
SoGCN: Second-Order Graph Convolutional Networks

Anonymous Authors¹

Abstract

Graph Convolutional Networks (GCN) with multi-hop aggregation is more expressive than one-hop GCN but suffers from higher model complexity. Finding the shortest aggregation range that achieves comparable expressiveness and minimizes this side effect remains an open question. We answer this question by showing that multi-layer second-order graph convolution (SoGC) is sufficient to attain the ability of expressing polynomial spectral filters with arbitrary coefficients. Compared to models with one-hop aggregation, multi-hop propagation, and jump connections, SoGC possesses filter representational completeness while being lightweight, efficient, and easy to implement. Thereby, we suggest that SoGC is a simple design capable of forming the basic building block of GCNs, playing the same role as 3×3 kernels in CNNs. We build our Second-Order Graph Convolutional Networks (SoGCN) with SoGC and design a synthetic dataset to verify its filter fitting capability to validate these points. For real-world tasks, we present the state-of-the-art performance of SoGCN on the benchmark of node classification, graph classification, and graph regression datasets.

1. Introduction

Graph Convolutional Networks (GCNs) has gained popularity in recent years. Researchers have shown that non-localized multi-hop GCNs (Liao et al., 2019; Luan et al., 2019; Abu-El-Haija et al., 2019) have better performance than localized one-hop GCNs (Defferrard et al., 2016; Kipf & Welling, 2017; Wu et al., 2019). However, in Convolutional Neural Networks (CNNs), the localized 3×3 kernels’ expressiveness has been shown in image recognition both experimentally (Simonyan & Zisserman, 2014) and theo-

¹Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Kernel	Expressiveness	Localized	Complexity
Vanilla GCN	Very Low	✓	$O(ms)$
GIN	Medium	✓	$O(ms)$
Multi-hop	Full	×	$O(Kms)$
SoGCN (Ours)	Full	✓	$O(ms)$

Table 1. Comparison of different GC kernels in terms of expressiveness, localization, and time complexity. In this table, m represents the number of neighborhoods around a graph node, s is the dimension of input features, and K denotes the aggregation length of multi-hop GCs. We compute the time complexity with respect to the method of Abu-El-Haija et al. (2019).

retically (Zhou, 2020). These contradictory observations motivate us to search for a maximally localized Graph Convolution (GC) kernel with guaranteed feature expressiveness.

Most existing GCN layers adopt localized graph convolution based on one-hop aggregation scheme as the basic building block (Kipf & Welling, 2017; Hamilton et al., 2017; Xu et al., 2019). The effectiveness of these one-hop models is based on the intuition that a richer class of convolutional functions can be recovered by stacking multiple one-hop layers. However, extensive works (Li et al., 2018; Oono & Suzuki, 2019; Cai & Wang, 2020) have shown performance limitations of such design, which indicates this hypothesis may not hold. Liao et al. (2019); Luan et al. (2019); Abu-El-Haija et al. (2019) observed that multi-hop aggregation run in each layer could lead to significant improvement in prediction accuracy. However, a longer-range aggregator introduces extra hyperparameters and higher computational cost. This design also contradicts the compositionality principle of deep learning that neural networks benefit from deep connections and localized kernels (LeCun et al., 2015).

Recent studies point out that one-hop GCNs suffer from filter incompleteness (Hoang & Maehara, 2019). By relating low-pass filtering on the graph spectrum with over-smoothing, one can show one-hop filtering could lead to performance limitations. One natural solution of adding a more complex graph kernel, such as multi-hop connections, seems to work in practical settings (Abu-El-Haija et al., 2019). The question is: “what is the simplest graph kernel with the full expressive power of the graph convolution?” We show that with a second-order graph kernel of “two-hop” connection, we could approximate any complicated graph

relationships. Intuitively, it means we should extract a contrast between graph nodes that almost-connected (via their neighbor) vs. directly-connected.

We construct the two-hop graph kernel with second-order polynomials in an adjacency matrix and call it the Second-Order GC (SoGC). We show this Second-Order GC (SoGC) is the “sweet spot” balancing localization and fitting capability. To justify our conclusion, we introduce a Layer Spanning Space (LSS) framework to quantify the filter representation power of multi-layer GCs. Our LSS works by mapping GC filters’ composition with arbitrary coefficients to polynomial multiplication (Section 3.1).

Under this LSS framework, we can show that *SoGCs can approximate any linear GCNs in channel-wise filtering* (Theorem 1). Vanilla GCN and GIN (first-order polynomials in adjacency matrix) cannot represent all polynomial filters in general; multi-hop GCs (higher-order polynomials in adjacency matrix) do not contribute more expressiveness (Section 3.2). In this sense, SoGC is the most localized GC kernel with the full representation power.

To validate our theory, we build our Second-Order Graph Convolutional Networks (SoGCN) by layering up SoGC layers (Section 3.3). We reproduce our theoretical results on a synthetic datasets for filtering power testing (Section 4.1). On the public benchmark datasets (Dwivedi et al., 2020), SoGCN using simple graph topological features consistently boosts the performance of our baseline model (i.e., vanilla GCN) comparable to the state-of-the-art GNN models (with more complex attention/gating mechanisms). We also verify that our SoGCN fits extensive real-world tasks, including network node classification, super-pixel graph classification, and molecule graph regression (Section 4.3).

To our best knowledge, this work is the first study that identifies the distinctive competence of the two-hop neighborhood in the context of expressing a polynomial filter with arbitrary coefficients. Our SoGC is a special but non-trivial case of polynomial approximated graph filters (Defferrard et al., 2016). Kipf & Welling (2017) conducted an ablation study with GC kernels of different orders but missed the effectiveness of the second-order relationships. The work of Abu-El-Haija et al. (2019) talked about multi-hop graph kernels; however, they did not identify the critical importance of the two-hop form. In contrast, we clarify the prominence of SoGCs in theories and experiments.

Our research on GCN using pure topologically relationship is orthogonal to those using geometric relations (Monti et al., 2017; Fey et al., 2018; Pei et al., 2020), or those with expressive edge features (Li et al., 2016; Gilmer et al., 2017), and hyper-edges (Morris et al., 2019; Maron et al., 2018; 2019). It is also independent with graph sampling procedures (Rong et al., 2019; Hamilton et al., 2017).

