuracy and runtime. An alternative is to view r as a meta-parameter to be determined via cross validation. We investigate its influence in Sect. 4.6. Note that it is very common for GNNs to depend on a few parameters that control some level of approximation.

3 Fast setup in special cases

Certain application settings allow us to exploit intrinsic structure of the adjacency matrix to speed up the required eigenvalue computations, as we discuss now.

3.1 Three-dimensional point clouds

One source of dense graphs are collections of points in three-dimensional space, called point clouds, which may be produced by LiDAR scans or other applications. The task of semi-supervised point classification occurs when a segmentation of such a scan is desired based on labels assigned manually to a few points. Other important tasks associated with this type of data are supervised point cloud segmentation and classification, which both revolve around transferring knowledge from fully labeled point clouds to new unlabeled point clouds.

One popular approach to working with this data, especially in robotics, is to turn the point cloud into a graph, e.g., by a k -Nearest Neighbor (KNN) setup (Nguyen and Le 2013; Golovinskiy and Funkhouser 2009). A promising alternative is to form a fully connected graph with Gaussian edge weights, leading to an adjacency matrix with entries

$$A_{ij} = \exp\left(\frac{-\|x_i - x_j\|^2}{\sigma^2}\right) \quad (12)$$

for all $i \neq j$, where the $x_i \in \mathbb{R}^3$ denote the point coordinates and $\sigma > 0$ is a localization parameter. This graph setup is often used for Spectral Clustering (Ng et al. 2002). A smaller value for σ leads to a sparser graph. However, it may be beneficial for the graph to incorporate more non-local information by means of a larger value for σ , which leads to a dense graph.

In this setting, the pseudoinverse assembly can be accelerated considerably. The smallest Laplacian eigenvalues and corresponding eigenvectors can be approximated accurately and efficiently by exploiting a fast summation scheme based on the Non-Equispaced Fast Fourier Transform (NFFT; see Alfke et al. 2018). This method has the remarkable property that it avoids assembling the n^2 adjacency entries altogether and that its computational effort for computing a small number of eigenvalues only scales linearly in n instead of the usual quadratic behaviour. By combining the NFFT algorithm with our low-rank approximation scheme, the computational effort is significantly reduced, especially for large n . Even though the amount of similarity information we look at is in $\mathcal{O}(n^2)$ and we do not discard any structure, the total complexity of our method with constant r scales only like $\mathcal{O}(n)$. For all details on the NFFT method, we refer to Alfke et al. (2018). Note that the NFFT-based fast multiplication for point cloud data Laplacians is applicable beyond our proposed filter whenever such a graph Laplacian matrix vector product is required.

3.2 Hypergraphs for categorical data

While the edges of traditional graphs connect exactly two nodes with each other, the *hyperedges* of a hypergraph may connect any number of nodes (Bretto 2013). Hypergraphs are most commonly described by their incidence matrix $H \in \{0, 1\}^{n \times |E|}$, where $|E|$ denotes the number of hyperedges and $H_{ie} = 1$ if and only if node i is a member of hyperedge e . Optional hyperedge weights can be given in a diagonal matrix

$W_E = \text{diag}(w_1, \dots, w_{|E|})$. In addition to the node degree matrix $D \in \mathbb{R}^{n \times n}$ with entries $D_{ii} = \sum_{e=1}^{|E|} H_{ie} w_e$, we also set up the hyperedge degree matrix $B \in \mathbb{R}^{|E| \times |E|}$ with entries $B_{ee} = \sum_{i=1}^n H_{ie}$.

3.2.1 The hypergraph Laplacian operator

The hypergraph Laplacian can be defined in multiple ways. We will use the linear definition introduced by Zhou et al. (2006), given as

$$\mathcal{L}_{\text{sym}} = I - D^{-1/2} H W_E B^{-1} H^T D^{-1/2}. \quad (13)$$

This is identical to the graph Laplacian of a classical graph with weighted adjacency matrix $H W_E B^{-1} H^T$, which is referred to as the *clique expansion* of the hypergraph. This graph contains a specific set of loops and is in most applications dense or even fully connected. As a consequence, naive computations with $\mathcal{L}_{\text{sym}}$ may become quite expensive. This problem also affects Hypergraph Neural Networks (Feng et al. 2019), which essentially apply the standard GCN architecture (including self loops) to the clique expansion graph.

