

# What Can We Learn From MIMO Graph Convolutions?

Andreas Roth<sup>1</sup>, Thomas Liebig<sup>1,2</sup>

<sup>1</sup> Artificial Intelligence, TU Dortmund University

<sup>2</sup> Lamarr Institute for ML and AI, TU Dortmund University

{andreas.roth, thomas.liebig}@tu-dortmund.de

## Abstract

Most graph neural networks (GNNs) utilize approximations of the general graph convolution derived in the graph Fourier domain. While GNNs are typically applied in the multi-input multi-output (MIMO) case, the approximations are performed in the single-input single-output (SISO) case. In this work, we first derive the MIMO graph convolution through the convolution theorem and approximate it directly in the MIMO case. We find the key MIMO-specific property of the graph convolution to be operating on multiple computational graphs, or equivalently, applying distinct feature transformations for each pair of nodes. As a localized approximation, we introduce localized MIMO graph convolutions (LMGCs), which generalize many linear message-passing neural networks. For almost every choice of edge weights, we prove that LMGs with a single computational graph are injective on multisets, and the resulting representations are linearly independent when more than one computational graph is used. Our experimental results confirm that an LMG can combine the benefits of various methods.

## 1 Introduction

Graph neural networks have emerged as an effective method for many challenging applications involving graph-structured data, e.g., molecular prediction [Hu *et al.*, 2021]. These utilize convolutional operations typically derived from the general graph convolution obtained in the Fourier domain, as given by the convolution theorem [Hammond *et al.*, 2011; Bruna *et al.*, 2014]. Initially, approximations of the general graph convolution were based on polynomials, e.g., Chebyshev polynomials [Hammond *et al.*, 2011]. The graph convolutional network (GCN) [Kipf and Welling, 2017] approximates these polynomials as a first-order localization. Many other message-passing approaches are derived from the GCN [Velickovic *et al.*, 2018; Bo *et al.*, 2021]. However, these approximations are based on the single-output single-input (SISO) case, where the input and output contain a single feature for each node. GNNs are typically applied in the multi-input multi-output (MIMO) case, where each

node has multiple feature channels assigned, and the output also contains multiple features. Extending from the SISO to the MIMO case is achieved by applying these methods for each input and output channel combination and learning distinct parameters [Bruna *et al.*, 2014; Defferrard *et al.*, 2016; Kipf and Welling, 2017].

Instead of first approximating the graph convolution in the SISO and then extending to the MIMO, we propose directly performing the approximation in the MIMO case to benefit from MIMO-specific properties. We first derive the general graph convolution in the MIMO case through the convolution theorem and the graph Fourier transform. We find the key property that allows the MIMO-GC to represent arbitrary transformations to be operating on multiple computational graphs or, equivalently, applying distinct linear feature transformations between each pair of nodes. This form allows a direct approximation in the MIMO case by localizing the aggregation step. The resulting localized MIMO-GC (LMGC) presents a general framework for linear message-passing neural networks (MPNN) that inherits the beneficial properties for multi-channel learning. While we show that the LMG can represent most MPNNs, the LMG cannot represent the graph isomorphism network (GIN) [Xu *et al.*, 2019] due to its non-linear feature transformation. However, we show that LMGs are injective on multisets for almost every choice of edge weights even for a single computational graph. When further utilizing multiple computational graphs as motivated by the MIMO-GC, we prove that representations are linearly independent for almost every choice of edge weights. We summarize our main contributions as follows:

- Based on the convolution theorem, we derive the MIMO graph convolution (MIMO-GC) for node representations with multiple feature channels. A key property of MIMO-GCs is to operate on multiple computational graphs, or equivalently, to apply distinct linear feature transformations for each pair of nodes (Section 3).
- We introduce the framework of localized MIMO-GCs (LMGCs) by localizing the aggregation step of the MIMO-GC. It merges the key idea of operating on multiple computational graphs with the efficient message-passing scheme (Section 4).
- We prove that LMGs are injective on multisets for a single computational graph and produce linearly inde-

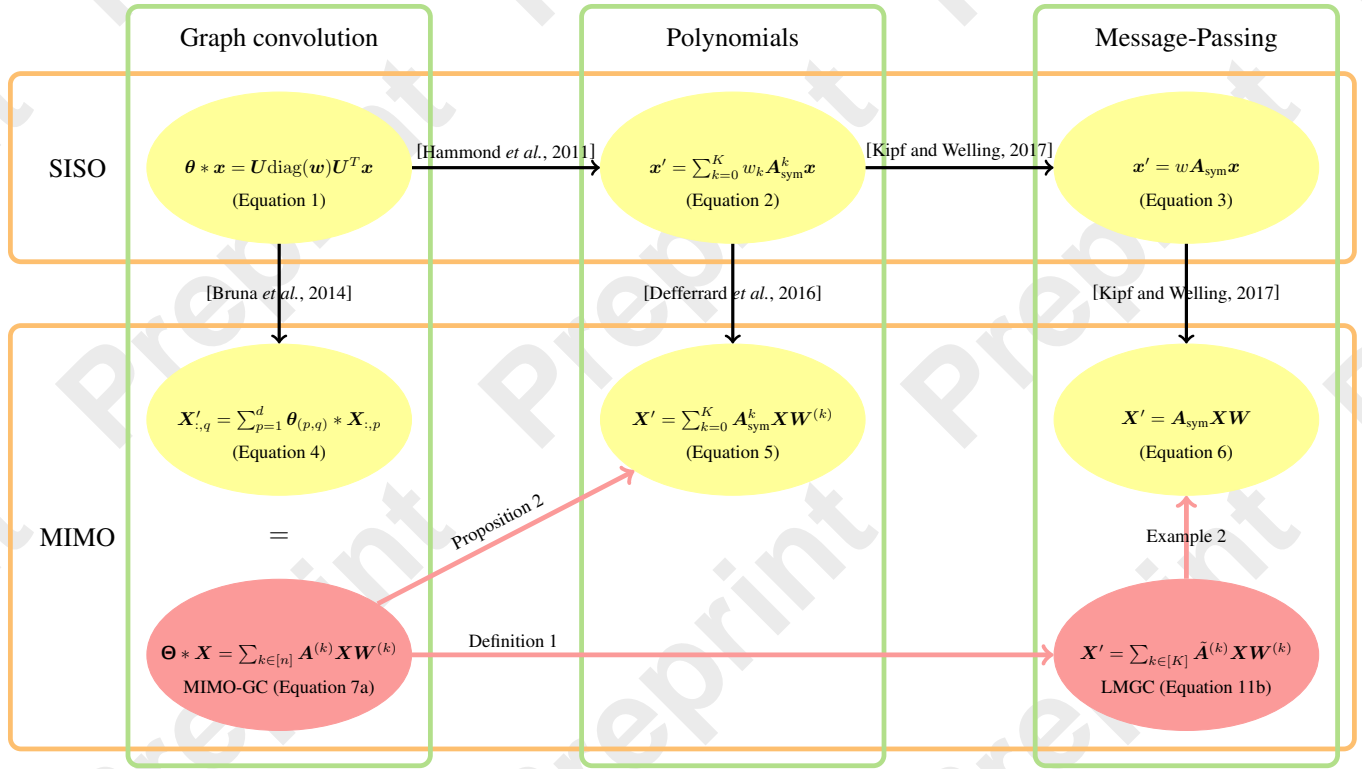


Figure 1: Connections between the graph convolution, polynomial filters, and message-passing approaches in the SISO and the MIMO case. Parts in yellow (■) indicate existing contributions, parts in pink (■) our contributions.

pendent representations when more than one computational graph is used for almost every choice of edge weights (Section 4).

