

RESEARCH

Open Access



An end-to-end graph convolutional kernel support vector machine

Padraig Corcoran

Correspondence:

corcoranp@cardiff.ac.uk

School of Computer Science & Informatics, Cardiff University, Queen's Buildings, 5 The Parade, Roath, CF24 3AA Cardiff, UK

Abstract

A novel kernel-based support vector machine (SVM) for graph classification is proposed. The SVM feature space mapping consists of a sequence of graph convolutional layers, which generates a vector space representation for each vertex, followed by a pooling layer which generates a reproducing kernel Hilbert space (RKHS) representation for the graph. The use of a RKHS offers the ability to implicitly operate in this space using a kernel function without the computational complexity of explicitly mapping into it. The proposed model is trained in a supervised end-to-end manner whereby the convolutional layers, the kernel function and SVM parameters are jointly optimized with respect to a regularized classification loss. This approach is distinct from existing kernel-based graph classification models which instead either use feature engineering or unsupervised learning to define the kernel function. Experimental results demonstrate that the proposed model outperforms existing deep learning baseline models on a number of datasets.

Keywords: Graph neural network, Kernel method, Support vector machine

Introduction

The world contains much implicit structure which can be modelled using a graph. For example, an image can be modelled as a graph where objects (e.g. person, chair) are modelled as vertices and their pairwise relationships (e.g. sitting) are modelled as edges (Krishna et al. 2017). This representation has led to useful solutions for many vision problems including image captioning and visual question answering (Chen et al. 2019). Similarly, a street network can be modelled as a graph where locations are modelled as vertices and street segments are modelled as edges. This representation has led to useful solutions for many transportation problems including the placement of electrical vehicle charging stations (Gagarin and Corcoran 2018).

Given the ubiquity of problems which can be modelled in terms of graphs, performing machine learning on graphs represents an area of great research interest. Advances in the application of deep learning or neural networks to sequence spaces in the context of natural language processing and fixed dimensional vector spaces in the context of computer vision has led to much interest in applying deep learning to graphs. There exist many types of machine learning tasks one may wish to perform on graphs. These

include vertex classification, graph classification, graph generation (You et al. 2018) and learning implicit/hidden structures (Franceschi et al. 2019). In this work we focus on the task of graph classification. Examples of graph classification tasks include human activity recognition where human pose is modelled using a skeleton graph (Yan et al. 2018), visual scene understanding where the scene is modelled using a *scene graph* (Xu et al. 2017) and semantic segmentation of three dimensional point clouds where the point cloud is modelled as a graph of geometrically homogeneous elements (Landrieu and Simonovsky 2018).

Graph convolutional is the most commonly used deep learning architecture applied to graphs. This architecture consists of a sequence of convolutional layers where each layer iteratively updates a vector space representation of each vertex. In their seminal work, Gilmer et al. (2017) demonstrated that many different convolutional layers can be formulated in terms of a framework containing two steps. In the first step, message passing is performed where each vertex receives messages from adjacent vertices regarding their current representation. In the second step, each vertex performs an update of its representation which is a function of its current representation and the messages it received in the previous step. In order to perform graph classification given a sequence of convolutional layers, the set of vertex representations output from this sequence must be integrated to form a graph representation. This graph representation can subsequently be used to predict a corresponding class label. We refer to this task of integrating vertex representations as *vertex pooling* and it represents the focus of this article. Note that, Gilmer et al. (2017) refers to this task as *readout*.

Performing vertex pooling is made challenging by the fact that different sets of vertex representations corresponding to different graphs may contain different numbers of elements. Furthermore, the elements in a given set are unordered. Therefore one cannot directly apply a feed-forward or recurrent architecture because these require an input lying in a vector space or sequence space respectively. To overcome this challenge most solutions involve mapping the sets of vertex representations to either a vector or sequence space which can then form the input to a feed-forward or recurrent architecture respectively. There exists a wide array of such solutions ranging from computing simple summary statistics such as mean vertex representation to more complex clustering based methods (Ying et al. 2018).

In this article we propose a novel binary graph classification model which performs vertex pooling by mapping a set of vertex representations to an element in a reproducing kernel Hilbert space (RKHS). A RKHS is a function space for which there exists a corresponding kernel function equalling the dot product in this space. Being a function space where the domain of functions in this space is a Euclidean Space, the RKHS in question is of infinite dimension and in turn has high model capacity. However, the infinite nature of this space makes it challenging to work directly in this space. To overcome this challenge, we use the corresponding kernel function which allows us to implicitly compute the dot product in this space without explicitly mapping to the space in question. This is a commonly used strategy known as the *kernel trick*. More specifically, the kernel corresponding to the RKHS is used within a support vector machine (SVM) to perform binary graph classification. A useful feature of the proposed pooling method is that the mapping to a RKHS is parameterized by a scale parameter which controls the degree to which different sets of vertex representations can be discriminated.

