

KerGNNs: Interpretable Graph Neural Networks with Graph Kernels

Aosong Feng¹, Chenyu You¹, Shiqiang Wang², Leandros Tassioulas¹

¹ Yale University, New Haven, CT, USA

² IBM T. J. Watson Research Center, Yorktown Heights, NY, USA

aosong.feng@yale.edu, chenyu.you@yale.edu, wangshiq@us.ibm.com, leandros.tassioulas@yale.edu

Abstract

Graph kernels are historically the most widely-used technique for graph classification tasks. However, these methods suffer from limited performance because of the hand-crafted combinatorial features of graphs. In recent years, graph neural networks (GNNs) have become the state-of-the-art method in downstream graph-related tasks due to their superior performance. Most GNNs are based on Message Passing Neural Network (MPNN) frameworks. However, recent studies show that MPNNs can not exceed the power of the Weisfeiler-Lehman (WL) algorithm in graph isomorphism test. To address the limitations of existing graph kernel and GNN methods, in this paper, we propose a novel GNN framework, termed *Kernel Graph Neural Networks* (KerGNNs), which integrates graph kernels into the message passing process of GNNs. Inspired by convolution filters in convolutional neural networks (CNNs), KerGNNs adopt trainable hidden graphs as graph filters which are combined with subgraphs to update node embeddings using graph kernels. In addition, we show that MPNNs can be viewed as special cases of KerGNNs. We apply KerGNNs to multiple graph-related tasks and use cross-validation to make fair comparisons with benchmarks. We show that our method achieves competitive performance compared with existing state-of-the-art methods, demonstrating the potential to increase the representation ability of GNNs. We also show that the trained graph filters in KerGNNs can reveal the local graph structures of the dataset, which significantly improves the model interpretability compared with conventional GNN models.

Introduction

In recent years, the machine learning research community has devoted substantial energy to applying graph neural networks (GNNs) to numerous downstream graph-related tasks (Ying et al. 2018; Kipf and Welling 2016; Zhang and Chen 2018; Chen, Li, and Bruna 2017). Considering graph-structured tasks, the commonalities between different variants of GNNs are Message Passing Neural Networks (MPNNs) (Gilmer et al. 2017). MPNNs consist of two stages, including neighborhood aggregation and graph-level readout. Specifically, for neighborhood aggregation, there are three steps for each node to generate embeddings: (1)

Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

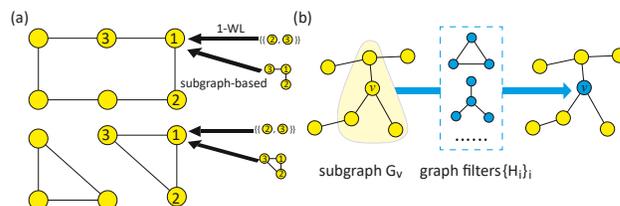


Figure 1: (a) 1-WL graph isomorphism test cannot distinguish one hexagon and two triangles because of same neighborhood multisets, while subgraph-based method can find the difference based on different subgraph topologies. (b) The yellow shadow represents the subgraph of node v . After interacting with graph filters, the updated node is colored in blue.

receiving messages from its neighbors, (2) aggregating messages, and (3) updating its own features to encode the local structural information. For graph-level tasks, a permutation-invariant readout function is used to extract feature representations from the entire graph. In fact, MPNNs have been motivated and derived as a continuous and differentiable analog of the Weisfeiler-Lehman (WL) algorithm (Leman and Weisfeiler 1968) which is known to successfully test graph isomorphism for a broad class of graphs. However, recent studies (Xu et al. 2018; Morris et al. 2019) show that MPNNs are at most as powerful as the WL kernel (Shervashidze et al. 2011) and WL algorithm regarding the graph isomorphism tests. This demonstrates theoretical limits in the expressivity of popular GNNs. For example, Figure 1(a) shows two graphs which cannot be distinguished by 1-WL algorithm, and therefore are also indistinguishable by MPNNs.

Before the advent of GNNs, graph kernels were the most widely-used techniques for solving graph classification tasks (Kriege, Johansson, and Morris 2020). Graph kernels measure the similarity between graphs, and can be applied into a kernel machine (e.g., support vector machine). Kernel functions remove the need of learning node embeddings in high dimensions, and enable us to operate in a high dimensional feature space by simply computing the kernel value in the low-dimensional feature space, which is more computationally efficient than computing in the high-dimensional space directly. Because of the empirical success of kernel-based methods and the increasing availability of graph-structured datasets, numerous graph kernel meth-

ods have been proposed, including walks and paths kernels (Gärtner, Flach, and Wrobel 2003; Kashima, Tsuda, and Inokuchi 2003; Borgwardt and Kriegel 2005), subgraph kernels (Shervashidze et al. 2009), and WL kernels (Shervashidze et al. 2011). However, graph kernels still have limitations due to their hand-crafted features and fixed feature construction scheme, which may not effectively capture high-dimensional information (e.g., complex node interactions) on large graphs.

In this paper, to address the above-mentioned issues and increase the expressivity of GNNs, we propose a **subgraph-based node aggregation algorithm** by combining GNNs and graph kernels into one framework, and thus the advantages of both methods can be leveraged. On one hand, for neighborhood aggregation, we apply graph kernels which use the subgraph induced by node neighbors, so that the expressivity will not be limited by 1-WL isomorphism test which uses the multiset of neighboring nodes. An example is shown in Figure 1(a), where we note that node 1 in both graphs has the same neighborhood multiset but induce different subgraph topologies with its neighbors which can be distinguished by graph kernels. On the other hand, we make the feature construction scheme of graph kernels trainable following the standard GNN training framework, possibly allowing for greater adaptability.