## 2 Preliminaries

Let  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  be a connected and undirected graph consisting of a set of  $n$  nodes  $\mathcal{V}$  and a set of edges  $\mathcal{E}$ . Let  $A \in \{0, 1\}^{n \times n}$  be the corresponding adjacency matrix with  $A_{i,j} = 1$  if  $(i, j) \in \mathcal{E}$  and 0 otherwise. The diagonal degree matrix is  $D \in \mathbb{N}^{n \times n}$ . The symmetrically normalized adjacency matrix is given by  $A_{\text{sym}} = D^{-1/2} A D^{-1/2}$  and the graph Laplacian by  $L_{\text{sym}} = I_n - A_{\text{sym}}$ . Its eigendecomposition is  $L = U \Lambda U^T$  where  $\Lambda \in \mathbb{R}^{n \times n}$  is a diagonal matrix containing its eigenvalues, and  $U \in \mathbb{R}^{n \times n}$  is an orthonormal matrix containing the corresponding eigenvectors as columns. We refer to a vector  $x \in \mathbb{R}^n$  as a single-channel graph signal and to a matrix  $X \in \mathbb{R}^{n \times d}$  as a multi-channel graph signal. These can be initial features or expressive and informative node embeddings. In the graph domain, the Fourier base is given by the eigenvectors  $U^T$  of the graph Laplacian. Thus, the Fourier transformation  $F = U^T$  is performed by projecting a graph signal onto the eigenvectors, and its inverse transformation is given by  $F^{-1} = U$ . We further refer to  $U_{i,:} x \in \mathbb{R}$  as the component of  $x$  corresponding to vector  $U_{i,:}$ .

### 2.1 Graph Convolutions

Given a graph signal, the graph convolution or a similar method derived from it are designed to obtain a more informative graph signal. In the SISO case, the input and output are single-channel graph signals, while in the MIMO case, they are multi-channel graph signals. Graph neural networks (GNNs) are typically constructed by interleaving these operations with non-linear activation functions. The following derivations and approximations of the graph convolution are visualized in Fig. 1.

The general graph convolution is defined in the SISO case through the convolution theorem using the graph Fourier transform as

$$\theta * x = U \text{diag}(w) U^T x \quad (1)$$

where  $w = U^T \theta \in \mathbb{R}^n$  [Hammond et al., 2011; Bruna et al., 2014]. As  $x$  and  $\theta * x$  are single-channel signals, we will refer to this as the SISO graph convolution (SISO-GC).

Due to the runtime and memory complexity and inability to apply the same graph Fourier transform across graphs, most GNNs utilize approximations of the SISO-GC. Polynomials in  $A_{\text{sym}}$  (or equivalently  $L_{\text{sym}}$ ) provide a  $K$ -localized approximation

$$\theta * x \approx \sum_{k=0}^K w_{(k)} A_{\text{sym}}^k x \quad (2)$$

of the SISO-GC where  $w_k \in \mathbb{R}$  are scalars for  $k \in [K]$  [Hammond et al., 2011]. Examples of such approximations are

Chebyshev [Hammond *et al.*, 2011] and Cayley polynomials [Levie *et al.*, 2019].

Similarly, the graph convolutional network (GCN) [Kipf and Welling, 2017] was derived as a first-order localization

$$\theta * x \approx w A_{\text{sym}} x \quad (3)$$

of SISO polynomials using a single parameter  $w \in \mathbb{R}$ .

The graph convolution has not yet been derived for the MIMO case. Instead, following Bruna *et al.* [2014], the graph convolution and the described approximations are extended to the MIMO case by applying it to each combination of input channel  $p \in [d]$  and output channel  $q \in [c]$ . For the graph convolution, the output

$$X'_{:,q} = \sum_{p=1}^d \theta_{(p,q)} * X_{:,p} \quad (4)$$

is obtained by defining distinct filters  $\theta_{(p,q)} \in \mathbb{R}^n$ .

SISO polynomials are equivalently extended to the MIMO case by applying distinct parameters  $W_{p,q}^{(k)} \in \mathbb{R}$  for each combination of input and output channels [Defferrard *et al.*, 2016]. Based on Equation 2, we have

$$X'_{:,q} = \sum_{p=1}^d \sum_{k=0}^K W_{p,q}^{(k)} A_{\text{sym}}^k X_{:,p} \quad (5)$$

where  $W^{(k)} \in \mathbb{R}^{d \times c}$ .

Equivalently, the GCN is applied in the MIMO case using distinct parameters  $W_{p,q} \in \mathbb{R}$  for each combination of input channel  $p$  and output channel  $q$  [Kipf and Welling, 2017]. This led to the typical form of

$$X'_{:,q} = \sum_{p=1}^d W_{p,q} A_{\text{sym}} X_{:,p} = [A_{\text{sym}} X W]_{:,q}. \quad (6)$$

Most other message-passing methods were then further derived from the GCN. In this work, we show the advantages of directly obtaining the graph convolution and approximations in the MIMO case.

### 3 MIMO Graph Convolution

We now consider the MIMO case. Let  $X \in \mathbb{R}^{n \times d}$  be a multi-channel graph signal with  $d$  channels for each node. The multi-channel output signal  $Y \in \mathbb{R}^{n \times c}$  can have a different number of channels  $c$ . We first derive the general graph convolution for the MIMO through the convolution theorem [O’Neil, 1963] and the graph Fourier transform. The filter  $\Theta \in \mathbb{R}^{n \times c \times d}$  contains the necessary element-wise mappings from  $d$  to  $c$  dimensions. To the best of our knowledge, this has not yet been derived.