2. Related Work

Spectral GCNs. Graph convolution is defined as element-wise multiplication on graph spectrum (Hammond et al., 2011). Bruna et al. (2014) first proposed spectral GCN with respect to this definition. ChebyNet (Defferrard et al., 2016) approximates graph filters using Chebyshev polynomials. Vanilla GCN (Kipf & Welling, 2017; Wu et al., 2019) further reduces the GC layer to a degree-one polynomial with lumping of first-order and constant terms. GIN (Xu et al., 2019) disentangles the effect of self-connection and pairwise neighboring connections by adding a separate mapping for central nodes. However, these simplifications causes performance limitations (Oono & Suzuki, 2019; Cai & Wang, 2020). APPNP (Klicpera et al., 2019) uses Personalized PageRank to derive a fixed polynomial filter. Bianchi et al. (2019) proposes a multi-branch GCN architecture to simulate ARMA graph filters. GCNII (Ming Chen et al., 2020) incorporates identity mapping and initial mapping to relieve over-smoothing problem and deepen GCNs. However, these models are not easy to implement and introduce additional hyper-parameters.

Multi-Hop GCNs. To exploit multi-hop information, Liao et al. (2019) proposes to use Lanczos algorithm to construct low rank approximations of the graph Laplacian for graph convolution. Luan et al. (2019) devises two architectures Snowball GCN and Truncated Krylov GCN to capture neighborhoods at various distances. To simulate neighborhood delta functions, Abu-El-Haija et al. (2019) repeat mixing multi-hop features to identify more topological information. JKNet (Xu et al., 2018) combines all feature activation of previous layers to learn adaptive and structure-aware representations of different graph substructures. These models exhibit the strength of multi-hop GCNs over one-hop GCNs while leaving the propagation length as a hyper-parameter. In the meanwhile, long-range aggregation in each layer causes higher complexity (Table 1).

Expressiveness of GCNs. Most of the works on GCN’s expressiveness are restricted to the over-smoothing problem: Li et al. (2018) first poses the over-smoothing problem; Hoang & Maehara (2019) indicates GCNs are no more than low-pass filters; Luan et al. (2019); Oono & Suzuki (2019) demonstrate the asymptotic behavior of feature activation to a subspace; Cai & Wang (2020) examines the decreasing Dirichlet energy. These analytic frameworks do not provide constructive improvements on building more powerful GCN. Ming Chen et al. (2020) first proposed to assess GCN’s overall expressiveness by tackling the ability of expressing polynomial filters with arbitrary coefficients. But their theory is only applicable for transductive learning on a single graph, and do not upper bound the degree of the graph polynomial filters.




Figure 1. Vertex domain interpretations of vanilla GC, SoGC, and Multi-Hop GC. Denote A the first-hop aggregator, A^2 the second-hop aggregator, and A^K the K -th hop aggregator. Nodes in the same colored ring share the same weights. (a) Vanilla GC only aggregates information from the first-hop neighbor nodes. (b) SoGC incorporates additional information from the second-hop (almost-connected) neighborhood. (c) Multi-hop GC simply repeats mixing information from every neighborhood within K hops.

3. Second-Order Graph Convolution

We begin by introducing our notation. We are interested in learning on a finite graph set $\mathcal{G} = \{G_1, \dots, G_{|\mathcal{G}|}\}$. Assume each graph $G \in \mathcal{G}$ is simple and undirected, associated with a finite vertex set $\mathcal{V}(G)$, an edge set $\mathcal{E}(G) \subseteq \mathcal{V}(G) \times \mathcal{V}(G)$, and a symmetric normalized adjacency matrix $\mathbf{A}(G)$ (Chung & Graham, 1997; Shi & Malik, 2000). Without loss of generality and for simplicity, $|\mathcal{V}(G)| = N$ for every $G \in \mathcal{G}$. We denote single-channel signals supported in graph $G \in \mathcal{G}$ as $\mathbf{x} \in \mathbb{R}^N$, a vectorization of function $\mathcal{V}(G) \rightarrow \mathbb{R}$.

Graph Convolution (GC) is defined as Linear Shift-Invariant (LSI) operators to adjacency matrices (Sandryhaila & Moura, 2013). This property enables GC to extract features regardless of where the local structure falls. A single-channel GC can be written as a mapping $f : \mathcal{G} \times \mathbb{R}^N \rightarrow \mathbb{R}^N$. According to Defferrard et al. (2016), a GC can be approximated by a polynomial in adjacency matrix (Figure 1c)¹:

$$f_{\boldsymbol{\theta}}(G, \mathbf{x}) = \sum_{k=0}^K \theta_k \mathbf{A}(G)^k \mathbf{x}, \quad (1)$$

where $\boldsymbol{\theta} = [\theta_0 \ \dots \ \theta_K]^T \in \mathbb{R}^{K+1}$ represents the kernel weights. The Equation 1 indicates that graph convolution can be interpreted as a linear combination of features aggregated by $\mathbf{A}(G)^k$. Thereby, the hyperparameter K can reflect the localization of a GC kernel.

Previous work of Kipf & Welling (2017) simplified the Equation 1 to a one-hop kernel, so-called vanilla GC (Figure 1a). We formulate vanilla GC as below:

$$f_1(G, \mathbf{x}) = \theta (\mathbf{A}(G) + \mathbf{I}) \mathbf{x}. \quad (2)$$

We are interested in the overall graph convolution networks' representation power of expressing a polynomial filter (cf. Equation 1) with arbitrary degrees and coefficients. At first glance, one-hop GC (cf. Equation 2) can approximate any

¹We can replace the Laplacian matrix L in Defferrard et al. (2016) with the normalized adjacency matrix A since $L = I - A$.

high-order GC kernels by stacking multiple layers. However, that is not the case (See formal arguments in Section 3.2). In contrast, when plugging the second-order term into vanilla GC, we will show that this approximation ability can be attained (See formal arguments in Theorem 1). We name this improved design Second-Order Graph Convolution (SoGC), as it can be written as the second-order polynomial in adjacency matrix:

$$f_2(G, \mathbf{x}) = (\theta_2 \mathbf{A}(G)^2 + \theta_1 \mathbf{A}(G) + \theta_0 \mathbf{I}) \mathbf{x}. \quad (3)$$

We illustrate its vertex-domain interpretation in Figure 1b. The critical insight is that graph filter approximation can be viewed as a polynomial factorization problem. It is known that *any univariate polynomial can be factorized into sub-polynomials of degree two*. Based on this fact, we show by stacking enough SoGCs (and varying their parameters) can achieve decomposition of any polynomial filters.