Based on the subgraph-based node aggregation, we propose a novel GNN framework, termed *KerGNNs*. Specifically, we first introduce a set of trainable hidden graphs, named graph filters, in each layer. Each node within the input graph is associated with a subgraph capturing its local topological information. We then adopt graph kernel functions to compare the similarity of graph filters and input subgraphs, and use the computed kernel values to update the respective node’s feature representations (as shown in Figure 1(b)). We show that KerGNNs provide a new kernel perspective to extend the standard CNN structure into the graph domain and generalize most MPNNs. The proposed model is then evaluated with various real-world graph and node classification tasks, and the results show superior performance of KerGNNs compared with many existing state-of-the-art models. To better understand the predictions of GNN-based methods, KerGNNs can further visualize the trained graph filters, similar to visualizing filters in CNNs, and thus provide better human-interpretable explanations for a variety of graph-related tasks, compared to existing GNNs. **Our main contributions** are summarized as follows:

1. We use neighborhood subgraph topology combined with kernel methods for GNN neighborhood aggregation, and show with proof that the expressivity of this approach is not limited by the 1-WL algorithm.
2. We provide a new perspective to generalize CNNs into the graph domain, by showing that both 2-D convolution and graph neighborhood aggregation can be interpreted using the language of kernel methods.
3. Besides envisioning the output graphs of the model, KerGNNs can further reveal the local structure of the input graphs by visualizing the topology of trained graph filters, which significantly improves the model in-

terpretability and transparency compared with standard MPNNs.

Related Work

Expressivity

Several works have been devoted to improving the expressivity of GNNs by introducing spatial, hierarchical, and higher-order GNN variants. For example, Abu-El-Haija et al. (2019) proposed the mix-hop structure which can learn a more general class of neighborhood mixing relationships. Sato, Yamada, and Kashima (2019) proposed to use Consistent Port Numbering GNN to augment the neighborhood aggregation, but port orderings are not unique and different orderings may lead to different expressivity. Klicpera, Groß, and Günnemann (2020) leveraged the atom coordinate information in the molecular graph to improve the expressivity, but the notion of direction is hard to generalize to more general graphs. Nguyen and Maehara (2020) used the graph homomorphism numbers as updated embeddings and show the expressivity of such graph classifiers with universality property, which unfortunately lacks neural network structure. Higher-order GNN variants have been studied in Morris et al. (2019) and Maron et al. (2019), which is more powerful than the 1-WL graph isomorphism test. However, higher-order methods always involve heavy computation and KerGNNs introduce a different way to break this 1-WL limit.

Combination of Graph Kernel and GNNs

Graph kernels and GNNs can be combined in the same framework. Some works apply graph kernels and neural networks at different stages (Navarin, Tran, and Sperduti 2018; Nikolentzos et al. 2018). There are also works on using GNN architecture to design new kernels. For example, Du et al. (2019) proposed a graph kernel equivalent to infinitely wide GNNs which can be trained using gradient descent.

A different line of research focuses on integrating kernel methods into GNNs. Lei et al. (2017) mapped inputs to RKHS by comparing inputs with reference objects. However, the reference objects they use lack graph structure and may not be able to capture the structural information. Chen, Jacob, and Mairal (2020) proposed GCKN which maps the input into a subspace of RKHS using walk and path kernel. While GCKN utilizes the local walk and path only starting from the central node, our model considers any walks (up to a maximal length) within the subgraph around the central node, and can thus explore more topological structures. Another recent work by Nikolentzos and Vazirgianis (2020) focused on improving model transparency by calculating the graph kernels between trainable hidden graphs and the entire graph. However, the method only supports a single-layer model and lacks theoretical interpretation. Our KerGNN model generalizes their scenario by applying hidden graphs to extract local structural information instead of the entire graph, and therefore constructs a multi-layer structure with better graph classification performance.

Explainability

Both graph structures and feature information lead to complex GNN models, making it hard for a human-intelligible explanation of the prediction results. Therefore, the transparency and explainability of GNN models are important issues to address. Baldassarre and Azizpour (2019) compared two main classes of explainability methods using infection and solubility problems. Pope et al. (2019) introduced explainability methods for the popular graph convolutional neural networks and demonstrated the extended methods on visual scene graphs and molecular graphs. Ying et al. (2019) proposed a model-agnostic approach that can identify a compact subgraph that has a crucial role in GNN’s prediction. In addition to visualizing output graphs as in regular GNNs, our KerGNN provides trained hidden graphs as a byproduct of training without additional computations, which contain useful structural information showing the common characteristics of the whole dataset instead of one specific graph, and can be helpful for interpreting the predictions of GNNs.

Background: Graph Kernels

Graph kernels have been proposed to solve the problem of assessing the similarity between graphs, and therefore making it possible to perform classification and regression with graph-structured data. Most graph kernels can be written as the sum of several pair-wise base kernels, following the \mathcal{R} -convolution framework (Haussler 1999):

$$K(G, G') = \sum_{v \in \mathcal{V}} \sum_{v' \in \mathcal{V}'} k_{\text{base}}(v, v'), \quad (1)$$

where $G = (\mathcal{V}, \mathcal{E})$, $G' = (\mathcal{V}', \mathcal{E}')$ are two input graphs with node attributes, and k_{base} can be any positive definite kernel defined on the node attributes. In this paper, we mainly consider random walk kernel which will be integrated into our proposed model in the next section.

Random walk kernels are one of the most studied graph kernels. They count the number of walks that two graphs have in common, and were initially proposed by Gärtner, Flach, and Wrobel (2003) and Kashima, Tsuda, and Inokuchi (2003). Among numerous variations of the random walk kernel, we deploy the P -step random walk kernel which compares random walks up to length P in two graphs.

Following Equation 1, we can write the base kernel of random walks with length p as

$$k_{\text{base}}^p(v, v') = \begin{cases} \langle a(v), a(v') \rangle, & \text{if } p = 0 \\ \langle a(v), a(v') \rangle \cdot \lambda \sum_{u \in \mathcal{N}(v)} \sum_{u' \in \mathcal{N}(v')} k_{\text{base}}^{(p-1)}(u, u'), & \text{if } p > 0 \end{cases}$$

where λ is the coefficient, $\mathcal{N}(v)$ denotes neighbors of v , p denotes the length of random walks which we compare in two graphs. If $p = 0$, the random walk kernel is equivalent to the simple node-pair kernel. To efficiently compute the random walk kernels, we follow the generalized framework of computing walk-based kernel (Vishwanathan et al. 2006), and utilize the direct product graph defined as below.