**Theorem 1** (MIMO Graph Convolution (MIMO-GC)). *Let  $X \in \mathbb{R}^{n \times d}$ ,  $\Theta \in \mathbb{R}^{n \times c \times d}$ , and the Fourier transform  $F = U^T \in \mathbb{R}^{n \times n}$  be given by the eigenvectors of the graph Laplacian  $\Lambda$ . Then, their convolution is given as*

$$(\Theta * X)(i) = \left[ \sum_{k=1}^n A^{(k)} X W^{(k)} \right]_{i,:} \quad (7a)$$

$$= \sum_{j=1}^n W_{(i,j)} X_{j,:} \in \mathbb{R}^c \quad (7b)$$

$$\text{where } A^{(k)} = U_{:,k} (U_{:,k})^T \in \mathbb{R}^{n \times n}, \\ W^{(k)} = (F(\Theta)_{k,:})^T \in \mathbb{R}^{d \times c} \text{ and } W_{(i,j)} = \left( \sum_{k=1}^n U_{i,k} U_{j,k} W^{(k)} \right)^T \in \mathbb{R}^{c \times d}.$$

We provide all detailed proofs as supplementary material<sup>1</sup>. The MIMO-GC is unique because it does not require additional definitions from us. As such, MIMO operations on graphs should closely approximate the MIMO-GC. We note that the MIMO-GC is equivalent to extending the SISO-GC to multi-channel signals by applying it to every pair of input and output channels, as introduced by Bruna *et al.* [2014]. The MIMO-GC can be interpreted in two ways.

Based on Eq. (7a), each  $A^{(k)} \in \mathbb{R}^{n \times n}$  can be seen as a fully connected computational graph with edge weights  $A_{i,j}^{(k)} = U_{i,k} \cdot U_{j,k} \in \mathbb{R}$  given by the corresponding Fourier basis vector  $U_{:,k}$ . This form is also similar to multi-head self-attention [Vaswani *et al.*, 2017]. However, they normalize edge weights by the softmax activation, preventing them from being orthogonal across heads. The corresponding parameter matrix  $W^{(k)}$  specifies how much this component is amplified or damped from each input channel to each output channel. Utilizing  $n$  computational graphs allows the MIMO-GC to amplify distinct components for each output channel. Assuming all components are present in the input signal, the MIMO-GC can produce any output signal:

**Proposition 1** (Universality of the MIMO-GC). *For any  $X \in \mathbb{R}^{n \times d}$  with  $U^T X \neq_{em} 0$  element-wise non-zero and any  $Y \in \mathbb{R}^{n \times c}$ , there exists a  $\Theta \in \mathbb{R}^{n \times c \times d}$ , such that*

$$\Theta * X = Y. \quad (8)$$

Based on Eq. (7b), the MIMO-GC can also be interpreted as applying distinct feature transformation  $W_{(i,j)}$  for each pair of nodes. Each  $W_{(i,j)}$  is a unique linear combination of a shared set of  $n$  feature transformations. Relatedly, utilizing distinct feature transformations was recently popularized as Neural Sheaf Diffusion [Hansen and Gebhart, 2020; Bodnar *et al.*, 2022]. The MIMO-GC provides an additional theoretical justification for such methods.

However, computing the MIMO-GC exactly is typically not desirable, as with the SISO-GC. It is inherently transductive, as the graph Fourier transform is graph-dependent, and thus, a learned filter cannot be applied to novel or changed graphs. Most importantly, the computational complexity of the graph convolution scales quadratically with the number of nodes:

**Computational Complexity.** Equivalently to computing the SISO-GC exactly, the total complexity of the MIMO-GC is dominated by the graph Fourier transform as it requires dense matrix multiplications. The overall complexity is thus  $O(n^2 \cdot c \cdot d)$ .

**Benefiting from the MIMO-GC.** Instead of directly computing the MIMO-GC, we aim to improve the approximations previously derived from the SISO-GC, which were then extended to the MIMO case. We first confirm that these MIMO polynomials are also approximations of the MIMO-GC with constraints on the allowed filters  $\Theta$ :

<sup>1</sup>Our appendix is available at <https://arxiv.org/abs/2505.11346>

**Proposition 2** (Every MIMO polynomial filter is a MIMO-GC with a specific filter). *Let  $\mathbf{X} \in \mathbb{R}^{n \times d}$  for some  $d \in \mathbb{N}$ . For any  $\mathbf{V}^{(0)}, \dots, \mathbf{V}^{(K)} \in \mathbb{R}^{d \times c}$  with  $c, K \in \mathbb{N}$ , there exists a  $\Theta_{\text{poly}} \in \mathbb{R}^{n \times c \times d}$ , such that*

$$\sum_{k=0}^K A_{\text{sym}}^k \mathbf{X} \mathbf{V}^{(k)} = \Theta_{\text{poly}} * \mathbf{X}. \quad (9)$$

As one such example of a first-degree polynomial, the GCN is a MIMO-GC with specific constraints on  $\Theta$ :

**Example 1** (GCN is a MIMO-GC). *Let  $\mathbf{X} \in \mathbb{R}^{n \times d}$ ,  $\mathbf{V} \in \mathbb{R}^{d \times c}$ . Then,*

$$\begin{aligned} A_{\text{sym}} \mathbf{X} \mathbf{V} &= \sum_{k=1}^n \lambda_j \mathbf{U}_{:,k} (\mathbf{U}_{:,k})^T \mathbf{X} \mathbf{V} \\ &= \sum_{k=1}^n \mathbf{U}_{:,k} (\mathbf{U}_{:,k})^T \mathbf{X} \mathbf{W}^{(k)} \\ &= \Theta_{\text{GCN}} * \mathbf{X} \end{aligned} \quad (10)$$

where  $\mathbf{W}^{(k)} = \lambda_j \mathbf{V}$  and corresponding  $\Theta_{\text{GCN}} \in \mathbb{R}^{n \times c \times d}$ .

As a first step, the MIMO-GC helps us with the understanding of properties of various approximations and can consequently improve these approximations. Based on Example 1, the GCN utilizes a single shared parameter matrix  $\mathbf{V}$  across all components. Each component is then amplified according to its respective eigenvalue, which is shared across all combinations of input and output channels. Other message-passing operations may utilize a different matrix  $\tilde{\mathbf{A}}$  instead of  $A_{\text{sym}}$ . However, as using any single computational graph  $\tilde{\mathbf{A}}$  can be similarly decomposed, the amplification of components is fixed and shared across all feature channels for any given  $\tilde{\mathbf{A}}$ . We refer to this phenomenon as shared component amplification (SCA). When repeatedly applying such filters or message-passing operations, SCA leads to the well-known phenomenon of over-smoothing and, more generally, rank collapse [Oono and Suzuki, 2020; Roth and Liebig, 2023]. We provide further details on this phenomenon in our appendix.

Contrarily, the MIMO-GC requires multiple computational graphs to amplify different components across feature channels. Equivalently, applying distinct feature transformations for each node pair can improve approximations. Developing approximations with these properties can lead to more effective learning on graph-structured data.