3.1. Representation Power Justification

To be more precise, we introduce our Layer Spanning Space (LSS) framework in this subsection and mathematically prove that arbitrary GC kernel can be decomposed into finite many SoGC kernels.

First, to illustrate overall graph filter space and filters expressed by a single layer, we define a *graph filter space*, in which every polynomial filter has degree no more than K :

Definition 1. (*Graph filter space*) Suppose the parameter space is \mathbb{R} . The Linear Shift-Invariant (LSI) graph filter space² of degree $K > 0$ with respect to a finite graph set \mathcal{G} is defined as $\mathcal{F}_K = \{f_{\boldsymbol{\theta}} : \mathcal{G} \times \mathbb{R}^N \rightarrow \mathbb{R}^N, \forall \boldsymbol{\theta} \in \mathbb{R}^{K+1}\}$, where $f_{\boldsymbol{\theta}}$ follows the definition in Equation 1.

We further provide Definition 2 and Lemma 1 to discuss the upper limit of \mathcal{F}_K 's dimension.

Definition 2. (*Spectrum capacity*) Let spectrum set $\mathcal{S}(\mathcal{G}) = \{\lambda : \lambda \in \mathcal{S}(\mathbf{A}(G)), \forall G \in \mathcal{G}\}$, where $\mathcal{S}(\mathbf{A})$ denotes the

²See Supp. A for a remark on how \mathcal{F}_K relates with the term "Linear Shift-Invariant".




Figure 2. Visualizing output activation in graph spectrum domain for vanilla GCN, SoGCN, and GRU variants. The test is conducted on a graph from the ZINC dataset. The spectrum is defined as a projection of activation functions on the graph eigenvectors. SoGCN preserved higher-order spectrum, while vanilla GCN shows over-smoothing. See Supp. F for more visualizations on the ZINC dataset.

eigenvalues of \mathbf{A} . Spectrum capacity $\Gamma = |\mathcal{S}(\mathcal{G})|$ is the cardinality of all distinct graph eigenvalues. In particular, $\Gamma = (N - 1)|\mathcal{G}|$ if every graph adjacency matrix has no common eigenvalues other than 1.

Lemma 1. Filter space \mathcal{F}_K with degree $K > 0$ has dimension $\min\{K + 1, \Gamma\}$ as a vector space.

Lemma 1 follows from Theorem 3 of Sandryhaila & Moura (2013). See the complete proof in Supp. C. \square

According to Lemma 1, one can define an ambient filter space \mathcal{A} of degree $\Gamma - 1$ (i.e., $\mathcal{F}_{\Gamma-1}$). Suppose a GCN consists of K -hop convolutional kernels $f^{(l)}, l = 1, \dots, L$, where the superscript indicates the layer number, and L denotes the network depth. We can consider each layer is sampled from \mathcal{F}_K . We intend to justify this GCN’s filter representation power via its Layer Spanning Space (LSS):

$$\mathcal{F}_K^L = \left\{ F : F = f^{(1)} \circ \dots \circ f^{(L)}, \forall f^{(l)} \in \mathcal{F}_K \right\}, \quad (4)$$

where the whole LSS is constructed by varying parameters of $f^{(l)}$ over \mathbb{R}^{K+1} . When $\mathcal{A} \subseteq \mathcal{F}_K^L$, the LSS covers the entire ambient space, then we say the GCN composed of L GC kernels in \mathcal{F}_K has full filter representation power.

As \mathcal{F}_K is a function space, \mathcal{F}_K^L can be analytically tricky. To investigate this space, we define a mapping $\tau : \mathcal{F}_K \rightarrow \mathbb{R}_K[x]$, where $\mathbb{R}_K[x]$ denotes a polynomial vector space of degree at most K :

$$\tau : \sum_{k=0}^K \theta_k \mathbf{A}(G)^k \mapsto \sum_{k=0}^K \theta_k x^k, \quad (5)$$

We provide Lemma 2 to reveal a good property of τ :

Lemma 2. τ is a ring isomorphism when $K \leq \Gamma - 1$.

The proof can be found in Supp. D. \square

This ring isomorphism signifies that the composition of filters in \mathcal{F}_K is identical to polynomial multiplication. Therefore, one can study the LSS through the polynomials that can be factorized into the corresponding sub-polynomials. For example, Equation 4 is identical to:

$$\mathcal{F}_K^L \simeq \left\{ p(x) : p(x) = \prod_{l=1}^L \sum_{k=0}^K \theta_k^{(l)} x^k, \theta_k^{(l)} \in \mathbb{R} \right\}. \quad (6)$$

In the rest of this subsection, we will show that $\lceil \frac{\Gamma-1}{2} \rceil$ -layer SoGCs (i.e., \mathcal{F}_2) can attain full representation power. That is, \mathcal{F}_2^L covers the whole ambient filter space \mathcal{A} when L is as large as $\lceil (\Gamma - 1)/2 \rceil$. We summarize a formal argument in the following theorem:

Theorem 1. For any $f \in \mathcal{A}$, there exists $f_2^{(l)} \in \mathcal{F}_2$ with coefficients $\theta_0^{(l)}, \theta_1^{(l)}, \theta_2^{(l)} \in \mathbb{R}, l = 1, \dots, L$ such that $f = f_2^{(L)} \circ \dots \circ f_2^{(1)}$ where $L \leq \lceil (\Gamma - 1)/2 \rceil$.

Proof. Proving Theorem 1 requires a fundamental polynomial factorization theorem, rephrased as below:

Lemma 3. (Fundamental theorem of algebra) Over the field of reals, the degree of an irreducible non-trivial univariate polynomial is either one or two.