Definition 1 (Direct Product Graph). For two labeled graphs $G = (\mathcal{V}, \mathcal{E})$ and $G' = (\mathcal{V}', \mathcal{E}')$, the direct product graph is defined as $G_{\times} = G \times G' = (\mathcal{V}_{\times}, \mathcal{E}_{\times})$, defined as $\mathcal{V}_{\times} = \{(v, v') : v \in \mathcal{V} \wedge v' \in \mathcal{V}'\}$ and $\mathcal{E}_{\times} = \{ \{(v, v'), (u, u')\} : \{v, u\} \in \mathcal{E} \wedge \{v', u'\} \in \mathcal{E}' \}$.

Performing a random walk on the direct product graph G_{\times} is equivalent to performing the simultaneous random walks on graphs G and G' . The P -step random walk kernel can be calculated as

$$K(G, G') = \sum_{p=0}^P K_p(G, G') = \sum_{p=0}^P \lambda_p \sum_{i,j=1}^{|\mathcal{V}_{\times}|} [A_{\times}^p]_{ij}, \quad (2)$$

where A_{\times} is the adjacency matrix of G_{\times} and $\lambda = (\lambda_0, \lambda_1, \dots)$ is a sequence of weights. It should be noted that the (i, j) -th element of A_{\times}^p (i.e., A_{\times} to the power of p) represents the number of common walks of length p between the i -th and j -th node in G_{\times} .

To generalize the above formula into the continuous and multi-dimensional scenario, we first define the vertex attributes of the direct product graph G_{\times} . Given the node attribute matrix $\mathbf{X} \in \mathbb{R}^{n \times d}$ for a graph with n nodes and each node attribute is of dimension d , the node attribute matrix \mathbf{S} of the direct product graph $G_{\times} = G_1 \times G_2$ is calculated as $\mathbf{S} = \mathbf{X}_1 \mathbf{X}_2^T$, where $\mathbf{X}_1 \in \mathbb{R}^{n_1 \times d}$ and $\mathbf{X}_2 \in \mathbb{R}^{n_2 \times d}$ are the node attribute matrices for G_1 and G_2 , respectively, and $\mathbf{S} \in \mathbb{R}^{n_1 \times n_2}$. The (i, j) -th element of matrix \mathbf{S} encodes the similarity between the i -th-node of G_1 and the j -th node of G_2 . We flatten \mathbf{S} into vector $\mathbf{s} \in \mathbb{R}^{n_1 n_2}$ for ease of notation, and then integrate the encoded pair-wise similarity into Equation 2

$$K_p(G, G') = \sum_{i,j=1}^{|\mathcal{V}_{\times}|} \mathbf{s}_i \mathbf{s}_j [A_{\times}^p]_{ij} = \mathbf{s}^T A_{\times}^p \mathbf{s}. \quad (3)$$

Based on this equation, we can calculate the kernel value between two input graphs using the similarity of common walks as the metric. The details of calculating Equation 3 are included in Appendix.

In practice, we also consider a slight variation of Equation 1 by adding trainable weights to each base kernel term, and we call it deep random walk kernel:

$$K(G, G') = \sum_{v \in \mathcal{V}} \sum_{v' \in \mathcal{V}'} w_{(v, v')} k_{\text{base}}(v, v'), \quad (4)$$

where $w_{(v, v')}$ represents the trainable weight assigned to the base kernel.

Proposed Model

In this section, we first discuss the framework of the proposed KerGNN model. Then we introduce the concept of subgraph-based neighborhood aggregation, and use it to analyze the expressivity of KerGNNs. Next, we show that KerGNNs are inspired by CNNs and compare them from the kernel perspective. Finally we argue that KerGNNs can generalize MPNN architecture and analyze the time complexity.

KerGNN Framework

In this subsection, we introduce the KerGNN model which updates each node’s embedding according to the subgraph centered at this node instead of the rooted subtree patterns in MPNNs, as shown in Figure 1(b). Unless otherwise specified, we refer to the subgraph as the vertex-induced subgraph formed from a node and all its 1-hop neighbors.

We first define the embeddings of nodes and subgraphs, which are mapping functions from graphs to the feature space and from nodes to the feature space.

Definition 2 (Feature mapping). *Given a graph $G = (\mathcal{V}, \mathcal{E})$, a node feature mapping is a node-wise mapping function $\phi : \mathcal{V} \rightarrow \mathbb{R}^d$, which maps every node $v \in \mathcal{V}$ to a point $\phi(v)$ in \mathbb{R}^d , and $\phi(v)$ is called the feature map for node v . A graph feature mapping is a function $\Phi : \mathcal{G} \rightarrow \mathbb{R}^d$, where \mathcal{G} is the set of graphs, and $\Phi(G)$ is called the feature map for graph $G \in \mathcal{G}$.*

For an L -layer neural network, we call the input layer the 0-th layer. At each hidden layer l , the input to this layer is an undirected graph $G = (\mathcal{V}, \mathcal{E})$, and each node $v \in \mathcal{V}$ has a feature map $\phi_{l-1}(v) \in \mathbb{R}^{d_{l-1}}$. The output of layer l is the same graph G , because we do not consider graph pooling here, and each node $v \in \mathcal{V}$ in the output graph has a feature map $\phi_l(v) \in \mathbb{R}^{d_l}$. For example, G can be a graph in the dataset, and $\phi_0(v)$ is the node attributes with dimension d_0 . For graphs with discrete node labels, the attributes can be represented as the one-hot encodings of labels and the dimension of attributes corresponds to the total number of classes. For graphs without node labels, we use the degree of the node as the node attribute.

Inspired by the filters in CNN, we define a set of graph filters at each KerGNN layer to extract the local structural information around each node in the input graph (see Figure 1(b)).