3.2.2 Efficient techniques for the special case

One application of hypergraphs is categorical data where each sample is described by a few categorical attributes. We can simply construct one hyperedge for each possible value of an attribute, connecting all the samples which share that particular attribute value. This leads to the number of hyperedges being significantly smaller than the number of samples, $|E| \ll n$. For other automatically generated hypergraphs, there is precedence for the benefits of generating fewer, larger hyperedges as well (Purkait et al. 2017).

In this special case, the Laplacian definition directly exhibits a useful structure. The matrix subtracted from the identity in (13) has rank $|E|$, so $\mathcal{L}_{\text{sym}}$ is a linear combination of the identity and a low-rank matrix and can be written as

$$\mathcal{L}_{\text{sym}} = I - \tilde{H} \tilde{H}^T, \quad \tilde{H} = D^{-1/2} H W_E^{1/2} B^{-1/2} \in \mathbb{R}^{n \times |E|}. \quad (14)$$

Assume that $\tilde{H}$ has full rank $|E|$ and that its *thin* singular value decomposition is given by

$$\tilde{H} = \tilde{U} \tilde{\Sigma} \tilde{V}^T, \quad (15)$$

where $\tilde{U} \in \mathbb{R}^{n \times |E|}$ and $\tilde{V} \in \mathbb{R}^{|E| \times |E|}$ have orthogonal columns and $\tilde{\Sigma} \in \mathbb{R}^{|E| \times |E|}$ holds the singular values on its diagonal. Then $n - |E|$ eigenvalues of $\mathcal{L}_{\text{sym}}$ are 1 and the remaining eigenvalues are given by $\tilde{\Lambda} = I - \tilde{\Sigma}^2$. Consequently, the exact pseudoinverse of $\mathcal{L}_{\text{sym}}$ is the identity plus a matrix of rank $|E|$ and has the structure

$$\mathcal{L}_{\text{sym}}^\dagger = I + \tilde{U} (\tilde{\Lambda}^\dagger - I) \tilde{U}^T, \quad \tilde{\Lambda}^\dagger = \text{diag} \left(0, \frac{1}{\lambda_1}, \dots, \frac{1}{\lambda_{|E|-1}} \right). \quad (16)$$

For our low-rank approach, this means that we cannot choose $r > |E| - 1$ because that would require singling out a few eigenvectors to the eigenvalue 1, which are indistinguishable. However, computing the first $r + 1 \leq |E|$ eigenvalues of the hypergraph Laplacian becomes much cheaper, since we only need the singular value decomposition of $\tilde{H}$ or equivalently the eigenvalue decomposition of the $|E| \times |E|$ matrix $\tilde{H}^T \tilde{H}$. Thus the setup cost is significantly reduced. Furthermore, we can recreate the full-rank filter (2) within the low-rank setting of (3) by fixing $r = |E| - 1$ and $\gamma = \beta$.

4 Experimental results

4.1 Network architecture and training setup

Our code is available online,¹ implemented in Python using PyTorch and PyTorch Geometric (Fey and Lenssen 2019). The input features of each dataset are propagated through the network architecture (10) with the hidden width set to $h = 32$ in all experiments. In the first layer, the products $\mathcal{K}^{(k)} X^{(0)}$ are precomputed as proposed by Chen et al. (2018), and in the second layer we employ the efficient scheme from (11). Afterwards, the rows of the output $X^{(2)}$ are transformed via the softmax function to gain the predicted class probabilities for each sample. The parameters are trained using the average cross entropy loss of the predicted probabilities for the true class of the training samples. We use the Adam optimizer (Kingma and Ba 2015) with learning rate 0.01. During training, we use a dropout rate of 0.5 between the layers. For the weight matrices $W^{(l,k)}$, we use Glorot initialization (Glorot and Bengio 2010) and a weight decay factor of 0.0005, while for the bias vectors $b^{(l)}$, we use zero initialization and no weight decay.

For each run, we set a fixed seed for random number generation, build the model, run 500 training epochs, and finally evaluate the classification accuracy on the non-training samples. We generally perform 100 of these runs for each experimental setting and report averages and standard deviations. Only in a few slow baseline experiments did we reduce the number of runs to 10. All experiments are run on a laptop with an Intel Core i7 processor and an NVIDIA GeForce RTX 2060.