## 4 Localized MIMO Graph Convolutions

Based on Eq. (7b), we localize the MIMO-GC by aggregating over the neighboring nodes instead of all nodes of a given graph:

**Definition 1.** *We define the Localized MIMO Graph Convolution (LMGC) as:*

$$\mathbf{x}'_{(i)} = \sum_{v_j \in N_i} \mathbf{W}_{(i,j)} \mathbf{x}_{(j)} \quad (11a)$$

$$= \left[ \sum_{k \in [K]} \tilde{\mathbf{A}}^{(k)} \mathbf{X} \mathbf{W}^{(k)} \right]_{i,:} \quad (11b)$$

where  $K \in \mathbb{N}$  and each  $\mathbf{W}_{(i,j)} = \sum_{k \in [K]} \alpha_{(k)}^{(i,j)} \mathbf{W}^{(k)} \in \mathbb{R}^{c \times d}$  is linear combination based on  $\alpha_{(1)}^{(i,j)}, \dots, \alpha_{(K)}^{(i,j)} \in \mathbb{R}$  and  $\mathbf{W}^{(1)}, \dots, \mathbf{W}^{(K)} \in \mathbb{R}^{d \times c}$ . The entries  $A_{i,j}^{(k)} = \alpha_{(k)}^{(i,j)}$  are given by the corresponding coefficients.

In this definition, the number of terms  $K$  and the coefficients or edge weights  $\alpha_{(k)}^{(i,j)}$  can be freely chosen, which allows methods that do not use the expensive eigenvector computation. The LMGC is permutation equivariant if the coefficients  $\alpha_{(k)}^{(i,j)}$  are also equivariant, for example, when derived from a function of the nodes  $v_i$  and  $v_j$ . The LMGC can also be applied across different graphs and for directed graphs. As with MIMO-GCs, the LMGC can be equivalently restated as operating on  $K$  computational graphs. The edge weights of the  $k$ -th computational graph are given by  $\alpha_{(k)}^{(i,j)}$ . Consequently, the LMGC can represent many linear MPNNs for different values for  $\alpha_{(k)}^{(i,j)}$ . We provide three examples below:

**Example 2** (GCN [Kipf and Welling, 2017]). *Let  $\mathbf{V} \in \mathbb{R}^{c \times d}$  be a feature transformation. The update step*

$$\mathbf{x}'_{(i)} = \sum_{v_j \in N_i} \frac{1}{\sqrt{d_i} \sqrt{d_j}} \mathbf{V} \mathbf{x}_{(j)} \quad (12)$$

is an LMGC with  $K = 1$ ,  $\mathbf{W}^{(1)} = \mathbf{V}$ , and  $\alpha_{(1)}^{(i,j)} = \frac{1}{\sqrt{d_i} \sqrt{d_j}}$  where  $d_i, d_j \in \mathbb{N}$  are the degrees of nodes  $v_i$  and  $v_j$ , respectively.

As the MIMO-GC is similar to multi-head self-attention, the LMGC is related to local multi-head attention-based methods while allowing for more flexible attention scores, i.e., scores do not need to sum to one for every node:

**Example 3** (GAT [Velickovic et al., 2018]). *Let  $H$  be the number of heads,  $\mathbf{V}^{(h)}$  the linear transformation of head  $h \in [H]$ , and  $a_{(h)}^{(i,j)} \in \mathbb{R}$  the attention score between nodes  $v_i$  and  $v_j$ . The update step*

$$\mathbf{x}'_{(i)} = \sum_{h \in [H]} \sum_{v_j \in N_i} a_{(h)}^{(i,j)} \mathbf{V}^{(h)} \mathbf{x}_{(j)} \quad (13)$$

is an LMGC with  $K = H$ ,  $\mathbf{W}^{(h)} = \mathbf{V}^{(h)}$  and  $\alpha_{(h)}^{(i,j)} = a_{(h)}^{(i,j)}$ .

The LMGC can also represent gating mechanisms, e.g., the GatedGCN [Dwivedi et al., 2023] or neural sheaf diffusion [Hansen and Gebhart, 2020].

The general form of the LMGC allows for a more focused development of novel and powerful methods. With specific choices of  $\alpha_{(k)}^{(i,j)}$ , the LMGC can model a symmetric or directed flow of information and can construct anisotropic or isotropic messages.

**Theoretical Properties.** Studying theoretical properties of LMGCs reduces to studying the effects of coefficients  $\alpha_{(k)}^{(i,j)}$ . For example, the LMGC cannot represent non-linear feature transformations, which are typically used to ensure injectivity, e.g., by GIN [Xu et al., 2019]. This allows GNNs to match the expressivity of the Weisfeiler-Leman graph isomorphism test [Leman and Weisfeiler, 1968], a key property for

graph-level tasks. However, we find that any LMGC with  $K > 0$  computational graphs is also injective for almost every choice coefficients  $\alpha_{(k)}^{(i,j)}$  without requiring a non-linear feature transformation:

**Proposition 3** (Injectivity). *Let  $f(x_{(i)}, N_i) = \sum_{x_{(j)} \in N_i} W_{(i,j)} x_{(j)}$  be an LMGC with  $K \geq 1$  and  $\mathcal{X}$  a countable set. Then,  $f(x_{(p)}, \mathcal{X}_p)$  is injective for finite multisets  $\mathcal{X}_p \subset \mathcal{X}$  and elements  $x_{(p)} \in \mathcal{X}$  for a.e. choice of coefficients  $\alpha_{(k)}^{(i,j)}$  and a.e.  $W^{(k)}$  for all  $k \in [K]$ .*

Different components can be amplified across feature channels when further using  $K > 1$  computational graphs. The resulting node representations are linearly independent for almost every choice of coefficients  $\alpha_{(k)}^{(i,j)}$ . This prevents the shared component amplification of methods utilizing a single computational graph.