Definition 3 (Graph filter). *The i -th graph filter at layer l is a graph $H_i^{(l)}$ with $n_i^{(l)}$ nodes. It has a trainable adjacency matrix $A_i^{(l)} \in \mathbb{R}^{n_i^{(l)} \times n_i^{(l)}}$ and node attribute matrix $W_i^{(l)} \in \mathbb{R}^{n_i^{(l)} \times d_{l-1}}$.*

At layer l , there are d_l graph filters such that the output dimension is also d_l , and each node attribute in the graph filter, represented by each row of $W_i^{(l)}$, has the same dimension as the node feature map $\phi_{l-1}(v)$ in the input graph.

KerGNN Layer. Now we consider a single KerGNN layer. We assume the input is a graph-structured dataset with undirected graph $G = (\mathcal{V}, \mathcal{E})$, and each node $v \in \mathcal{V}$ has the attribute $a(v) \in \mathbb{R}^{d_0}$. Then the input node feature map is $\phi_0(v) = a(v)$.

Each node v in the graph is equipped with a subgraph $G_v = (\mathcal{V}_v, \mathcal{E}_v)$, and feature maps $\{\phi_0(u) : u \in \mathcal{V}_v\}$ are transformed to $\phi_1(v)$ in a way such that neighbors’ local information (topological information and node representations) contained in G_v will be aggregated to the central node v . We then rely on the graph filters $\{H_i^{(1)} : i = 1, \dots, d_1\}$ to obtain $\phi_1(v)$. Specifically, we calculate $\phi_1(v)$ by projecting subgraph feature map $\Phi_0(G_v)$ into the i -th dimension of $\phi_1(v)$ using the kernel function value between graph filter

$H_i^{(1)}$ and subgraph G_v , i.e.,

$$\phi_{1,i}(v) = K(G_v, H_i^{(1)}), \quad (5)$$

where we adopt a random walk kernel as $K(\cdot, \cdot)$, which is introduced in Equation 2. After calculating the kernel value of the subgraph G_v with respect to every graph filter $\{H_i^{(1)} : i = 1, \dots, d_1\}$, we obtain every dimension of node v ’s feature map $\phi_0(v)$, which forms the output of the KerGNN layer.

It should be noted that using graphs G_v and $H_i^{(1)}$ to calculate the kernel value is equivalent to performing inner product of $\phi_1(G_v)$ and $\phi_1(H_i^{(1)})$ in an implicit high-dimensional space, and using feature map of G_v instead of the multiset of neighboring nodes (as used in MPNNs) improves expressivity, which is analyzed in the next subsection. Besides, if we use the output space \mathbb{R}^{d_1} to approximate the high-dimensional space introduced by the kernel method, the updating rule will correspond to the convolutional kernel network proposed by Mairal et al. (2014), and we will follow the same idea when we compare KerGNNs with CNNs in the later subsection.

Multiple-layer Model. Based on the single-layer analysis above, we can construct a multiple-layer KerGNN by stacking KerGNN layers followed by readout layers. Specifically, the input to layer l is the graph G with node feature map $\{\phi_{l-1}(v) : v \in G\}$. Layer l is parameterized by d_l graph filters $\{H_i^{(l)} : i = 1, \dots, d_l\}$. Each graph filter $H_i^{(l)}$ has a trainable adjacency matrix $A_i^{(l)}$ and node attributes $W_i^{(l)}$. Then the i -th dimension of the output feature map for node v in G can be explicitly calculated as

$$\phi_{l,i}(v) = K(G_v, H_i^{(l)}). \quad (6)$$

The forward pass of the l th-layer of KerGNNs is summarized in Algorithm 1.

For the graph classification, we then deploy the graph-level readout layer to generate the embedding for the entire graph. We obtain the graph representation at each layer by summing all the nodes’ representations. To leverage information from every layer of the model, we then concatenate the graph representations across all layers:

$$\Phi(G) = \text{concat} \left(\sum_{v \in G} \phi_l(v) \mid l = 0, 1, \dots, L \right). \quad (7)$$

Expressivity of Subgraph-based Aggregation

In this subsection, we first define the subgraph-based neighborhood aggregation, and discuss the requirements of the subgraph feature map to achieve higher expressivity than 1-WL algorithm, then we show that KerGNN is one of the models that satisfy these requirements.

To leverage the structural information contained in the subgraph, we aggregate the subgraph information by finding a proper subgraph feature map $\Phi(G_v)$, and update the node representation of v combining the subgraph feature map with v ’s own feature map. Formally, we define this aggregation process as follows.

Algorithm 1: Forward pass in l -th KerGNN layer

Input: Graph $G = (\mathcal{V}, \mathcal{E})$; Input node feature maps $\{\phi_{l-1}(v) : v \in \mathcal{V}\}$; Graph filters $\{H_i^{(l)} : i = 1, \dots, d_l\}$; Graph kernel function K

Output: Graph $G = (\mathcal{V}, \mathcal{E})$; Output node feature maps $\{\phi_l(v) : v \in \mathcal{V}\}$

```
for  $v \in \mathcal{V}$  do  
   $G_v = \text{subgraph}(\{v\} \cup \mathcal{N}(v))$ ;  
  for  $i = 1$  to  $d_l$  do  
     $\phi_{l,i}(v) = K(G_v, H_i^{(l)})$ ;  
  end for  
end for
```

Definition 4 (Subgraph-based aggregation). *The graph neural network at layer l deploying subgraph-based neighborhood aggregation updates feature mapping ϕ according to $\phi_l(v) = u(\phi_{l-1}(v), f(\Phi_{l-1}(G_v)))$, where u and f are update and aggregation functions, respectively.*

GNNs distinguish different graphs by mapping them to different embeddings, which resembles the graph isomorphism test. Xu et al. (2018) characterize the representational capacity of MPNNs using the WL graph isomorphism test criterion, and show that MPNNs can be as powerful as 1-WL graph isomorphism test if the node update, aggregation, and graph-level readout function are injective. We follow the similar approach and show in the following that subgraph-based GNNs like KerGNNs can be at least as powerful as the 1-WL graph isomorphism test.