4.2 Baselines

We refer to our own method as PinvGCN in plots and tables, where we compare it against the following methods on graphs:

- GCN (Kipf and Welling 2017), using the implementation from PyTorch Geometric.
- GraphSAGE (Hamilton et al. 2017), using the implementation from PyTorch Geometric with mean aggregation. We do not use neighbor sampling since that is designed for node batches in large datasets and it did not improve the results in any of our tests.
- GDC (Klicpera et al. 2019), using the implementation from PyTorch Geometric with $\alpha = 0.05$ and top-64 sparsification. Since the datasets were too large

¹ <https://github.com/dominikalfke/PinvGCN>.

Table 1 Information on the Oakland point clouds

Name	Nodes	Classes	Label rate	Diameter	Eigengap λ_1	
					$\sigma = 10$	$\sigma = 100$
Oakland 1	36 932	5	1.35 %	112.1	0.084	0.929
Oakland 2	91 515	5	0.55 %	126.9	0.002	0.872

for exact matrix inversion, we needed to use the “inexact” version, which is not supported in the original code published with the paper.²

- ARMA (Bianchi et al. 2019), using the implementation from PyTorch Geometric with parameters $K = 3$ and $T = 2$.

On hypergraphs we additionally tested the following two methods:

- HGNN* (Feng et al. 2019), which we mark with an asterisk because we use our own implementation that exploits the structure of (14) in a similar way to our own method, significantly speeding up the runtime without changing the output.
- HyperGCN (Yadati et al. 2019), using the code published with the paper.³ Since the fast and non-fast variant yield the same accuracy in our experiments, we only employ “FastHyperGCN with mediators”.

4.3 Semi-supervised point classification in 3D point clouds

In order to illustrate our method’s superior performance on very dense graphs, we employ two 3D point clouds as described in Sect. 3.1. We use a part of a subset of the popular Oakland dataset as used by Munoz et al. (2009). We obtained the subset of the Oakland dataset from the project website.⁴ The dataset consists of multiple point clouds, two of which are used for training and validation in the original setting. Since their original usage is different from our own training splits, we only refer to the clouds as Oakland 1 (original training cloud) and Oakland 2 (original validation cloud). The remaining test clouds are unused. This data is usually used for 3D point segmentation, where the task is to transfer knowledge from one cloud to other clouds (Nguyen and Le 2013). Instead, we look at the two clouds independently and perform 5-class semi-supervised point classification on each of them by splitting each cloud into its own sets of training points and test points. We use 100 training points from each of the five classes. For the input feature matrix $X^{(0)}$, we use the original 3D coordinates. Dataset information is given in Table 1, where the diameter denotes the maximum Euclidean distance between two points in the cloud.

For eigenvalue computation, we use the `fastadj` Python implementation⁵ of the NFFT-based fast summation scheme (Alfke et al. 2018) with default parameters, combined with the Krylov-Schur algorithm with tolerance 10^{-3} . These settings are

² <https://github.com/klicperajo/gdc>.

³ <https://github.com/malllabiisc/HyperGCN>.

⁴ http://www.cs.cmu.edu/~vmr/datasets/oakland_3d/cvpr09/doc/.

⁵ <https://github.com/dominikalfke/FastAdjacency>.

chosen to give fast, rough approximations of the eigenvalues because we found that higher quality did not have any influence on the PinvGCN results.

We conduct experiments for $\sigma \in \{10, 100\}$ in the Gaussian function (12), where the larger value amounts to a stronger inclusion of non-local information. Our experimental results are shown in Table 2. Since the baselines are not designed for such dense graphs and would have exploding time and memory requirements on the fully connected graph, we only employ them on a k -NN subgraph of the k nearest neighbors, where k is either 10 or 100. Note that it would be possible to employ the GCN baseline on the fully connected graph by utilizing the same NFFT-based fast summation, but providing such an implementation was beyond the scope of our experiments.

To summarize the results, our method produces accurate predictions in reasonable time, comfortably outperforming all baselines. The fact that our accuracy is substantially better for $\sigma = 100$ shows that our method greatly profits from non-local information in non-sparse Laplacians. As we see in Table 1, adding non-local information via a larger value of σ results in increasing λ_1 , confirming the connection between spatial and spectral properties.