For any $f \in \mathcal{A}$, apply $\tau : \mathcal{A} \rightarrow \mathbb{R}_{\Gamma-1}[x]$ (cf. Equation 5) to map kernel f to polynomial $h(x)$. By Lemma 1, $\deg h(x) \leq \Gamma - 1$. By Lemma 3, factorize $h(x)$ into series of polynomials with the degree at most two, and then merge first-order polynomials into second-order ones until one single or no first-order sub-polynomial remains. As a consequence, $h(x)$ can be written as $h(x) = \prod_{l=1}^L h_l(x)$, where $L = \lceil \deg h(x)/2 \rceil$. If $\deg h(x)$ is even, $\deg h_l(x) = 2$ for

every $l = 1, \dots, D$. Otherwise, except for at most one $h_l(x)$ whose degree is one, all terms have degree two.

The last step is to apply the inverse of morphism $\tau^{-1} : \mathbb{R}_{\Gamma^{-1}}[x] \rightarrow \mathcal{A}$ formulated as below:

$$\tau^{-1} : \sum_{k=0}^K \theta_k x^k \mapsto \sum_{k=0}^K \theta_k \mathbf{A}(G)^k.$$

Since τ^{-1} is also a ring isomorphism, we have:

$$\begin{aligned} \tau^{-1}(h(x)) &= \tau^{-1}(h_1(x) \cdots h_L(x)) \\ &= \tau^{-1}(h_1(x)) \circ \cdots \circ \tau^{-1}(h_L(x)) \\ &= f^{(1)} \circ \cdots \circ f^{(L)} \end{aligned}$$

where $f^{(l)} \in \mathcal{F}_2, l = 1, \dots, L$ by definition, which implies $f = \tau^{-1}(\tau(f)) = f^{(L)} \circ \cdots \circ f^{(1)}$. \square

Theorem 1 can be regarded as the universal approximation theorem of linear GCNs. Although nonlinear activation is not considered within our theoretical framework, we make reasonable hypothesis that achieving linear filter completeness can also boost GCNs with nonlinearity (Ming Chen et al., 2020). See our experiments in Section 4.3.

Theorem 1 implies that multi-layer SoGC can implement arbitrary filtering effects and extract features at any positions on the spectrum (Figure 2). Theorem 1 also coincides with Dehmamy et al. (2019) on how GCNs built on SoGC kernels could utilize depth to raise fitting accuracy (Figure 3).

3.2. Compared with Other Graph Convolution

In this subsection, we will show that vanilla GCN (and GIN) does not attain full expressiveness in terms of filter representation. We will also contrast our SoGC to multi-hop GCs (i.e., higher-order GCs in our terminology) to further reveal the prominence of SoGC. A brief comparison is summarized in Table 1.

Vanilla vs. second-order. Vanilla GCN (Kipf & Welling, 2017) is a typical one-hop GCN with lumbing of of graph node self-connection and pairwise neighboring connection. Compared with SoGC, vanilla GC is more localized and computationally cheaper. However, this design has huge performance limitations (Hoang & Maehara, 2019; Wu et al., 2019; Li et al., 2018; Oono & Suzuki, 2019; Cai & Wang, 2020). We illustrate this issue in terms of filter approximation power based on the LSS framework.

Suppose a GCN stacks L GC layers $f_1^{(l)}(G, \mathbf{x}) \in \mathcal{F}_1$, apply mapping τ to its spanned LSS, the isomorphic polynomial space is:

$$\mathcal{F}_1^L \simeq \left\{ p : p(x) = \Theta \sum_{l=0}^L \binom{l}{L} x^l, \Theta \in \mathbb{R} \right\}, \quad (7)$$

where $\Theta = \theta^{(L)} \cdots \theta^{(1)}$. According to Equation 7, one can see *no matter how large L is or how a optimizer tunes the parameters $\theta^{(l)}$, $\dim \mathcal{F}_1^L = 1$, which implies \mathcal{F}_1^L degenerates to a negligible subspace inside \mathcal{A}* . GIN (Xu et al., 2019) disentangles the the weights for neighborhoods and central nodes. We can write this GC layer as $f_{GIN}^{(l)}(G, \mathbf{x}) = (\theta_1 \mathbf{A}(G) + \theta_0 \mathbf{I})\mathbf{x}$. The LSS of GIN is isomorphic to the polynomial space:

$$\mathcal{F}_{GIN}^L \simeq \left\{ p : p(x) = \prod_{l=1}^L \left(\theta_1^{(l)} x + \theta_0^{(l)} \right), \theta_0^{(l)}, \theta_1^{(l)} \in \mathbb{R} \right\}. \quad (8)$$

This polynomial space represents all polynomials that can split over the real domain. However, $\mathcal{F}_{GIN}^L \subsetneq \mathcal{A}$ since *not all polynomials can be factorized into first-order polynomials*. The expectation number of real roots of a K -degree polynomial with zero-mean random coefficients is $\mathbb{E}(N) \sim (2/\pi) \log K$ (Ibragimov & Maslova, 1971). When the ambient dimension goes larger, all-real-root polynomials only occupy a small proportion in the ambient space (Li, 2011), which indicates GIN does not have full expressiveness in terms of filter representation either.

Higher-order vs. second-order. Higher-order GCs refer to those polynomial filters with degree larger than three (i.e., $\mathcal{F}_K, K \geq 3$). They can model multi-hop GCNs such as Luan et al. (2019); Liao et al. (2019); Abu-El-Haija et al. (2019). Compared to SoGCs, higher-order GCs have equivalent expressive power, since they can be reduced to SoGCs. However, we point out four limitations of adopting higher-order kernels: 1) From our polynomial factorization perspective, fitting graph filters using higher-order GC requires coefficient sparsity, which brings about learning difficulty. Abu-El-Haija et al. (2019) overcomes this problem by adding lasso regularization and extra training procedures. Adopting SoGC can avoid these troubles since decomposition into second-order polynomials results in at most one zero coefficient (See Section 3.1). 2) Eigenvalues of graph adjacency matrices diminish when powered. This leads to a decreasing numerical rank of $\mathbf{A}(G)^k$ and makes aggregating larger-scale information ineffective. SoGCs can alleviate this problem by preventing higher-order powering operations. 3) Higher-order GC lacks nonlinearity. SoGCN can bring a better balance between the expressive power of low-level layers and nonlinearity among them. 4) Multi-hop aggregation consumes higher computational resources (See Table 1). In contrast, SoGC matches the time complexity of vanilla GCN by fixing the kernel size to two.