Because we are comparing the model’s expressivity with the 1-WL algorithm which updates node labels based on the multiset of neighboring nodes, to achieve high expressivity, it is natural to think that $\Phi(G_v)$ should have a one-to-one relationship with respect to the multiset of nodes that subgraph G_v contains. We show in Lemma 1 that the graph feature map induced by the random walk kernel satisfies this condition.

Lemma 1 *if $\Phi(G)$ is the feature map of graph G induced by the random walk graph kernel, then $\Phi(G)$ is injective with respect to the multiset of all its contained nodes $\{\{a(v) : v \in \mathcal{V}(G)\}\}$, where $\{\cdot\}$ denotes the multiset and $a(v)$ is the label or attribute of node v .*

The proof follows directly from the random walk kernel definition in Gärtner, Flach, and Wrobel (2003), and we notice that the graph feature map induced by the WL graph kernel also satisfies this lemma. Based on this injective relationship between multiset and subgraph feature map, we can compare the expressivity of the subgraph-based GNN and 1-WL graph isomorphism test using the following theorem.

Theorem 1 *Let $\mathcal{A} : \mathcal{G} \rightarrow \mathbb{R}^d$ be a GNN with a sufficient number of GNN layers, if the following conditions hold at layer l :*

a) \mathcal{A} aggregates and updates node features iteratively with $\phi_l(v) = u(\phi_{l-1}(v), f(\Phi_{l-1}(G_v)))$, where function u and f are injective, and Φ_{l-1} is the feature mapping induced by the random walk kernel;

b) \mathcal{A} ’s graph-level readout, which operates on the multiset of node features $\{\{\phi_l(v)\}\}$, is injective;
then \mathcal{A} maps any graphs G and H that 1-WL test decides as non-isomorphic to different embeddings, and there exist graph G and H that 1-WL test decides as isomorphic, but can be mapped to different embeddings by \mathcal{A} .

The proof is shown in Appendix. This theorem shows that subgraph-based GNNs can be more expressive than the 1-WL isomorphism test and thus MPNNs. In the KerGNN model, we do not explicitly calculate the subgraph feature map $\Phi(G_v)$ which lives in the high-dimensional space. Instead, we apply the kernel trick and use the subgraph feature map as $K(G_v, H) = \langle \Phi(G_v), \Phi(H) \rangle$. Then, the graph kernel function $K(\cdot, H)$ can be seen as a composition of functions u and f . Therefore, according to Theorem 1, to achieve high representational power, the graph kernel function needs to be injective with respect to the subgraph feature map $\Phi(G_v)$, and we introduce the following lemma to show that the KerGNN model satisfies this requirement.

Lemma 2 *There exists a feature map $\Phi(H)$ so that $K(H, G_v) = \langle \Phi(H), \Phi(G_v) \rangle$ is unique for different $\Phi(G_v)$.*

The proof is shown in Appendix. Besides, as shown in the definition of graph filters, in the KerGNN model we parameterize the node feature and adjacency matrix of graph filter H instead of directly parameterizing $\Phi(H)$.

Connections to CNNs

Standard CNN models update the representation of each pixel by convolving filters with the patch centered at it, and in GNNs, a natural analog of the patch in the graph domain is the subgraph. While many MPNNs draw connections with CNNs by extending 2-D convolution to the graph convolution, we show in this subsection that both 2-D image convolution and KerGNN aggregation process can be viewed as applying kernel tricks to the input image or graph, and therefore, the KerGNN model naturally extends the CNN architecture into the graph domain, from a new kernel perspective.

We first show in Appendix that under suitable assumptions, the 2-D image convolution can be viewed as applying kernel functions between input patches and filters. The basic idea is that we can rethink the 2-D convolution as projecting the input image patch into the kernel-induced Hilbert space. The projection is done by performing inner product between the patch and basis vectors, which can be calculated using the kernel trick, and the projected representation in the output space will be the output of the CNN layer.

Then we can extend the same philosophy to the graph domain, by introducing subgraphs and topology-aware graph filters as the counterpart of patches and filters in CNNs, and KerGNN will adopt the kernel trick to project the input subgraph representation into the output space (detailed in Appendix). Based on these two observations, we can see that KerGNNs generalize CNNs into the graph domain by replacing the kernel function for vectors with the graph kernel function, which provides a new insight into designing GNN

architecture, different from the spatial and spectral convolution perspectives.

Connections to Existing GNNs

As the subgraph of one node can be a more fruitful source of information than just the multiset of its neighbors, we show in this subsection that KerGNNs can generalize the standard MPNNs. From the point of view of KerGNNs, MPNNs deploy a simple graph filter with one node, and an appropriate kernel function can be chosen within KerGNN framework, such that KerGNNs iteratively update nodes' representations using neighborhood multiset aggregation like in MPNNs. For example, we show in Appendix that the node update rule of Graph Convolutional Network (GCN) (Kipf and Welling 2016) can be treated as using one-node graph filters with properly-defined \mathcal{R} -convolution graph kernel. Our model generalizes most MPNN structures by deploying more complex graph filters with multiple nodes and learnable adjacency matrix, and using more expressive and efficient graph kernels.

Time Complexity Analysis

Most MPNNs incur a time complexity of $\mathcal{O}(n^2)$, or $\mathcal{O}(m)$ if the adjacency matrix is sparse containing m non-zero entries, because updating the embedding of node v involves n_v neighbors, where n_v is the degree of node v . In KerGNNs, we apply graph kernel with the subgraph G_v instead of the whole graph, so the computational complexity would be related to the complexity of each subgraph. For the subgraph G_v with $n_v + 1$ nodes and adjacency matrix with m_v non-zero entries, we update the representation of node v by calculating the random walk kernel. This calculation takes a computation time of $\mathcal{O}(Pd(d'n_{GF}(n_{GF} + n_v + 1) + m_v))$, where P is the maximum length of the random walk, d and d' are the node dimensions of the current layer and next layer, n_{GF} is the number of nodes in each graph filter. In an undirected subgraph, m_v represents the number of edges and will be greater than n_v and smaller than $n_v(n_v - 1)/2$. If we sum up the computation time for all the nodes in the entire graph, the time complexity of KerGNNs will range between $\mathcal{O}(n^2)$ and worst-case scenario (fully-connected graph) $\mathcal{O}(n^3)$. We experimentally compare the running time of the proposed model with several GNN benchmarks. As shown in Appendix, KerGNNs achieve better or similar running time compared to the fastest benchmark method, and much less running time than higher-order GNNs.