The proposed graph classification model is trained in a supervised end-to-end manner where the convolutional layers, the kernel function and SVM parameters are jointly optimized with respect to a regularized classification loss. This approach is distinct from existing kernel-based models which instead use feature engineering or unsupervised learning to define the kernel function and only optimize the parameters of the classification method in a supervised manner (Yanardag and Vishwanathan 2015). Using feature engineering can result in *diagonal dominance* whereby a graph is determined to only be similar to itself, but not to any other graph (Yanardag and Vishwanathan 2015). Although unsupervised learning can overcome this problem and improve performance, the kernel may not be optimal for the task at hand given it was learned in an unsupervised as opposed to supervised manner (Ivanov and Burnaev 2018). The proposed solution of optimizing in an end-to-end manner overcomes these limitations.

The remainder of this paper is structured as follows. “[Related work](#)” section reviews related work on graph kernels and vertex pooling methods. “[Methodology](#)” section describes the proposed graph classification model. “[Evaluation](#)” section presents an evaluation of this model through comparison to 12 baseline models on 4 datasets. Finally, “[Conclusions and future work](#)” section draws some conclusions from this work and discusses possible future research directions.

Related work

In this work we propose a novel vertex pooling method which performs vertex pooling by mapping to a RKHS. In the following two sections we review related work on vertex pooling methods and graph kernels.

Vertex pooling

As discussed in the introduction to this article, existing vertex pooling methods generally map the set of vertex representations to a fixed dimensional vector space or sequence space. The simplest methods for performing vertex pooling compute a summary statistic of the set of vertex representations. Commonly used summary statistics include mean, max and sum (Duvenaud et al. 2015). Despite the simple nature of these methods, a recent study by Luzhnica et al. (2019) demonstrated that in some cases they can outperform more complex methods. Zhang et al. (2018a) proposed a vertex pooling method which first performs a sorting of vertex representations based on the Weisfeiler-Lehman graph isomorphism algorithm. A subset of these vertex representations are then selected based on this ranking, where the size of this subset is a user specified parameter. Li et al. (2016) proposed a vertex pooling method which outputs an element in sequence space. Gilmer et al. (2017) proposed to perform vertex pooling by applying the set2set model from Vinyals et al. (2016). The set2set model maps the set of vertex representations to fixed dimensional vector space representation which is invariant to the order of elements in the set. Ying et al. (2018) proposed a vertex pooling method which uses clustering to iteratively integrate vertex representations and outputs an element in a fixed dimensional vector space. Kearnes et al. (2016) proposed a vertex pooling method which creates a fuzzy histogram of the vertex representations and outputs an element in a fixed dimensional vector space.

Graph kernels

As described in the introduction to this article, existing kernel-based graph classification methods use either feature engineering or unsupervised learning to define the kernel. We now review each of these approaches in turn.

The most common approach for feature engineering kernels is the \mathcal{R} -convolution framework where the kernel function of two graphs is defined in terms of the similarity of their respective substructures (Haussler 1999). This framework is similar to the bag-of-words framework used in natural language processing. Substructures used in the \mathcal{R} -convolution framework to define kernels include graphlets (Shervashidze et al. 2009), shortest path properties (Borgwardt and Kriegel 2005) and random walk properties (Sugiyama and Borgwardt 2015).

The Weisfeiler-Lehman framework is a framework for feature engineering kernels which is inspired by the Weisfeiler-Lehman test of graph isomorphism. In this framework the vertex representations of a given graph are iteratively updated in a similar manner to graph convolution to give a sequence of graphs. A kernel is then defined with respect to this sequence by summing the application of a given kernel, known as the base kernel, to each graph in the sequence. Shervashidze et al. (2011) proposed a family of kernels using this framework by considering a set of base kernels including one which measures the similarity of shortest path properties. Rieck et al. (2019) proposed a kernel using this framework by considering a base kernel which measures the similarity of topological properties.

Kriege et al. (Kriege et al. 2016) proposed another framework for feature engineering kernels known as *assignment kernels* which computes an optimal assignment between graph substructures and sums over a kernel applied to each correspondence in the assignment. The authors proposed a number of kernels using this framework including one based on the Weisfeiler-Lehman graph isomorphism algorithm. Kondor and Pan (2016) proposed a multiscale kernel which considers vertex features plus topological information through the graph Laplacian. Zhang et al. (2018c) proposed a kernel based on the return probabilities of random walks. The authors used an approximation of the kernel function so that the method can be applied to large datasets (Rahimi and Recht 2008).

To overcome the limitations of feature engineering and improve performance, recent works in the field of graph kernels have considered unsupervised learning techniques. These methods generally learn a graph representation in an unsupervised manner and subsequently use this representation to define a kernel. Yanardag and Vishwanathan (2015) proposed a kernel which uses the \mathcal{R} -convolution framework to define a set of substructures and subsequently learns an embedding of these substructures in an unsupervised manner using a word2vec type model. Ivanov and Burnaev (2018) proposed a kernel which determines two graphs to be similar if their vertices have similar neighbourhoods measured in terms of anonymous walks which are a generalization of random walks. Learning is performed in an unsupervised manner using a word2vec type model. Nikolentzos et al. (2017) proposed a graph kernel which first computes sets of vertex representations corresponding to the graphs in question in an unsupervised manner. The similarity of these sets are then computed using the earth mover's distance. The authors noted that these similarities do not yield a positive semidefinite kernel matrix preventing it from being used in some kernel-based classification methods. To overcome this issue the authors use a version of the support vector machine for indefinite kernel matrices.