3.3. Implementation of Second-Order Graph Convolutional Networks

In this section, we introduce other building blocks to establish our Second-Order Graph Convolutional Networks

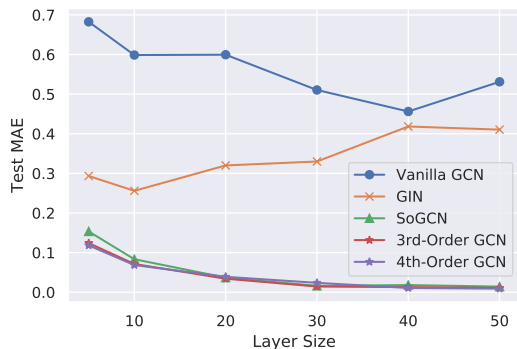


Figure 3. Relations between test MAE and layer size. Experiments are conducted on synthetic *Band-Pass* dataset. Each model has 16 channels per hidden layer with varying layer size.

(SoGCN) following the latest trends of GCN. We promote SoGC to its multi-channel version analogous to Kipf & Welling (2017). Then we prepend a feature embedding layer, cascade multiple SoGC layers, and append a readout module. Suppose the network input is $\mathbf{X} \in \mathbb{R}^{N \times D}$ supported in graph $G \in \mathcal{G}$, denote the output of l -th layer as $\mathbf{X}^{(l)} \in \mathbb{R}^{N \times E}$, the final node-level output as $\mathbf{Y} \in \mathbb{R}^{N \times F}$, or graph-level output as $\mathbf{Y} \in \mathbb{R}^E$, we formulate our novel deep GCN built with SoGC (cf. Equation 3) as follows:

$$\mathbf{X}^{(0)} = \phi(\mathbf{X}; \Phi), \quad (9)$$

$$\mathbf{X}^{(l)} = \sigma\left(f_2^{(l)}\left(\mathbf{X}^{(l-1)}; \Theta^{(l)}\right)\right), \quad (10)$$

$$\mathbf{Y} = \psi\left(\mathbf{X}^{(L)}; \Psi\right), \quad (11)$$

where $\Theta^{(l)} = \{\Theta_i^{(l)} \in \mathbb{R}^{E \times E}\}_{i=0,1,2}$ are trainable weights for linear filters $f_2^{(l)} \in \mathcal{F}_2$; $\phi: \mathbb{R}^{N \times D} \rightarrow \mathbb{R}^{N \times E}$ is an equivariant embedder (Maron et al., 2018) with parameters Φ ; $\sigma: \mathbb{R}^{N \times E} \rightarrow \mathbb{R}^{N \times E}$ is an activation function. For node-level readout, $\psi: \mathbb{R}^{N \times E} \rightarrow \mathbb{R}^{N \times F}$ can be a decoder (with parameters Ψ) or an output activation (e.g., softmax) in place of the prior layer. For graph-level output, $\psi: \mathbb{R}^{N \times E} \rightarrow \mathbb{R}^E$ should be an invariant readout function (Maron et al., 2018), e.g., channel-wise sum, mean or max. In practice, we adopt ReLU as nonlinear activation (i.e., $\sigma = \text{ReLU}$), a multi-layer perceptron (MLP) as the embedding function ϕ , another MLP for node regression readout, and sum (Xu et al., 2019) for graph classification readout.

We also provide a variant of SoGCN integrated with Gated Recurrent Unit (GRU) (Girault et al., 2015), termed SoGCN-GRU. According to Cho et al. (2014), GRU can utilize gate mechanism to preserve and forget information. We hypothesize that a GRU can be trained to remove redundant signals and retain lost features on the spectrum. Similar to Li et al. (2016); Gilmer et al. (2017), we append a shared GRU module after each GC layer, which takes the signal

Table 2. The performance of graph node signal regression with *High-Pass*, *Low-Pass*, and *Band-Pass* filters (over graph spectral space) as learning target. Each model has 16 GC layers and 16 channels of hidden layers.

Model	#Param	Test MAE		
		High-Pass	Low-Pass	Band-Pass
Vanilla	4611	0.308	0.317	0.559
GIN	4627	0.344	0.096	0.274
SoGCN	12323	0.021	0.023	0.050
3rd-Order	16179	0.021	0.022	0.045
4th-Order	20035	0.021	0.022	0.049

before the GC layer as the hidden state, after the GC layer as the current input. We show by our experiment that GRU can facilitate SoGCN in avoiding noises and enhancing features on the spectrum (Figure 2). Our empirical study in Table 5 also indicates the effectiveness of GRU for spectral GCNs is general. Hence, we suggest including this recurrent module as another basic building block of our SoGCNs.

4. Experiments

4.1. Synthetic Graph Spectrum Dataset for Node Regression

To validate the expressiveness of SoGCN, and its power to fit arbitrary graph filters, we build a Synthetic Graph Spectrum (SGS) dataset for the node signal filtering regression task. We construct SGS dataset with random graphs. The learning task is to simulate three types of hand-crafted filtering functions: high-pass, low-pass, and band-pass on the graph spectrum (defined over the graph eigenvectors). There are 1k training graphs, 1k validation graphs, and 2k testing graphs for each filtering function. Each graph is undirected and comprises 80 to 120 nodes. Supp. E covers more details of our SGS dataset. We choose Mean Absolute Error (MAE) as evaluation metric.

Experimental Setup. We compare SoGCN with vanilla GCN (Kipf & Welling, 2017), GIN (Xu et al., 2019), and higher-order GCNs on the synthetic dataset. To evaluate each model’s expressiveness purely on the GC kernel design, we remove ReLU activations for all tested models. We adopt the Adam optimizer (Kingma & Ba, 2015) in our training process, with a batch size of 128. The learning rate begins with 0.01 and decays by half once the validation loss stagnates for more than 10 training epochs.