Experiments

We evaluate the proposed model on graph classification task and node classification task (discussed in Appendix), and we also show the model interpretability by visualizing the graph filters in the trained models as well as the output graphs.

Experiment Settings

Datasets. We evaluate our proposed KerGNN model on 8 publicly available graph classification datasets. Specifically, we use DD (Dobson and Doig 2003), PROTEINS (Borgwardt et al. 2005), NCI1 (Schomburg et al. 2004),

ENZYMES (Schomburg et al. 2004) for binary and multi-class classification of biological and chemical compounds, and we also use the social datasets IMDB-BINARY, IMDB-MULTI, REDDIT-BINARY, and COLLAB (Yanardag and Vishwanathan 2015).

Setup. To make a fair comparison with state-of-the-art GNNs, we follow the cross-validation procedure described in Errica et al. (2019). We use a 10-fold cross-validation for model assessment and an inner holdout technique with a 90%/10% training/validation split for model selection, following the same dataset index splits as Errica et al. (2019). Besides, we use Adam optimizer with an initial learning rate of 0.01 and decay the learning rate by half in every 50 epochs. For the four social datasets, we use node degrees as the input attributes for each node, and for the four bio/chemical datasets, we use node labels or attributes as the input feature for each node.

Hyper-parameters. The hyper-parameters that we tune for each dataset include the learning rate, the dropout rate, the number of layers of KerGNNs and MLP, the number of graph filters at each layer, the number of nodes in each graph filter, the number of nodes for each subgraph, and the hidden dimension of each KerGNN layer. For the random walk kernel, we also tune the length of random walks.

Baseline Models. We consider the KerGNN model with single and multiple KerGNN layers, namely KerGNN- L , corresponding to KerGNN model with L layers, and KerGNN- L -DRW representing the model deploying the deep random walk kernel. We also compare our models with widely-used GNNs: DGCNN (Zhang et al. 2018), Diff-Pool (Ying et al. 2018), ECC (Simonovsky and Komodakis 2017), GIN (Xu et al. 2018), GraphSAGE (Hamilton, Ying, and Leskovec 2017), RWGNN (Nikolentzos and Vazirgianis 2020), GCKN (Chen, Jacob, and Mairal 2020), and two high-order GNNs: 1-2-3 GNN (Morris et al. 2019) and Powerful GNN (Maron et al. 2019). Part of the results for these baseline GNNs are taken from Errica et al. (2019), and we run GCKN, 1-2-3 GNN and Powerful GNN using the official implementations. In addition, we also compare the proposed KerGNN model with three popular GNN-unrelated graph kernels: shortest path (SP) kernel (Borgwardt and Kriegel 2005), propagation (PK) kernel (Neumann et al. 2016), the Weisfeiler-Lehman subtree (WL-sub) kernel (Shervashidze et al. 2011) and GNN-related GNTK (Du et al. 2019). We use the GraKeL library (Siglidis et al. 2020) to implement these graph kernels and run GNTK using the official implementation.

Results

The graph classification results are shown in Table 1, with the best results highlighted in bold. We can see that the proposed models achieve superior performance than conventional GNNs with 1-WL limits, and achieve similar performance compared with high-order GNNs, with less running time. The single-layer KerGNN model performs well on small graphs like IMDB social datasets. For larger graphs, deeper models with more layers or with deep random walk kernel perform better. We show more experimental results,

	DD	NCI1	PROTEINS	ENZYMES	IMDB-B	IMDB-M	REDDIT-B	COLLAB
# GRAPHS	1178	4110	1113	600	1000	1500	2000	5000
# CLASSES	2	2	2	6	2	3	2	3
AVG. # NODES	284	30	39	33	20	13	430	74
SP	78.7±3.8	66.3±2.6	71.9±6.1	25.0±5.6	57.5±5.4	40.5±2.8	75.5±2.1	58.4±1.3
PK	78.0±3.8	72.3±2.8	59.7±0.3	61.0±6.7	73.9±4.3	51.1±5.8	68.5±2.9	77.3±2.4
WL-SUB	77.5±3.5	79.5±3.3	74.8±3.2	51.2±5.3	72.5±4.6	51.5±5.8	67.2±4.2	77.5±2.4
GNTK	OOB	83.5±1.2	75.5±2.2	48.2±2.4	75.9±3.1	52.2±4.2	OOB	OOB
DGCNN	76.6±4.3	76.4±1.7	72.9±3.5	38.9±5.7	69.2±3.0	45.6±3.4	87.8±2.5	71.2±1.9
DIFFPOOL	75.0±3.5	76.9±1.9	73.7±3.5	59.5±5.6	68.4±3.3	45.6±3.4	89.1±1.6	68.9±2.0
ECC	72.6±4.1	76.2±1.4	72.3±3.4	29.5±8.2	67.7±2.8	43.5±3.1	OOB	OOB
GIN	75.3±2.9	80.0±1.4	73.3±4.0	59.6±4.5	71.2±3.9	48.5±3.3	89.9±1.9	75.6±2.3
GRAPHSAGE	72.9±2.0	76.0±1.8	73.0±4.5	58.2±6.0	68.8±4.5	47.6±3.5	84.3±1.9	73.9±1.7
RWGNN	77.6±4.7	73.9±1.3	74.7±3.3	57.6±6.3	70.8±4.8	48.8±2.9	90.4±1.9	71.9±2.5
GCKN	77.3±4.0	79.2±1.2	76.1±2.8	59.3±5.6	74.5±1.2	51.0±3.9	OOB	74.3±2.8
1-2-3 GNN	OOB	72.7±2.9	74.5±5.6	OOB	70.7±3.4	50.2±2.2	91.1±2.1	OOB
POWERFUL GNN	OOB	83.4±1.8	75.9±3.3	54.8±5.5	73.0±4.9	50.5±3.2	OOB	75.4±1.4
KERGNN-1	77.6±3.7	74.3±2.2	75.8±3.5	62.1±5.5	74.4±4.3	51.6±3.1	81.5±1.9	70.5±1.6
KERGNN-2	78.9±3.5	76.3±2.6	75.5±4.6	55.0±5.0	73.7±4.0	50.9±5.1	82.0±2.5	72.7±2.1
KERGNN-3	75.5±3.1	80.5±1.9	76.5±3.9	54.1±4.3	72.1±4.6	50.1±4.5	82.0±1.9	71.1±2.0
KERGNN-2-DRW	77.0±4.4	82.8±1.8	76.1±4.1	59.5±4.5	71.1±4.1	50.5±3.1	89.5±1.6	75.1±2.3