4.4 Hypergraph datasets from categorical attributes

We finally employ our method on three hypergraphs based on the *Mushroom* and *Covertype* datasets from the UCI machine learning repository (Dua and Graff 2017), which are common benchmarks for semi-supervised classification on hypergraphs (Hein et al. 2013; Yadati et al. 2019).

The Mushroom dataset⁶ contains 8124 samples from two classes, described by 21 categorical attributes (ignoring one with missing values). For each attribute, we create as many hyperedges as there are attribute values present, where each hyperedge connects all samples with a specific value.

The Covertype dataset⁷ contains 581012 samples from 7 classes, described by 10 continuous and 44 binary attributes. We follow the setup process from (Hein et al. 2013), dividing the value range of each continuous attribute into 10 equally sized bins and creating hyperedges that connect all samples with values in the same bin. For each binary attribute, we only create one hyperedge for those samples with a “true” value. All hyperedges have weight $w_e = 1$. Afterwards, we create two subhypergraphs by using only samples from classes 4 and 5, or 6 and 7. Because we remove all hyperedges with less than two nodes, the resulting hypergraphs have less than the original 144 hyperedges.

As in (Yadati et al. 2019), we use the hypergraph incidence matrix as input features for all three data sets, $X^{(0)} = H$. Table 3 gives the full hypergraph specifications as well as their smallest nonzero Laplacian eigenvalue.

Classification results on all three datasets are listed in Table 4. We only employ our Pseudoinverse GCN with the maximum rank $r = |E| - 1$, as will be supported in Sect. 4.6. For the graph-based methods, we use KNN sparsification of the clique expansion graph, so each node is connected to the $k = 10$ other nodes with highest

⁶ <https://archive.ics.uci.edu/ml/datasets/Mushroom>.

⁷ <https://archive.ics.uci.edu/ml/datasets/Covertype>.

Table 2 Average results on Oakland datasets

Method	σ	Oakland 1		Oakland 2		
		Time	Accuracy ($\pm$ SD)	Time	Accuracy ($\pm$ SD)	
10-NN	GCN	10	4.23 s	58.16 % ($\pm$ 4.39)	5.93 s	68.70 % ($\pm$ 18.28)
		100	4.24 s	58.42 % ($\pm$ 4.10)	5.91 s	68.64 % ($\pm$ 3.28)
	GraphSAGE	–	2.61 s	56.22 % ($\pm$ 5.57)	3.50 s	71.32 % ($\pm$ 4.26)
	GDC	–	112.08 s	48.61 % ($\pm$ 8.50)	1160.33 s	63.25 % ($\pm$ 25.15)
100-NN	ARMA	100	29.23 s	49.63 % ($\pm$ 2.55)	451.49 s	64.29 % ($\pm$ 4.41)
	GCN	10	17.99 s	35.44 % ($\pm$ 11.87)	44.04 s	66.47 % ($\pm$ 20.54)
		100	17.67 s	54.17 % ($\pm$ 6.49)	43.66 s	70.89 % ($\pm$ 3.27)
	GraphSAGE	–	6.63 s	59.34 % ($\pm$ 4.94)	13.76 s	70.55 % ($\pm$ 3.98)
	GDC	–	3446.95 s	37.84 % ($\pm$ 5.09)	Out of memory	
	ARMA	100	76.02 s	56.04 % ($\pm$ 4.51)		
PinvGCN, $r = 10$		10	6.70 s	84.64 % ($\pm$ 4.56)	14.86 s	74.26 % ($\pm$ 6.39)
		100	5.79 s	92.58 % ($\pm$ 1.42)	10.82 s	93.25 % ($\pm$ 1.02)
PinvGCN, $r = 30$		10	11.02 s	81.58 % ($\pm$ 4.64)	24.62 s	74.18 % ($\pm$ 6.61)
		100	11.08 s	93.13 % ($\pm$ 1.68)	24.03 s	94.91 % ($\pm$ 0.93)
PinvGCN, $r = 100$		10	29.80 s	81.70 % ($\pm$ 5.02)	69.54 s	73.29 % ($\pm$ 6.18)
		100	28.93 s	93.50 % ($\pm$ 1.83)	67.70 s	95.38 % ($\pm$ 0.90)

Table 3 Hypergraph datasets

Name	n	$ E $	Classes	λ_1
Mushroom	8 124	112	2	0.67
Covertime45	12 240	104	2	0.58
Covertime67	37 877	125	2	0.59

Fig. 1 Misclassification rate development over different split sizes for hypergraph datasets

shared hyperedge membership. On the Mushroom hypergraph, our results are in the same order of magnitude as the graph baselines, which is better than the other hypergraph methods, but the ARMA network gives superior results. The graph methods perform remarkably well considering the fact that this process of clique expansion sparsification has, to our knowledge, never been discussed in the hypergraph literature. On both Covertime datasets, however, our Pseudoinverse GCN yields the best accuracies.