**Proposition 4** (Linear Independence). *Let  $f(x_{(i)}, N_i) = \sum_{x_{(j)} \in N_i} W_{(i,j)} x_{(j)}$  be an LMGC with  $K > 1$  and  $\mathcal{X}$  a countable set. Then,  $f(x_{(i)}, \mathcal{X}_1)$  is linearly independent to  $f(x_{(j)}, \mathcal{X}_2)$  for all finite multisets  $\mathcal{X}_1, \mathcal{X}_2 \subset \mathcal{X}$  with  $\mathcal{X}_1 \neq c \cdot \mathcal{X}_2$  for any  $c \in \mathbb{N}$  and elements  $x_{(i)}, x_{(j)} \in \mathcal{X}$  for a.e. choice of coefficients  $\alpha_{(k)}^{(i,j)}$  and a.e.  $W^{(k)}$  for all  $k \in [K]$ .*

This result aligns with previous findings that identified cases where multiple computational graphs can ensure linearly independent representations [Roth *et al.*, 2024]. Importantly, each  $\alpha_{(k)}^{(i,j)}$  can be independently obtained, e.g., by a function  $\alpha_{(k)}^{(i,j)} = \phi_k(x_{(i)}, x_{(j)}) \in \mathbb{R}$  of the corresponding node states. Many functions  $\phi_k$  satisfy Proposition 3 and Proposition 4. A neural network can then approximate such a function. As a negative example of such a functions, softmax-activated attention scores do not satisfy the a.e. condition as the space of scores forms a measure-zero set, e.g., for GAT [Velickovic *et al.*, 2018] and the more powerful GATv2 [Brody *et al.*, 2022]. As has been pointed out by several works [Xu *et al.*, 2019], such methods cannot distinguish multisets of different multiplicities, e.g., when  $\mathcal{X}_1 = \{\{x_1\}\}$  and  $\mathcal{X}_2 = \{\{x_1, x_1\}\}$ . Other methods, such as FAGCN [Bo *et al.*, 2021] and GGCN [Yan *et al.*, 2022], proposed to apply the tanh activation function instead, which does not constrain the outputs to a measure-zero set.

Thus, an LMGC can incorporate the advantages of attention-based by filtering incoming messages and preventing the shared component amplification across feature channels by utilizing multiple computational graphs. At the same time, it applies linear feature transformations and can be injective on multisets, as in GIN.

**An LMGC Instantiation.** When constructing an LMGC instantiation, only the number of computational graphs  $K$  and the coefficients  $\alpha_{(k)}^{(i,j)}$  for all  $k \in [K]$  need to be defined. For our empirical study, we define a simple LMGC instantiation as a mix of GATv2 and FAGCN. We define the coefficients as

$$\alpha_{(k)}^{(i,j)} = \phi_k(x_{(i)}, x_{(j)}) := \sigma_2(v_{(k)}^T \sigma_1(W^{(1)} x_{(i)} || \dots || W^{(K)} x_{(i)} || W^{(1)} x_{(j)} || \dots || W^{(K)} x_{(j)})) \quad (14)$$

Method	MSE
GATv2	$0.12 \pm 0.04$
FAGCN	$0.68 \pm 0.02$
ACM	$0.49 \pm 0.02$
GIN	<u><math>0.08 \pm 0.03</math></u>
LMGC	<b><math>25 \cdot 10^{-9} \pm 86 \cdot 10^{-11}</math></b>

Table 1: Results for the universality task. Given representations  $\mathbf{X}, \mathbf{Y}$  and a graph  $\mathbf{A}$ , one layer of each method is optimized to approximate the function  $f(\mathbf{X}, \mathbf{A}) = \mathbf{Y}$ . Average and standard deviation of the minimal mean-squared error (MSE) during optimization. Best MSE in **bold**, second-best underlined.

where  $v_{(k)} \in \mathbb{R}^{2 \cdot K \cdot c}$  are learnable vectors for  $k \in [K]$ ,  $\sigma_1$  is the LeakyReLU activation and  $\sigma_2$  is the tanh activation function. The execution time is slightly favorable compared to GATv2, as we do not normalize the messages.

## 5 Related Work

We now describe previous works related to various parts of the MIMO-GC and the LMGC.

**Graph Convolutions.** Bruna *et al.* [2014] extend the SISO-GC to the MIMO case by utilizing a filter between all pairs of input and output channels. This extension is equivalent to the MIMO-GC directly derived through the convolution theorem. Approximations are derived in the SISO case and mapped to the MIMO case using the same procedure afterward. Hammond *et al.* [2011] propose to approximate the SISO-GC using Chebyshev polynomials in the SISO case. Defferrard *et al.* [2016] employ separate filters for pairs of input and output channels to extend Chebyshev polynomials to the MIMO case. Sandryhaila [2013] define general polynomial graph filters for the SISO case. Using the same procedure, Gama *et al.* [2018] extend these polynomial graph filters to the MIMO case. Kipf and Welling [2017] derive the GCN as a 1-localized approximation of the SISO Chebyshev polynomials. They equivalently extend it to the MIMO case afterward by applying separate parameters for each combination of input and output channels. Most other MPNNs are derived from the GCN to mitigate various shortcomings [Velickovic *et al.*, 2018; Xu *et al.*, 2019; Roth and Liebig, 2022]. Directly approximating the MIMO-GC allows us to benefit from MIMO-specific properties of the graph convolution.

**MIMO Improvements.** While most MPNNs are applied to the MIMO case, many of these are well-known to be unable to amplify distinct components across channels, a phenomenon known as over-smoothing [Oono and Suzuki, 2020], over-correlation [Jin *et al.*, 2022], or rank collapse [Roth and Liebig, 2023; Roth, 2024]. Various methods have been proposed to improve multi-channel learning within MPNNs. Luan *et al.* [2022] propose to apply separate graph filters for different feature channels. In ADR-GNNs [Eliasof *et al.*, 2023], feature channels are separately aggregated using channel-specific edge weights. Other works similarly propose to apply distinct filters across channels [Liu *et al.*, 2025]. Zhou *et al.* [2020] propose the multi-channel graph neural

Method	Basic	+ LapPE	+ Jumping Knowledge	+ Residual	+ All three
GATv2	0.377 ± 0.024	0.341 ± 0.040	0.388 ± 0.017	0.311 ± 0.016	0.294 ± 0.019
FAGCN	0.365 ± 0.018	0.349 ± 0.038	0.352 ± 0.042	0.289 ± 0.019	0.232 ± 0.012
ACM	0.278 ± 0.006	0.281 ± 0.019	0.288 ± 0.008	0.266 ± 0.017	0.238 ± 0.006
GIN	0.272 ± 0.009	0.259 ± 0.012	0.267 ± 0.020	0.240 ± 0.005	0.228 ± 0.014
LMGC	<b>0.241</b> ± 0.018	<b>0.234</b> ± 0.009	<b>0.233</b> ± 0.019	<b>0.215</b> ± 0.006	<b>0.203</b> ± 0.004

Table 2: Test MAE results on ZINC12k. LapPE indicates that a Laplacian position encoding is concatenated to the initial features. For Jumping Knowledge, the channel-wise maximum value after each iteration is used for each after the message-passing steps. Residual indicates that the input to each message-passing step is added to its output. With + All three, these three techniques are simultaneously applied. Best scores in **bold**, second-best underlined.