Table 1: Test set classification accuracies (%). The mean accuracy and standard deviation are reported. Best performances are highlighted in bold. OOB means Out of Resources, either time or GPU memory.

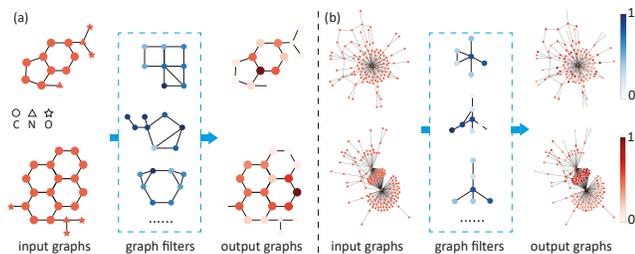


Figure 2: Model visualization. Input graphs are drawn from (a) MUTAG and (b) REDDIT-B datasets with different node shapes corresponding to different atom types. In both graph filters and output graphs, node color represents relative attribute value.

model parameter studies, and node classification results in Appendix. The optimal parameters of the graph filter are different for different datasets, depending on the local structures of different types of graphs, e.g., the star patterns in graphs of REDDIT-B and the ring and chain patterns in graphs of NCI1.

Model Interpretability

Visualizing the filters in CNNs gives insights into what features CNNs focus on. Following the same idea, we can also visualize the trained graph filters, which indicate some key structures of the input dataset. We visualize the graph filters trained with MUTAG and REDDIT-B dataset in Figure 2. The MUTAG dataset consists of 188 chemical compounds divided into two classes according to their mutagenic effect on a bacterium. As shown in the input graphs in Figure 2(a), most of the MUTAG graphs in the dataset consist of ring structures with 6 carbon atoms.

Compared to standard GNN variants, KerGNNs have graph filters as extra information to help explain the predictions of the model. To visualize the graph filters, we extract the adjacency matrix and the attribute matrix for each graph filter from the trained KerGNN layer. We then adopt the *ReLU* functions to prune the unimportant edges. In Figure 2, we use different sizes of nodes to denote the relative importance of nodes. For the MUTAG dataset, we can see most of the graph filters have ring structures, similar to the carbon rings at the input graphs, and some graph filters have small connected rings, similar to the concatenated carbon rings. It should be noted that the number of nodes in the rings of graph filters may not be equal to 6 because we limit the total number of nodes to be 8. KerGNN layers utilize these rings in the graph filter to match against the local structural patterns (e.g., carbon rings) in the input graphs, using the graph kernels. This indicates the importance of carbon rings in the mutagenic effect, which also corresponds to our observations in the output graphs.

Conclusion

In this paper, we have proposed Kernel Graph Neural Networks (KerGNNs), a new graph neural network framework that is not restricted to the theoretical limits of the message passing aggregation. KerGNNs are inspired by several characteristics of CNNs and can be seen as a natural extension of CNNs in the graph domain, from the viewpoint of the kernel methods. KerGNNs achieve competitive performance on a variety of datasets compared with several GNNs and graph kernels, and can provide improved explainability and transparency by visualizing the graph filters and output graphs.

Acknowledgments

This work was partially supported by the U.S. Office of Naval Research under Grant N00173-21-1-G006, the U.S. National Science Foundation AI Institute Athena under Grant CNS-2112562, and the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001.