In addition, we also emulated the full-rank pseudoinverse filter (2) via fixing $\beta = \gamma$ in (3) with $r = |E| - 1$. This method is named *PinvGCN without high-pass* in Table 4 and it produces comparable results. This shows that in the case that all informative eigenvalues can be computed, the high-pass part is no longer obligatory. However, manually removing this part also does not increase accuracy, so we recommend always using all three filter basis functions.

For further investigation, Fig. 1 presents the dependency of the misclassification rate (i.e., the complement of the accuracy, here plotted logarithmically) on the number of random training samples. The plots confirm the results of Table 4 in principle. A remarkable finding is that our method has a significantly better convergence rate on the Covertime45 hypergraph.

Table 4 Results for UCI categorical dataset with 10 training samples per class

Method		Mushroom Runtime	Accuracy ($\pm$ SD)
Sparsifiedgraph	GCN	2.93 s	91.46 % ($\pm$ 2.99)
	GraphSAGE	2.61 s	92.23 % ($\pm$ 3.71)
	GDC	31.98 s	92.24 % ($\pm$ 3.17)
	ARMA	8.16 s	92.36 % ($\pm$ 3.19)
HGNN*		0.61 s	83.38 % ($\pm$ 11.23)
FastHyperGCN		2.93 s	76.13 % ($\pm$ 16.38)
PinvGCN, $r = E - 1$		3.22 s	91.38 % ($\pm$ 3.92)
PinvGCN without high-pass		2.77 s	91.58 % ($\pm$ 3.33)
Method		Covertype45 Runtime	Accuracy ($\pm$ SD)
Sparsified graph	GCN	1.49 s	96.68 % ($\pm$ 2.43)
	GraphSAGE	1.30 s	97.80 % ($\pm$ 1.76)
	GDC	138.49 s	98.16 % ($\pm$ 1.76)
	ARMA	9.79 s	96.96 % ($\pm$ 3.39)
HGNN*		0.62 s	94.83 % ($\pm$ 15.75)
FastHyperGCN		3.51 s	89.81 % ($\pm$ 8.72)
PinvGCN, $r = E - 1$		3.33 s	99.58 % ($\pm$ 0.71)
PinvGCN without high-pass		2.71 s	99.56 % ($\pm$ 0.81)
Method		Covertype67 Runtime	Accuracy ($\pm$ SD)
Sparsified graph	GCN	2.80 s	90.01 % ($\pm$ 2.63)
	GraphSAGE	3.14 s	92.91 % ($\pm$ 2.33)
	GDC	417.18 s	94.20 % ($\pm$ 2.72)
	ARMA	30.39 s	93.53 % ($\pm$ 2.56)
HGNN*		1.09 s	91.07 % ($\pm$ 10.78)
FastHyperGCN		9.00 s	80.57 % ($\pm$ 13.30)
PinvGCN, $r = E - 1$		3.55 s	96.33 % ($\pm$ 1.41)
PinvGCN without high-pass		3.02 s	96.27 % ($\pm$ 1.43)

Table 5 Results for sparse citation networks

Method	Citeseer	Cora	Pubmed
Number of nodes n	3327	2708	19717
Number of connected components k	438	78	1
Eigengap λ_k	0.0013	0.0036	0.0095
GCN (as reported in Kipf and Welling 2017)	70.3 %	81.5 %	79.0 %
PinvGCN, rank 50	61.14 %	70.79 %	70.76 %
PinvGCN, rank 200	60.90 %	71.15 %	71.53 %
PinvGCN, rank 500	61.62 %	71.41 %	71.58 %
PinvGCN, rank 500, best regularization	68.33 %	82.34 %	78.26 %

4.5 Limitations: sparse graph datasets

For comparison and transparency, we also apply our method to the standard citation networks Cora, Citeseer, and Pubmed. These graph datasets are typical benchmarks for GNNs. However, they all are examples of sparse graphs that are exactly the opposite of what our method is designed to handle. The main challenge of these examples for our method are the very small eigengaps. Moreover, the eigenvalue zero often has multiplicity larger than one, which requires putting more effort into the eigenvalue computation. Network statistics and full results are given in Table 5. Unsurprisingly, our method by itself produces poor results on these datasets and it does not come close to the performance of a simple GCN (Kipf and Welling 2017).