Method	Texas	Cornell	Wisconsin	Film	Chameleon	Squirrel
GATv2	71.6 ± 1.0	66.1 ± 0.6	79.1 ± 2.0	35.1 ± 0.2	<u>47.1</u> ± 0.3	<u>35.1</u> ± 0.2
FAGCN	<u>73.5</u> ± 1.8	<u>68.1</u> ± 1.9	<u>80.2</u> ± 1.8	<u>36.0</u> ± 0.3	46.9 ± 0.5	34.6 ± 0.3
ACM	72.3 ± 0.4	65.1 ± 0.7	74.2 ± 0.9	35.8 ± 0.3	45.5 ± 0.9	34.5 ± 0.1
GIN	70.5 ± 1.1	66.1 ± 1.0	79.0 ± 0.6	34.1 ± 0.3	46.1 ± 0.4	34.6 ± 0.5
LMGC	<b>74.2</b> ± 2.2	<b>68.9</b> ± 2.2	<b>81.4</b> ± 1.1	<b>36.3</b> ± 0.4	<b>49.8</b> ± 0.8	<b>35.9</b> ± 0.5

Table 3: Test accuracy on heterophilic node classification tasks. Best scores in **bold**, second-best underlined. All models contain at most 100 000 parameters and the same hyperparameter optimization was applied.

network that obtains multiple computational graphs through a pooling operation and learns interaction scores between graphs. Utilizing multiple computational graphs has been extensively studied in mitigating over-smoothing and representational rank collapse [Roth *et al.*, 2024]. Applying different linear transformations between pairs of nodes has also been derived within neural sheaf diffusion [Hansen and Gebhart, 2020; Bodnar *et al.*, 2022]. As the MIMO-GC and LMGC naturally allow multi-channel learning, these frameworks can be closer aligned as approximations of the MIMO-GC. The LMGC can equivalently be interpreted as message-passing on multigraphs. Butler *et al.* [2023] introduced convolutional multigraph neural networks that utilize polynomial filters on multigraphs.

**Approaches Related to the LMGC.** The LMGC is closely related to several existing methods. As described in Example 3, multi-head attention-based methods like GAT [Velickovic *et al.*, 2018] and GATv2 [Brody *et al.*, 2022] are LMGCs with constraints on the attention scores by applying the softmax activation. By lifting this constraint, LMGCs can be injective on multisets (Proposition 3). Several other methods have been proposed to replace the softmax activation. The FAGCN [Bo *et al.*, 2021] instead applies the tanh activation function to amplify high-frequencies or low-frequencies. Similarly, the GGCN [Yan *et al.*, 2022] allows learning of signed edge weights. Other studies considered replacing the softmax activation function within transformers and self-attention modules. Wortsman *et al.* [2023] apply the ReLU activation in vision transformers. Saratchandran *et al.* [2024] found empirical success using polynomial activation functions for self-attention. However, as self-attention typically considers a fully connected graph, these works did not study distinguishing structural differences. Contrarily, Proposition 3 shows that differences in the number of neighbors can

be distinguished without the softmax activation.

## 6 Experiments

We now want to confirm the beneficial properties of LMGCs. As the LMGC can match the expressive power of GIN, we want to evaluate whether it can match the performance of GIN for graph-level tasks. We also evaluate whether the LMGC can match the performance of attention-based methods for node-level tasks. All experiments are run on an H100 GPU. Additional details on all models, datasets, and hyperparameters are provided as supplementary material.<sup>2</sup>

### 6.1 Methods

We consider the following four message-passing methods across all experiments. We conduct all results ourselves using the same hyperparameter ranges across methods.

**GATv2.** This method extends GAT with dynamic attention [Brody *et al.*, 2022]. Attention-based methods are particularly effective for node-level prediction tasks due to their ability to filter information. We utilize the standard implementation that corresponds to an LMGC with  $\alpha_{(k)}^{(i,j)} = \sigma_2(\mathbf{v}_{(k)}^T \sigma_1(\mathbf{W}^{(k)} \mathbf{x}_{(i)} + \mathbf{W}^{(k)} \mathbf{x}_{(j)}))$  where  $\sigma_1$  is the LeakyReLU activation function and  $\sigma_2$  is the node-wise softmax activation function and  $\mathbf{v}^{(k)} \in \mathbb{R}^c$  is a learnable vector. We set the number of heads to  $K = 4$  for all experiments.

**FAGCN.** This method was designed for heterophilic node classification tasks by allowing for negative edge weights [Bo *et al.*, 2021]. Stated in the LMGC framework, we evaluate a method that sets  $K = 1$  and  $\alpha_{(1)}^{(i,j)} = \frac{\sigma(\mathbf{v}[\mathbf{x}_{(i)}] \|\mathbf{x}_{(j)}])}{\sqrt{d_i} \sqrt{d_j}}$  where  $\sigma$  is the tanh activation function,  $\mathbf{v} \in \mathbb{R}^{2 \cdot d}$  is a learnable vector

<sup>2</sup>Our implementation is available at <https://github.com/roth-andreas/mimo-graph-convolutions>.



and  $d_i, d_j$  are the degrees of nodes  $v_i$  and  $v_j$ , respectively. Despite the tanh activation, FAGCN is not always injective due to the degree normalization.

**ACM.** Written in the LMGC framework, the adaptive channel mixing (ACM) [Luan *et al.*, 2022] proposes to utilize  $\tilde{A}^{(1)} = A_{\text{sym}}$  for amplifying low-frequency components and  $\tilde{A}^{(2)} = L_{\text{sym}}$  for amplifying high-frequency components. They further propose a third computational graph  $\tilde{A}^{(3)} = I$ , which we utilize whenever residual connections are used in an experiment.

**GIN.** For graph-level tasks, the GIN [Xu *et al.*, 2019] is particularly effective as it can match the expressivity of the WL-test due to the non-linear feature transformation. As the non-linear feature transformation, we apply a two-layer MLP with ReLU activations for all experiments.

**LMGC.** As the LMGC can combine the favorable properties of the other three methods, we utilize the instantiation of the LMGC as described in Eq. (14). As with GATv2, we set the total number of heads to  $K = 4$  for all experiments.

## 6.2 Universality

Based on Proposition 1, the MIMO-GC can represent almost every mapping  $\Theta * X = Y$  with  $X \in \mathbb{R}^{n \times d}$  and  $Y \in \mathbb{R}^{n \times c}$ . We evaluate the ability of message-passing approaches to approximate such a transformation. We sample a random undirected and connected Erdős–Rényi graph [Erdős and Rényi, 1959] with  $n = 16$  nodes and an edge probability of  $p = 10\%$ . We set  $d = c = 16$ . Similarly,  $X \in \mathbb{R}^{n \times d}$  and  $Y \in \mathbb{R}^{n \times c}$  are randomly sampled with  $X_{i,j} \sim \mathcal{N}(0, 1)$  and  $Y_{i,j} \sim \mathcal{N}(0, 1)$ . We apply a single message-passing layer as  $f(X, A)$ . We minimize the mean-squared error (MSE) between  $f(X, A)$  and  $Y$  using the Adam optimizer for 40 000 steps. The learning rate is tuned in  $\{0.03, 0.01, 0.003\}$ .