References

- Abu-El-Haija, S.; Perozzi, B.; Kapoor, A.; Alipourfard, N.; Lerman, K.; Harutyunyan, H.; Ver Steeg, G.; and Galstyan, A. 2019. Mixhop: Higher-order graph convolutional architectures via sparsified neighborhood mixing. In *international conference on machine learning*, 21–29. PMLR.
- Baldassarre, F.; and Azizpour, H. 2019. Explainability techniques for graph convolutional networks. *arXiv preprint arXiv:1905.13686*.
- Borgwardt, K. M.; and Kriegel, H.-P. 2005. Shortest-path kernels on graphs. In *Fifth IEEE international conference on data mining (ICDM'05)*, 8–pp. IEEE.
- Borgwardt, K. M.; Ong, C. S.; Schönauer, S.; Vishwanathan, S.; Smola, A. J.; and Kriegel, H.-P. 2005. Protein function prediction via graph kernels. *Bioinformatics*, 21(suppl_1): i47–i56.
- Chen, D.; Jacob, L.; and Mairal, J. 2020. Convolutional kernel networks for graph-structured data. In *International Conference on Machine Learning*, 1576–1586. PMLR.
- Chen, Z.; Li, X.; and Bruna, J. 2017. Supervised community detection with line graph neural networks. *arXiv preprint arXiv:1705.08415*.
- Dobson, P. D.; and Doig, A. J. 2003. Distinguishing enzyme structures from non-enzymes without alignments. *Journal of molecular biology*, 330(4): 771–783.
- Du, S. S.; Hou, K.; Póczos, B.; Salakhutdinov, R.; Wang, R.; and Xu, K. 2019. Graph neural tangent kernel: Fusing graph neural networks with graph kernels. *arXiv preprint arXiv:1905.13192*.
- Errica, F.; Podda, M.; Bacciu, D.; and Micheli, A. 2019. A fair comparison of graph neural networks for graph classification. *arXiv preprint arXiv:1912.09893*.
- Gärtner, T.; Flach, P.; and Wrobel, S. 2003. On graph kernels: Hardness results and efficient alternatives. In *Learning theory and kernel machines*, 129–143. Springer.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *International Conference on Machine Learning*, 1263–1272. PMLR.
- Hamilton, W. L.; Ying, R.; and Leskovec, J. 2017. Inductive representation learning on large graphs. *arXiv preprint arXiv:1706.02216*.
- Haussler, D. 1999. Convolution kernels on discrete structures. Technical report, Technical report, Department of Computer Science, University of California . . .
- Kashima, H.; Tsuda, K.; and Inokuchi, A. 2003. Marginalized kernels between labeled graphs. In *Proceedings of the 20th international conference on machine learning (ICML-03)*, 321–328.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Klicpera, J.; Groß, J.; and Günnemann, S. 2020. Directional message passing for molecular graphs. *arXiv preprint arXiv:2003.03123*.
- Kriege, N. M.; Johansson, F. D.; and Morris, C. 2020. A survey on graph kernels. *Applied Network Science*, 5(1): 1–42.
- Lei, T.; Jin, W.; Barzilay, R.; and Jaakkola, T. 2017. Deriving neural architectures from sequence and graph kernels. In *International Conference on Machine Learning, 2024–2033*. PMLR.
- Leman, A.; and Weisfeiler, B. 1968. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsiya*, 2(9): 12–16.
- Mairal, J.; Koniusz, P.; Harchaoui, Z.; and Schmid, C. 2014. Convolutional kernel networks. *arXiv preprint arXiv:1406.3332*.
- Maron, H.; Ben-Hamu, H.; Serviansky, H.; and Lipman, Y. 2019. Provably powerful graph networks. *arXiv preprint arXiv:1905.11136*.
- Morris, C.; Ritzert, M.; Fey, M.; Hamilton, W. L.; Lenssen, J. E.; Rattan, G.; and Grohe, M. 2019. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 4602–4609.
- Navarin, N.; Tran, D. V.; and Sperduti, A. 2018. Pre-training graph neural networks with kernels. *arXiv preprint arXiv:1811.06930*.
- Neumann, M.; Garnett, R.; Bauckhage, C.; and Kersting, K. 2016. Propagation kernels: efficient graph kernels from propagated information. *Machine Learning*, 102(2): 209–245.
- Nguyen, H.; and Maehara, T. 2020. Graph Homomorphism Convolution. In *International Conference on Machine Learning*, 7306–7316. PMLR.
- Nikolentzos, G.; Meladianos, P.; Tixier, A. J.-P.; Skianis, K.; and Vazirgiannis, M. 2018. Kernel graph convolutional neural networks. In *International Conference on Artificial Neural Networks*, 22–32. Springer.
- Nikolentzos, G.; and Vazirgiannis, M. 2020. Random Walk Graph Neural Networks. In *Conference on Neural Information Processing System*. PMLR.
- Pope, P. E.; Kolouri, S.; Rostami, M.; Martin, C. E.; and Hoffmann, H. 2019. Explainability methods for graph convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10772–10781.
- Sato, R.; Yamada, M.; and Kashima, H. 2019. Approximation ratios of graph neural networks for combinatorial problems. *arXiv preprint arXiv:1905.10261*.

- Schomburg, I.; Chang, A.; Ebeling, C.; Gremse, M.; Heldt, C.; Huhn, G.; and Schomburg, D. 2004. BRENDA, the enzyme database: updates and major new developments. *Nucleic acids research*, 32(suppl.1): D431–D433.
- Shervashidze, N.; Schweitzer, P.; Van Leeuwen, E. J.; Mehlhorn, K.; and Borgwardt, K. M. 2011. Weisfeiler-lehman graph kernels. *Journal of Machine Learning Research*, 12(9).
- Shervashidze, N.; Vishwanathan, S.; Petri, T.; Mehlhorn, K.; and Borgwardt, K. 2009. Efficient graphlet kernels for large graph comparison. In *Artificial intelligence and statistics*, 488–495. PMLR.
- Siglidis, G.; Nikolentzos, G.; Limnios, S.; Giatsidis, C.; Skianis, K.; and Vazirgiannis, M. 2020. GraKeL: A Graph Kernel Library in Python. *Journal of Machine Learning Research*, 21(54): 1–5.
- Simonovsky, M.; and Komodakis, N. 2017. Dynamic edge-conditioned filters in convolutional neural networks on graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3693–3702.
- Vishwanathan, S.; Borgwardt, K. M.; Schraudolph, N. N.; et al. 2006. Fast computation of graph kernels. In *NIPS*, volume 19, 131–138. Citeseer.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2018. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*.
- Yanardag, P.; and Vishwanathan, S. 2015. Deep graph kernels. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining*, 1365–1374.
- Ying, R.; Bourgeois, D.; You, J.; Zitnik, M.; and Leskovec, J. 2019. Gnnexplainer: Generating explanations for graph neural networks. *Advances in neural information processing systems*, 32: 9240.
- Ying, R.; You, J.; Morris, C.; Ren, X.; Hamilton, W. L.; and Leskovec, J. 2018. Hierarchical graph representation learning with differentiable pooling. *arXiv preprint arXiv:1806.08804*.
- Zhang, M.; and Chen, Y. 2018. Link prediction based on graph neural networks. *arXiv preprint arXiv:1802.09691*.
- Zhang, M.; Cui, Z.; Neumann, M.; and Chen, Y. 2018. An end-to-end deep learning architecture for graph classification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.