However, since these citation networks come with categorical node features, it is possible to use similar techniques as for the categorical hypergraphs in order to obtain a more dense structure. For a first step, we transform the graph into a hypergraph with one hyperedge per node, following Feng et al. (2019). We then “regularize” the sparse graphs by adding one hyperedge for each column of the node feature matrix X , connecting all nodes with a nonzero feature value. Those feature-based hyperedges are weighted with a *regularization factor* $R > 0$. Similarly to the hypergraphs from Section 4.4, the newly formed hypergraph Laplacian is considerably more dense and the eigengap grows monotonously with R . For each dataset, the results for various values of R are visualized in Fig. 2 and the best obtained accuracy is listed in Table 5. The regularization helps our method catch up with state of the art, but in practice, this approach presents challenges in the high sensitivity and lack of a heuristic regarding R . It appears that for these citation networks, the sparse representation contains more useful clustering information than the node features. Our method is better suited to intrinsically dense data instead of relying on densification.

4.6 Rank dependency

Since the target rank r is the only metaparameter of our method, it is important to discuss its impact. As is typical for low-rank approximations, the choice of rank is

Fig. 2 Results for graph datasets transformed into centroid hypergraphs with categorical regularization

Fig. 3 Misclassification rate and runtime development over different ranks r

Table 6 Average absolute entries in weight matrices for different filter basis functions

Dataset	Rank	Part 1 (zero-impulse)	Part 2 (pseudoinverse)	Part 3 (high-pass)
Oakland 1, $\sigma = 100$	100	0.085	0.488	0.193
Oakland 2, $\sigma = 100$		0.247	0.464	0.191
Mushroom		0.075	0.137	0.046
Coverttype45	$ E - 1$	0.053	0.112	0.013
Coverttype67		0.034	0.114	0.009

subject to a trade-off between runtime and accuracy. Figure 3 depicts the development of the misclassification rate (plotted logarithmically) and runtime.

We note that the accuracy depends nontrivially on the rank. For almost all datasets, the best performance is achieved with the highest tested rank, but the dependency is not monotonous. Especially the hypergraphs show behavior where increasing the rank over a short range slightly deteriorates the accuracy. The point clouds are closer to monotony in this regard. However, the method appears to be very robust with respect to the choice of r , as long as it is not too small ($r < 20$).

The runtime development, on the other hand, depends on the exploited structure of the adjacency matrix. The total effort is governed by the setup cost, while the cost of layer operations does not seem to have a big impact. For point clouds, the number of computed eigenvalues influences the number of Krylov-Schur iterations, which leads to an almost linear dependency on r . For hypergraphs, the SVD of the normalized adjacency matrix is cheap enough to avoid any increase in runtime, leading to almost constant times.

For the best performance, we recommend choosing the maximum $r = |E| - 1$ on hypergraphs without worrying about this parameter. For point clouds, we encourage choosing r as large as possible while keeping a practically feasible computational cost.

4.7 Analysis of learned weight entries

As described in Sect. 2.3, our neural network learns the parameters of the filter functions (3) in training. The individual learned filters (4) may put varying focus on each of the three parts (6)–(8) as determined by the magnitudes of $W_{ij}^{(1,l)}$, $W_{ij}^{(2,l)}$, $W_{ij}^{(3,l)}$. By forming the averages of the absolute weight entries, we can quantify the importance of each filter part for the trained network. To account for the different weight matrix sizes in each layer, we use the formula

$$\mu_k = \frac{1}{2} \left(\frac{1}{dh} \sum_{i=1}^d \sum_{j=1}^h |W_{ij}^{(k,1)}| + \frac{1}{hm} \sum_{i=1}^h \sum_{j=1}^m |W_{ij}^{(k,2)}| \right) \quad (k = 1, 2, 3). \quad (17)$$

These numbers are furthermore averaged over all 100 runs. Table 6 lists these values for multiple PinvGCN instances. We clearly see that the pseudoinverse part consis-

tently is the most important filter basis function, which supports the notion that these eigenvectors carry the most clustering information. The weights of the other two parts are smaller, but not by a large margin, which supports the intuition that the other eigenvectors are still beneficial for the classification result. At the same time, we observe that the entry ratios differ quite a bit between datasets, which implies that is indeed hard to manually choose parameters for one suitable filter function as in (3) a priori.