Results and Discussion. Table 2 summarizes the quantitative comparisons. SoGCN achieves the superior performance on all of the 3 tasks outperforming vanilla GCN and GIN, which implies that SoGC graph convolutional kernel does benefit from explicit disentangling of the second-hop

Table 3. Results and comparison with other GNN models on ZINC, CIFAR10, MNIST, CLUSTER and PATTERN datasets. For ZINC dataset, the parameter budget is set to 500k. For CIFAR10, MNIST, CLUSTER and PATTERN datasets, the parameter budget is set to 100k. **Red**: the best model, **Green**: good models.

Model	Test MAE \pm s.d.	Test ACC \pm s.d. (%)			
	ZINC	MNIST	CIFAR10	CLUSTER	PATTERN
Vanilla GCN	0.367 \pm 0.011	90.705 \pm 0.218	55.710 \pm 0.381	53.445 \pm 2.029	63.880 \pm 0.074
Vanilla GCN + GRU	0.295 \pm 0.005	96.020 \pm 0.090	61.332 \pm 0.849	57.932 \pm 0.168	70.194 \pm 0.216
GAT	0.384 \pm 0.007	95.535 \pm 0.205	64.223 \pm 0.455	57.732 \pm 0.323	75.824 \pm 1.823
MoNet	0.292 \pm 0.006	90.805 \pm 0.032	65.911 \pm 2.515	58.064 \pm 0.131	85.482 \pm 0.037
GraphSage	0.398 \pm 0.002	97.312\pm0.097	65.767 \pm 0.308	50.454 \pm 0.145	50.516 \pm 0.001
GIN	0.387 \pm 0.015	96.485 \pm 0.252	55.255 \pm 1.527	58.384 \pm 0.236	85.590\pm0.011
GatedGCN	0.350 \pm 0.020	97.340\pm0.143	67.312\pm0.311	60.404 \pm 0.419	84.480 \pm 0.122
3WLGNN	0.407 \pm 0.028 ³	95.075 \pm 0.961	59.175 \pm 1.593	57.130 \pm 6.539	85.661\pm0.353
SoGCN	0.238\pm0.017	96.785\pm0.113	66.338\pm0.155	68.167\pm1.164	85.735\pm0.037
SoGCN-GRU	0.201\pm0.006	97.729\pm0.159	68.208\pm0.271	67.994\pm2.619	85.711\pm0.047

neighborhoods. Our results also show that higher-order (3rd-order and 4th-order) GCNs do not improve the performance further, even though they incorporate much more parameters. SoGCN is more expressive and does a better trade-off between performance and model size.

Figure 3 plots MAE results as we vary the depth of GC layers for each graph kernel type. Vanilla GCN and GIN can not benefit from depth while SoGC and higher-order GCs can leverage depth to span larger LSS, contributing to the remarkable filtering results. SoGC and higher-order GCs have very close performance after increasing the layer size, which suggests higher-order GCs do not obtain more expressiveness than SoGC.

Table 4. The performance of graph multi-label classification on ogb-protein dataset. We compare each model condering the following dimensions: the number of parameters, training time (in seconds) per epoch (ep.), and final test ROC-AUC (%).

Model	ogb-protein		
	#Param	Time / Ep.	ROC-AUC \pm s.d.
Vanilla GCN	96880	3.47 \pm 0.40	72.16 \pm 0.55
GIN	128512	4.33 \pm 0.27	76.77 \pm 0.20
GCNII	227696	4.96 \pm 0.29	74.79 \pm 1.17
GCNII*	424304	5.09 \pm 0.17	72.50 \pm 2.49
APPNP	96880	6.56 \pm 0.37	65.37 \pm 1.15
GraphSage	193136	6.51 \pm 0.13	77.53 \pm 0.30
SoGCN	192512	4.88 \pm 0.36	79.28 \pm 0.47
4th-Order GCN	320512	8.89 \pm 0.82	78.95 \pm 0.57
6th-Order GCN	448512	9.76 \pm 0.64	78.61 \pm 0.42

³ This is the result of 3WLGNN with 100k parameters. The test MAE of 3WLGNN with 500k parameters is increased to 0.427 \pm 0.011.

4.2. OGB Benchmarks

We choose Open Graph Benchmark (OGB) (Hu et al., 2020) to compare our SoGC with other GCNs in terms of the parameter numbers, train time per epoch, and test ROC-AUC. We only demonstrate the results for predicting presence of protein functions (multi-label graph classification). We refer interested reader to Supp. G for more results on OGB.

Experiment Setup The chosen models mainly include spectral-domain models: vanilla GCN, GIN, APPNP (Klicpera et al., 2019), GCNII (Ming Chen et al., 2020), our SoGCN, and two high-order GCNs. We also obtain the performance of GraphSage, the vertex-domain GNN baseline, for a reference. We build GIN, GCNII, GraphSage, and APPNP based on official implementations in PyTorch Geometric. Every model consists of three layers and shares the same node embedder and readout modules. We borrow the method of Dwivedi et al. (2020) to compute the number of paramters. We run an exclusive training program on an Nvidia Quadro P6000 GPU to test the training time per epoch for each model. We follow the same training and evaluation procedures on the OGB benchmarks to ensure fair comparisons. We train each model until convergence (~1k epochs for Vanilla GCN, GIN, GraphSage, and ~3k epochs for SoGCN, higher-order GCNs, APPNP, and GCNII).

Results and Discussion. Table 4 demonstrates the ROC-AUC score for each model on ogb-protein dataset. Our SoGCN achieves the best performance among all presented GCNs but its parameter number and time complexity is only slightly higher than GIN (consistent with Table 1). SoGC is more expressive than other existing graph filters (such as APPNP and GCNII), and also outperforms message-passing

Table 5. Results of ablation study on ZINC, MNIST and CIFAR10 datasets. Vanilla GCN is the comparison baseline and the number in the (\uparrow ·) and (\downarrow ·) represents the performance gain compared with the baseline.