The minimum achieved approximation error averaged over three runs is presented in Table 1. LMGC achieves a significantly lower error than GATv2, FAGCN, and GIN. As the MIMO-GC can represent such a function exactly, LMGCs benefit from this property as a close approximation. While these improved capabilities come with the risk of overfitting, we expect LMGCs to be particularly beneficial for challenging tasks.

## 6.3 Graph-Level Prediction

GIN is typically used for graph-level tasks due to its expressive power. Based on Proposition 3, we now want to validate that the LMGC can match these results empirically. We consider the challenging ZINC12k dataset [Sterling and Irwin, 2015]. It consists of around 12 000 molecular graphs, with the task being to predict the constrained solubility of each molecule. We integrate all models into the implementation of GraphGPS [Rampásek *et al.*, 2022] and the Long Range Graph Benchmark [Dwivedi *et al.*, 2022]. Based on Toenshoff *et al.* [2024], we optimize the number of layers in  $\{6, 8, 10\}$  and the learning rate in  $\{0.001, 0.0003, 0.0001\}$  using a grid search. Each model utilizes at most 100 000 parameters to ensure fairness.

In Table 2, we present results of a detailed study in which we combine these base message-passing methods with various other established techniques. These techniques are Laplacian positional encoding (LapPE) [Kreuzer *et al.*, 2021], jumping knowledge [Xu *et al.*, 2018] and residual connections [He *et al.*, 2016]. We find all methods to benefit from these techniques, with the LMGC achieving the best results in all cases. We provide additional results, including runtimes and training losses, as supplementary material.

## 6.4 Node Classification

While expressivity is a key property for graph-level tasks, attention-based methods typically outperform GIN on node-level tasks due to their ability to filter messages [Brody *et al.*, 2022]. Thus, we also evaluate whether the LMGC can match the performance of GATv2 and FAGCN on these tasks. We consider six heterophilic benchmark datasets for node classification: Texas, Cornell, Wisconsin, Film, Chameleon, and Squirrel. We use the ten splits into train, validation, and test sets proposed by Pei *et al.* [2020]. We integrate all models into the implementation from Rusch *et al.* [2023]. As with ZINC, each model uses at most 100 000 parameters. For each method, we tune the learning rate in  $\{0.01, 0.003, 0.001\}$  and dropout ratio in  $\{0.0, 0.25, 0.5\}$  using a grid search, as these affected the results the most. Based on the optimal hyperparameters for the validation set, we rerun each method five times for all ten splits and report average test results.

These average test accuracies are presented in Table 3. GIN achieves the lowest accuracy, and LMGC achieves the highest accuracy across all tasks. While the differences are only a few percentage points, these experiments confirm that the LMGC can combine the benefits of GATv2, FAGCN, GIN, and the MIMO-GC into a single model.

## 7 Conclusion

This work derives the MIMO graph convolution (MIMO-GC) using the convolution theorem and emphasizes the advantages of approximating the graph convolution in the MIMO case rather than the SISO case. A key property of the MIMO-GC is operating on multiple computational graphs or equivalently applying distinct linear transformations for each node pair. We have proven that the localized form is injective and results in linearly independent representations for almost every choice of edge weights. Due to our direct theoretical derivation from the MIMO-GC and the generality of the LMGC framework, studying properties of message-passing operations can now focus on analyzing the coefficients  $\alpha_{(k)}^{(i,j)}$ . This allows the development of more effective methods within a well-defined framework. While we have confirmed the advantages and potential of the LMGC framework, identifying optimal instantiations of LMGCs for specific tasks remains open.

## Acknowledgments

Part of this research has been funded by the Federal Ministry of Education and Research of Germany and the state of North-Rhine Westphalia as part of the Lamarr-Institute

for Machine Learning and Artificial Intelligence and by the Federal Ministry of Education and Research of Germany under grant no. 01IS22094E WEST-AI. Simulations were performed with computing resources granted by WestAI under project rwth1631.

## References

- [Bo *et al.*, 2021] Deyu Bo, Xiao Wang, Chuan Shi, and Huawei Shen. Beyond low-frequency information in graph convolutional networks. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, Virtual Event, February 2-9*, pages 3950–3957. AAAI Press, 2021.
- [Bodnar *et al.*, 2022] Cristian Bodnar, Francesco Di Giovanni, Benjamin Paul Chamberlain, Pietro Lió, and Michael M. Bronstein. Neural sheaf diffusion: A topological perspective on heterophily and oversmoothing in gnns. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- [Brody *et al.*, 2022] Shaked Brody, Uri Alon, and Eran Yahav. How attentive are graph attention networks? In *The Tenth International Conference on Learning Representations, Virtual Event, April 25-29, 2022*.
- [Bruna *et al.*, 2014] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. Spectral networks and locally connected networks on graphs. In *2nd International Conference on Learning Representations, Banff, AB, Canada, April 14-16, Conference Track Proceedings, 2014*.
- [Butler *et al.*, 2023] Landon Butler, Alejandro Parada-Mayorga, and Alejandro Ribeiro. Convolutional learning on multigraphs. *IEEE Trans. Signal Process.*, 71:933–946, 2023.
- [Defferrard *et al.*, 2016] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, Barcelona, Spain*, pages 3837–3845, 2016.
- [Dwivedi *et al.*, 2022] Vijay Prakash Dwivedi, Ladislav Rampásek, Michael Galkin, Ali Parviz, Guy Wolf, Anh Tuan Luu, and Dominique Beaini. Long range graph benchmark. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- [Dwivedi *et al.*, 2023] Vijay Prakash Dwivedi, Chaitanya K. Joshi, Anh Tuan Luu, Thomas Laurent, Yoshua Bengio, and Xavier Bresson. Benchmarking graph neural networks. *J. Mach. Learn. Res.*, 24:43:1–43:48, 2023.
- [Eliasof *et al.*, 2023] Moshe Eliasof, Eldad Haber, and Eran Treister. ADR-GNN: advection-diffusion-reaction graph neural networks. *CoRR*, abs/2307.16092, 2023.
- [Erdős and Rényi, 1959] P Erdős and A Rényi. On random graphs i. *Publicationes Mathematicae Debrecen*, 6:290–297, 1959.
- [Gama *et al.*, 2018] Fernando Gama, Antonio G. Marques, Alejandro Ribeiro, and Geert Leus. MIMO graph filters for convolutional neural networks. In *19th IEEE International Workshop on Signal Processing Advances in Wireless Communications, Kalamata, Greece, June 25-28*, pages 1–5. IEEE, 2018.
- [Hammond *et al.*, 2011] David K Hammond, Pierre Vandergheynst, and Rémi Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129–150, 2011.
- [Hansen and Gebhart, 2020] Jakob Hansen and Thomas Gebhart. Sheaf neural networks. *CoRR*, abs/2012.06333, 2020.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, June 27-30*, pages 770–778. IEEE Computer Society, 2016.
- [Hu *et al.*, 2021] Weihua Hu, Matthias Fey, Hongyu Ren, Maho Nakata, Yuxiao Dong, and Jure Leskovec. OGB-LSC: A large-scale challenge for machine learning on graphs. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks 1, December 2021, virtual*, 2021.
- [Jin *et al.*, 2022] Wei Jin, Xiaorui Liu, Yao Ma, Charu C. Aggarwal, and Jiliang Tang. Feature overcorrelation in deep graph neural networks: A new perspective. In *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 14 - 18*, pages 709–719. ACM, 2022.
- [Kipf and Welling, 2017] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings, 2017*.
- [Kreuzer *et al.*, 2021] Devin Kreuzer, Dominique Beaini, William L. Hamilton, Vincent Létourneau, and Prudencio Tossou. Rethinking graph transformers with spectral attention. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, December 6-14, virtual*, pages 21618–21629, 2021.
- [Leman and Weisfeiler, 1968] Andrei Leman and Boris Weisfeiler. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Tekhnicheskaya Informatsiya*, 2(9):12–16, 1968.
- [Levie *et al.*, 2019] Ron Levie, Federico Monti, Xavier Bresson, and Michael M. Bronstein. Cayleynets: Graph convolutional neural networks with complex rational spectral filters. *IEEE Trans. Signal Process.*, 67(1):97–109, 2019.
- [Liu *et al.*, 2025] Xun Liu, Alex Hay-Man Ng, Fangyuan Lei, Yikuan Zhang, and Zhengming Li. Gpnet: Simpli-