5 Conclusion

We here presented Pseudoinverse GCN, a new type of Graph Convolutional Network designed for dense graphs and hypergraphs. The feature maps are based on a novel three-part filter space motivated by a low-rank approximation of the Laplacian pseudoinverse. The method yielded strong experimental results in a setting where popular GNNs struggle. A further advantage of our method is the robustness with respect to its only parameter. Future work might include extensions towards supervised 3D point cloud segmentation.

Funding Open Access funding enabled and organized by Projekt DEAL. D. Alfke gratefully acknowledges partial funding by the Sächsische Aufbaubank – Förderbank – (SAB) 100378180.

Declarations

Availability of data and material All used datasets are publicly available.

Code availability The authors have made their code available online at <https://github.com/dominikalfke/PinvGCN>.

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alfke D, Potts D, Stoll M, Volkmer T (2018) NFFT meets Krylov methods: fast matrix-vector products for the graph Laplacian of fully connected networks. *Front Appl Math Stat*. <https://doi.org/10.3389/fams.2018.00061>
- Bai S, Zhang F, Torr PH (2019) Hypergraph convolution and hypergraph attention [arxiv:1901.08150](https://arxiv.org/abs/1901.08150)
- Bauer F, Jost J (2013) Bipartite and neighborhood graphs and the spectrum of the normalized graph Laplacian. *Commun Anal Geom* 21(4):787–845
- Bianchi FM, Grattarola D, Alippi C, Livi L (2019) Graph neural networks with convolutional ARMA filters [arXiv:1901.01343](https://arxiv.org/abs/1901.01343)