Model	Test MAE \pm s.d.		Test ACC \pm s.d. (%)	
	ZINC	MNIST	CIFAR10	
Vanilla GCN	0.367 \pm 0.011 (Baseline)	90.705 \pm 0.218 (Baseline)	55.710 \pm 0.381 (Baseline)	
SoGCN	0.238 \pm 0.017 (\downarrow 0.129)	96.785 \pm 0.113 (\uparrow 6.080)	66.338 \pm 0.155 (\uparrow 10.628)	
4th-Order GCN	0.243 \pm 0.009 (\downarrow 0.124)	96.167 \pm 0.198 (\uparrow 5.462)	64.230 \pm 0.212 (\uparrow 8.520)	
6th-Order GCN	0.261 \pm 0.014 (\downarrow 0.106)	96.292 \pm 0.134 (\uparrow 5.587)	63.687 \pm 0.151 (\uparrow 7.977)	
Vanilla GCN + GRU	0.295 \pm 0.005 (\downarrow 0.072)	96.020 \pm 0.090 (\uparrow 5.315)	61.332 \pm 0.381 (\uparrow 5.622)	
SoGCN + GRU	0.201 \pm 0.006 (\downarrow 0.166)	97.729 \pm 0.159 (\uparrow 7.024)	68.208 \pm 0.271 (\uparrow 12.498)	
4th-Order GCN + GRU	0.204 \pm 0.004 (\downarrow 0.163)	97.304 \pm 0.296 (\uparrow 6.599)	64.697 \pm 0.341 (\uparrow 8.987)	
6th-Order GCN + GRU	0.218 \pm 0.005 (\downarrow 0.149)	97.325 \pm 0.218 (\uparrow 6.620)	64.523 \pm 0.251 (\uparrow 8.813)	

GNN baseline GraphSage. Compared with higher-order (4th and 6th) GCNs, the ROC-AUC score of SoGCN surpasses all of them while reducing model complexity significantly.

4.3. GNN Benchmarks

We follow the benchmarks outlined in Dwivedi et al. (2020) for evaluating GNNs on several datasets across a variety of artificial and real-world tasks. We choose to evaluate our SoGCN on a real-world chemistry dataset (ZINC molecules) for the graph regression task, two semi-artificial computer vision datasets (CIFAR10 and MNIST superpixels) for the graph classification task, and two artificial social network datasets (CLUSTER and PATTERN) for node classification.

Experimental Setup. We compare our proposed SoGCN and SoGCN-GRU with state-of-the-art GNNs: vanilla GCN, GIN, GraphSage, GAT (Veličković et al., 2018), MoNet (Monti et al., 2017), GatedGCN (Bresson & Laurent, 2017) and 3WL-GNN (Maron et al., 2019). To ensure fair comparisons, we follow the same training and evaluation pipelines (including optimizer settings) and data splits of benchmarks. Furthermore, we adjust our model’s depth and width to ensure it satisfies parameter budgets as specified in the benchmark. Note that we do not use any geometrical information to encode rich graph edge relationship, as in models such as GatedGCN-E-PE. We only employ graph connectivity information for all tested models.

Results and Discussion. Table 3 reports the benchmark results. Our model SoGCN makes small computational changes to GCN by adopting second-hop neighborhood, and it outperforms models with complicated message-passing mechanisms, such as GAT and GraphSage. With GRU module, SoGCN-GRU tops almost all state-of-the-art GNNs on the ZINC, MNIST and CIFAR10 datasets. In Figure 2, we visualize a spectrum of the last layer’s feature activation on ZINC dataset. One can see our SoGC can extract features on high-frequency bands and GRU can further sharpen these patterns. However, GRU does not lift accuracy on

CLUSTER and PATTERN datasets for node classification task. According to Li et al. (2018), that GRU suppresses low-frequency band results in the slight performance drop on the CLUSTER and PATTERN datasets.

Ablation Study. To contrast the performance gain produced by different aggregation ranges and GRU on the benchmarks, we evaluate vanilla GCN, SoGCN, 4th-Order GCN, 6th-Order GCN as well as their GRU variants on the ZINC, MNIST and CIFAR10 datasets. Table 5 presents the results of our ablation study, which are consistent with our observation on Section 4.1 and 4.2. As shown by our ablation study, adopting the second-hop aggregation makes huge performance gain (Vanilla GCN vs. SoGCN). However, high-order GCNs are not capable of boosting the performance further over SoGCN. On the contrary, higher-order GCs can even lead to the performance drop (4th-Order GCN vs. 6th-Order GCN vs. SoGCN). We also testify GRU’s effectiveness for each presented model. But the gain brought by GRU is not as large as adding second-hop aggregation. Figure 2 shows our SoGC can extract patterns on the spectrum alone. GRU plays a role of enhancing the features.

5. Conclusion

What should be the basic convolutional blocks for GCNs? To answer this, we seek the most localized graph convolution kernel (GC) with full expressiveness. We establish our LSS framework to assess GC layers of different aggregation ranges. We show the second-order graph convolutional filter, termed SoGC, possesses the full representation power than one-hop GCs. Hence, it becomes the efficient and simplest GC building blocks that we adopt to establish our SoGCN. Both synthetic and benchmark experiments exhibit the prominence of our theoretic design. We also make an empirical study on the GRU’s effects in spectral GCNs. Interesting directions for future work include analyzing two-hop aggregation schemes with message-passing GNNs and proving the universality of nonlinear GCNs.