- fying graph neural networks via multi-channel geometric polynomials. *Inf. Sci.*, 694:121696, 2025.
- [Luan *et al.*, 2022] Sitao Luan, Chenqing Hua, Qincheng Lu, Jiaqi Zhu, Mingde Zhao, Shuyuan Zhang, Xiao-Wen Chang, and Doina Precup. Revisiting heterophily for graph neural networks. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- [Oono and Suzuki, 2020] Kenta Oono and Taiji Suzuki. Graph neural networks exponentially lose expressive power for node classification. In *8th International Conference on Learning Representations Addis Ababa, Ethiopia, April 26-30, 2020*.
- [O’Neil, 1963] Richard O’Neil. Convolution operators and  $L(p, q)$  spaces. *Duke Mathematical Journal*, 30(1):129 – 142, 1963.
- [Pei *et al.*, 2020] Hongbin Pei, Bingzhe Wei, Kevin Chen-Chuan Chang, Yu Lei, and Bo Yang. Geom-gcn: Geometric graph convolutional networks. In *8th International Conference on Learning Representations, Addis Ababa, Ethiopia, April 26-30, 2020*.
- [Rampásek *et al.*, 2022] Ladislav Rampásek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini. Recipe for a general, powerful, scalable graph transformer. In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- [Roth and Liebig, 2022] Andreas Roth and Thomas Liebig. Transforming pagerank into an infinite-depth graph neural network. In *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2022, Grenoble, France, September 19-23, 2022, Proceedings, Part II*, volume 13714, pages 469–484. Springer, 2022.
- [Roth and Liebig, 2023] Andreas Roth and Thomas Liebig. Rank collapse causes over-smoothing and over-correlation in graph neural networks. In *Learning on Graphs Conference, 27-30 November 2023, Virtual Event*, volume 231, page 35. PMLR, 2023.
- [Roth *et al.*, 2024] Andreas Roth, Franka Bause, Nils M Kriege, and Thomas Liebig. Preventing representational rank collapse in mpnns by splitting the computational graph. *CoRR*, abs/2409.11504, 2024.
- [Roth, 2024] Andreas Roth. Simplifying the theory on over-smoothing. *CoRR*, abs/2407.11876, 2024.
- [Rusch *et al.*, 2023] T. Konstantin Rusch, Benjamin Paul Chamberlain, Michael W. Mahoney, Michael M. Bronstein, and Siddhartha Mishra. Gradient gating for deep multi-rate learning on graphs. In *The Eleventh International Conference on Learning Representations, Kigali, Rwanda, May 1-5, 2023*.
- [Sandryhaila and Moura, 2013] Aliaksei Sandryhaila and José MF Moura. Discrete signal processing on graphs. *IEEE transactions on signal processing*, 61(7):1644–1656, 2013.
- [Saratchandran *et al.*, 2024] Hemanth Saratchandran, Jian-qiao Zheng, Yiping Ji, Wenbo Zhang, and Simon Lucey. Rethinking softmax: Self-attention with polynomial activations. *CoRR*, abs/2410.18613, 2024.
- [Sterling and Irwin, 2015] Teague Sterling and John J. Irwin. ZINC 15 - ligand discovery for everyone. *J. Chem. Inf. Model.*, 55(11):2324–2337, 2015.
- [Tönshoff *et al.*, 2024] Jan Tönshoff, Martin Ritzert, Eran Rosenbluth, and Martin Grohe. Where did the gap go? reassessing the long-range graph benchmark. *Trans. Mach. Learn. Res.*, 2024, 2024.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, Long Beach, CA, USA*, pages 5998–6008, 2017.
- [Velickovic *et al.*, 2018] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph attention networks. In *6th International Conference on Learning Representations, Vancouver, BC, Canada, April 30 - May 3, 2018*.
- [Wortsman *et al.*, 2023] Mitchell Wortsman, Jaehoon Lee, Justin Gilmer, and Simon Kornblith. Replacing softmax with relu in vision transformers. *CoRR*, abs/2309.08586, 2023.
- [Xu *et al.*, 2018] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In *Proceedings of the 35th International Conference on Machine Learning, Stockholm, Sweden, July 10-15*, volume 80, pages 5449–5458. PMLR, 2018.
- [Xu *et al.*, 2019] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *7th International Conference on Learning Representations, New Orleans, LA, USA, May 6-9, 2019*.
- [Yan *et al.*, 2022] Yujun Yan, Milad Hashemi, Kevin Swersky, Yaoqing Yang, and Danai Koutra. Two sides of the same coin: Heterophily and oversmoothing in graph convolutional neural networks. In *IEEE International Conference on Data Mining, Orlando, FL, USA, November 28 - Dec. 1*, pages 1287–1292. IEEE, 2022.
- [Zhou *et al.*, 2020] Kaixiong Zhou, Qingquan Song, Xiao Huang, Daochen Zha, Na Zou, and Xia Hu. Multi-channel graph neural networks. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, pages 1352–1358. ijcai.org, 2020.