- Bosch J, Klamt S, Stoll M (2018) Generalizing diffuse interface methods on graphs: nonsmooth potentials and hypergraphs. *SIAM J Appl Math* 78(3):1350–1377
- Bretto A (2013) *Hypergraph theory*, 1st edn. Eng, Springer, Math
- Bronstein M, Bruna J, Lecun Y, Szlam A, Vandergheynst P (2017) Geometric deep learning: going beyond Euclidean data. *IEEE Signal Process Mag* 34(4):18–42. <https://doi.org/10.1109/MSP.2017.2693418>
- Bruna J, Zaremba W, Szlam A, LeCun Y (2014) Spectral networks and locally connected networks on graphs. In: *Proc Int Conf Learn Represent, ICLR* 14
- Chan THH, Louis A, Tang ZG, Zhang C (2018) Spectral properties of hypergraph Laplacian and approximation algorithms. *J ACM* 65(3):15:1–15:48
- Chen JJ, Ma T, Xiao C (2018) FastGCN: Fast learning with graph convolutional networks via importance sampling. In: *Proc Int Conf Learn Represent, ICLR* 14
- Coll B, Morel JM (2005) A non-local algorithm for image denoising. *Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit, CVPR* 05:60–65
- Defferrard M, Bresson X, Vandergheynst P (2016) Convolutional neural networks on graphs with fast localized spectral filtering. In: *Adv Neural Inf Process Syst* 29, NIPS 16
- Dua D, Graff C (2017) UCI machine learning repository. <http://archive.ics.uci.edu/ml>
- Feng Y, You H, Zhang Z, Ji R, Gao Y (2019) Hypergraph neural networks. In: *33rd AAAI Conf Artif Intell, AAAI* 19
- Fey M, Lenssen JE (2019) Fast graph representation learning with PyTorch Geometric. In: *ICLR workshop on representation learning on graphs and manifolds, ICLR* 19
- Gilboa G, Osher S (2008) Nonlocal operators with applications to image processing. *Multiscale Model Simul* 7(3):1005–1028
- Glorot X, Bengio Y (2010) Understanding the difficulty of training deep feedforward neural networks. In: *Proc 13th Int Conf Artif Intell Stat, AISTATS* 10, vol 9
- Golovinskiy A, Funkhouser T (2009) Min-cut based segmentation of point clouds. In: *Proc IEEE Int Conf Comput Vis, ICCV* 09, <https://doi.org/10.1109/ICCVW.2009.5457721>
- Golub GH, Van Loan CF (1996) *Matrix computations*, 3rd edn. The Johns Hopkins University Press
- Hamilton W, Ying Z, Leskovec J (2017) Inductive representation learning on large graphs. *Adv Neural Inf Process Sys, NIPS* 17:1024–1034
- Hein M, Setzer S, Jost L, Rangapuram SS (2013) The total variation on hypergraphs - learning on hypergraphs revisited. *Adv Neural Inf Process Sys, NIPS* 13:2427–2435
- Herbster M, Pontil M, Wainer L (2005) Online learning over graphs. *Proc Int Conf Mach Learn, New York, US, ICML* 05:305–312
- Horn RA, Johnson CR (1985) *Matrix analysis*. Cambridge University Press
- Kingma D, Ba JL (2015) Adam: a method for stochastic optimization. In: *Proc Int Conf Learn Represent, ICLR* 15
- Kipf TN, Welling M (2017) Semi-supervised classification with graph convolutional networks. In: *Proc Int Conf Learn Represent, ICLR* 17
- Klicpera J, Weissenberger S, Günnemann S (2019) Diffusion improves graph learning. *Adv Neural Inf Process Sys, NIPS* 19:13333–13345
- Mercado P, Gautier A, Tudisco F, Hein M (2018) The power mean laplacian for multilayer graph clustering. In: *Proc Int Conf Artif Intell Stat, PMLR, AISTATS* 18, vol 84, pp 1828–1838
- Munoz D, Bagnell JA, Vandapel N, Martial H (2009) Contextual classification with functional max-margin Markov networks. In: *Proc IEEE Comput Soc Conf Comput Vis Pattern Recognit, CVPR* 09, <https://doi.org/10.1109/CVPR.2009.5206590>
- Ng AY, Jordan MI, Weiss Y (2002) On spectral clustering: analysis and an algorithm. *Adv Neural Inf Process Sys, MIT Press, NIPS* 01:849–856
- Nguyen A, Le B (2013) 3D point cloud segmentation: a survey. *Proc IEEE Conf Robot Autom Mechatron, RAM* 13:225–230
- Purkait P, Chin TJ, Sadri A, Suter D (2017) Clustering with hypergraphs: the case for large hyperedges. *IEEE Trans Pattern Anal Mach Intell* 39:1697–1711
- Saad Y (2011) Numerical methods for large eigenvalue problems. *SIAM*. <https://doi.org/10.1137/1.9781611970739>
- Shuman D, Narang S, Frossard P, Ortega A, Vandergheynst P (2013) The emerging field of signal processing on graphs: extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Process Mag* 30:83–98

- Stewart GW (2002) A Krylov-Schur algorithm for large eigenproblems. *SIAM J Matrix Anal Appl* 23(3):601–614
- Tao Y, Sun Q, Du Q, Liu W (2018) Nonlocal neural networks, nonlocal diffusion and nonlocal modeling. *Adv Neural Inf Process Sys*, NIPS 18:496–506
- von Luxburg U (2007) A tutorial on spectral clustering. *Stat Comp* 17(4):395–416
- Wu Z, Pan S, Chen F, Long G, Zhang C, Yu PS (2019) A comprehensive survey on graph neural networks [arxiv:1901.00596](https://arxiv.org/abs/1901.00596)
- Yadati N, Nimishakavi M, Yadav P, Nitin V, Louis A, Talukdar P (2019) Hypergcnn: a new method for training graph convolutional networks on hypergraphs. *Adv Neural Inf Process Sys*, NIPS 19:1509–1520
- Zhang S, Tong H, Xu J, Maciejewski R (2019) Graph convolutional networks: a comprehensive review. *Comput Soc* 6(11) <https://doi.org/10.1186/s40649-019-0069-y>
- Zhou D, Huang J, Schölkopf B (2006) Learning with hypergraphs: clustering, classification, and embedding. In: *Adv Neural Inf Process Syst*, NIPS 06

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.