References

- 440
441
442 Abu-El-Haija, S., Perozzi, B., Kapoor, A., Alipourfard, N.,
443 Lerman, K., Harutyunyan, H., Steeg, G. V., and Galstyan,
444 A. Mixhop: Higher-order graph convolutional architec-
445 tures via sparsified neighborhood mixing. In *ICML*, 2019.
- 446 Bianchi, F. M., Grattarola, D., Livi, L., and Alippi, C. Graph
447 neural networks with convolutional arma filters. In *CoRR*,
448 2019.
- 449
450 Bresson, X. and Laurent, T. Residual gated graph convnets.
451 *arXiv:1711.07553*, 2017.
- 452
453 Bruna, J., Zaremba, W., Szlam, A., and LeCun, Y. Spec-
454 tral networks and locally connected networks on graphs.
455 2014.
- 456 Cai, C. and Wang, Y. A note on over-smoothing for graph
457 neural networks. In *ICML*, 2020.
- 458
459 Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau,
460 D., Bougares, F., Schwenk, H., and Bengio, Y. Learn-
461 ing phrase representations using rnn encoder-decoder
462 for statistical machine translation. *arXiv preprint*
463 *arXiv:1406.1078*, 2014.
- 464
465 Chung, F. R. and Graham, F. C. *Spectral graph theory*.
466 American Mathematical Soc., 1997.
- 467
468 Defferrard, M., Bresson, X., and Vandergheynst, P. Con-
469 volutional neural networks on graphs with fast localized
470 spectral filtering. In *NeurIPS*, 2016.
- 471
472 Dehmamy, N., Barabási, A.-L., and Yu, R. Understanding
473 the representation power of graph neural networks in
474 learning graph topology. In *NeurIPS*, 2019.
- 475
476 Dwivedi, V. P., Joshi, C. K., Laurent, T., Bengio, Y.,
477 and Bresson, X. Benchmarking graph neural networks.
478 *arXiv:2003.00982*, 2020.
- 479
480 Fey, M., Eric Lenssen, J., Weichert, F., and Müller, H.
481 Splinecnn: Fast geometric deep learning with continuous
482 b-spline kernels. In *CVPR*, 2018.
- 483
484 Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and
485 Dahl, G. E. Neural message passing for quantum chem-
486 istry. 2017.
- 487
488 Girault, B., Gonçalves, P., and Fleury, É. Translation on
489 graphs: An isometric shift operator. *SPL*, 2015.
- 490
491 Hamilton, W., Ying, Z., and Leskovec, J. Inductive repre-
492 sentation learning on large graphs. In *NeurIPS*, 2017.
- 493
494 Hammond, D. K., Vandergheynst, P., and Gribonval, R.
Wavelets on graphs via spectral graph theory. *ACHA*,
2011.
- Hoang, N. and Maehara, T. Revisiting graph neural net-
works: All we have is low-pass filters. *arXiv:1905.09550*,
2019.
- Hu, W., Fey, M., Zitnik, M., Dong, Y., Ren, H., Liu, B.,
Catasta, M., and Leskovec, J. Open graph benchmark:
Datasets for machine learning on graphs. arXiv preprint
arXiv:2005.00687, 2020.
- Ibragimov, I. A. and Maslova, N. B. The mean number
of real zeros of random polynomials. i. coefficients with
zero mean. *Theory of Probability & Its Applications*, 16:
228–248, 1971.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic
optimization. 2015.
- Kipf, T. N. and Welling, M. Semi-supervised classification
with graph convolutional networks. In *ICLR*, 2017.
- Klicpera, J., Bojchevski, A., and Günnemann, S. Predict
then propagate: Graph neural networks meet personalized
pagerank. In *ICLR*, 2019.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *Nature*,
2015.
- Li, Q., Han, Z., and Wu, X.-M. Deeper insights into graph
convolutional networks for semi-supervised learning. In
AAAI, 2018.
- Li, W. Probability of all real zeros for random polynomial
with the exponential ensemble. *Preprint*, 2011.
- Li, Y., Tarlow, D., Brockschmidt, M., and Zemel, R. Gated
graph sequence neural networks. In *ICLR*, 2016.
- Liao, R., Zhao, Z., Urtasun, R., and Zemel, R. S. Lanczos-
net: Multi-scale deep graph convolutional networks. In
ICLR, 2019.
- Luan, S., Zhao, M., Chang, X.-W., and Precup, D. Break the
ceiling: Stronger multi-scale deep graph convolutional
networks. In *NeurIPS*, 2019.
- Maron, H., Ben-Hamu, H., Shamir, N., and Lipman, Y.
Invariant and equivariant graph networks. In *ICLR*, 2018.
- Maron, H., Ben-Hamu, H., Serviansky, H., and Lipman, Y.
Provably powerful graph networks. In *NeurIPS*, 2019.
- Ming Chen, Z. W., Zengfeng Huang, B. D., and Li, Y. Sim-
ple and deep graph convolutional networks. In *ICML*,
2020.
- Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J.,
and Bronstein, M. M. Geometric deep learning on graphs
and manifolds using mixture model cnns. In *CVPR*, 2017.

- 495 Morris, C., Ritzert, M., Fey, M., Hamilton, W. L., Lenssen,
496 J. E., Rattan, G., and Grohe, M. Weisfeiler and leman
497 go neural: Higher-order graph neural networks. In *AAAI*,
498 2019.
- 499 Oono, K. and Suzuki, T. Graph neural networks exponen-
500 tially lose expressive power for node classification. In
501 *ICLR*, 2019.
- 503 Pei, H., Wei, B., Chang, K. C.-C., Lei, Y., and Yang, B.
504 Geom-gcn: Geometric graph convolutional networks. In
505 *ICLR*, 2020.
- 507 Rong, Y., Huang, W., Xu, T., and Huang, J. Dropedge:
508 Towards deep graph convolutional networks on node clas-
509 sification. In *ICLR*, 2019.
- 510 Sandryhaila, A. and Moura, J. M. Discrete signal processing
511 on graphs. *IEEE Trans. Signal Process.*, 2013.
- 513 Shi, J. and Malik, J. Normalized cuts and image segmenta-
514 tion. *TPAMI*, 2000.
- 516 Simonyan, K. and Zisserman, A. Very deep convo-
517 lutional networks for large-scale image recognition.
518 *arXiv:1409.1556*, 2014.
- 519 Veličković, P., Cucurull, G., Casanova, A., Romero, A., Lio,
520 P., and Bengio, Y. Graph attention networks. In *ICLR*,
521 2018.
- 523 Wu, F., Zhang, T., Souza Jr, A. H. d., Fifty, C., Yu, T.,
524 and Weinberger, K. Q. Simplifying graph convolutional
525 networks. In *ICML*, 2019.
- 527 Xu, K., Li, C., Tian, Y., Sonobe, T., Kawarabayashi, K.-i.,
528 and Jegelka, S. Representation learning on graphs with
529 jumping knowledge networks. In *ICML*, 2018.
- 531 Xu, K., Hu, W., Leskovec, J., and Jegelka, S. How powerful
532 are graph neural networks? In *ICLR*, 2019.
- 533 Zhou, D.-X. Universality of deep convolutional neural net-
534 works. *ACHA*, 2020.
- 535
536
537
538
539
540
541
542
543
544
545
546
547
548
549