Similar to Nikolentzos et al. (2017); Wu et al. (2019) proposed a graph kernel which first computes sets of vertex representations corresponding to the graphs in questions in an unsupervised manner. The resulting set of embeddings are in turn used to embed the graph in question by measuring the disturbance distance to sets of embeddings corresponding to random graphs. Finally, this graph representation is used to define a kernel. Nikolentzos et al. (2018) proposed a method that performs an unsupervised clustering of the input graph into components and subsequently learns a kernel function which takes as input these components.

Methodology

The proposed graph classification model consists of the following three steps. In the first step, a sequence of graph convolutional layers are applied to the graph in question to generate a corresponding set of vertex representations. In the second step, this set of vertex representations is mapped to a RKHS. In the final step, graph classification is performed using a SVM. Each of these three steps are described in turn in the first three subsections of this section. In the final subsection we describe how the parameters of each step are optimized jointly in an end-to-end manner. Before that, we first introduce some notation and formally define the problem of graph classification.

A graph is a tuple (V, E) where V is a set of vertices and $E \subseteq (V \times V)$ is a set of edges. Let \mathcal{G} denote the space of graphs. Let $l: V \rightarrow \Sigma$ denote a vertex labelling function. In this work we assume that Σ is a finite set. Let $\mathbb{G} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_n\}$ denote a set of n graphs and $\mathbb{Y} = \{\mathcal{Y}_1, \mathcal{Y}_2, \dots, \mathcal{Y}_n\}$ denote a corresponding set of graph labels. In this work we assume that graph labels take elements in the set $\{0, 1\}$. In this work we consider the problem of binary graph classification where given \mathbb{G} and \mathbb{Y} we wish to learn a map $\mathcal{G} \rightarrow \{0, 1\}$.

Graph convolution layers

A large number of different graph convolutional layers have been proposed. Broadly speaking a graph convolutional layer will update the representation of each vertex in a given graph where this update is a function of the current representation of that vertex plus the representations of its adjacent neighbours. In this section we only briefly review existing graph convolutional layers but the interested reader can find a more indepth analysis in the following review papers (Zhang et al. 2018b; Wu et al. 2019).

Gilmer et al. (2017) showed that many different convolutional layers may be reformulated in terms of a framework called *Message Passing Neural Networks* defined in terms of a message function M and an update function U . In this framework vertex representations are updated according to Eq. 1 where h_v^t denotes the representation of vertex v output from the t -th convolutional layer and $N(v)$ denotes the set of vertices adjacent to v . Each vertex representation h_v^t is an element of \mathbb{R}^m where the dimension m may vary from layer to layer. For the input layer, that is $t = 1$, vertex representations equal a one-hot encoding of the vertex labelling function l and therefore the corresponding dimension is $|\Sigma|$. For all subsequent layers the corresponding dimension is a model hyper-parameter.

$$\begin{aligned} m_v^{t+1} &= \sum_{w \in N(v)} M(h_v^t, h_w^t) \\ h_v^{t+1} &= U(h_v^t, m_v^{t+1}) \end{aligned} \tag{1}$$

In the proposed graph classification model we use the functions M and U originally proposed by Hamilton et al. (2017) and defined in Eq. 2. Here CONCAT is the horizontal vector concatenation operation, \mathbf{W}^t and \mathbf{b}^t are the weights and biases respectively for the t -th convolutional layer, and ReLU is the real valued rectified linear unit non-linearity.

$$\begin{aligned} M(h_v^t, h_w^t) &= h_w^t \\ U(h_v^t, m_v^{t+1}) &= \text{ReLU}(\mathbf{W}^t \cdot \text{CONCAT}(h_v^{t-1}, m_v^{t+1}) + \mathbf{b}^t) \end{aligned} \tag{2}$$

A sequence of two convolutional layers were used in the proposed model. A number of studies have found that the use of two layers empirically gives the best performance (Kipf and Welling 2017). This sequence of layers will map a graph $\mathcal{G}_i = (V, E)$ to a set of $|V|$ points in \mathbb{R}^m where m is the dimension of the final convolutional layer. Since the number of vertices in a graph may vary the number of points in \mathbb{R}^m may in turn vary. Let us denote by Set the space of sets of points in \mathbb{R}^m . Given this, the sequence of convolutions layers defines a map $\mathcal{G} \rightarrow \text{Set}$.

Mapping to RKHS

The output from the sequence of convolutional layers defined in the previous subsection is an element in the space Set . In this section we propose a method for mapping elements in this space to a reproducing kernel Hilbert space (RKHS). We in turn define a kernel between elements in this space.

A Hilbert space is a vector space with an inner product such that the induced norm turns the space into a complete metric space. A positive-semidefinite kernel on a set \mathcal{X} is a function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ such that there exists a feature space \mathcal{H} and a map $\phi : \mathcal{X} \rightarrow \mathcal{H}$ such that $k(x, y) = \langle \phi(x), \phi(y) \rangle$ where $x, y \in \mathcal{X}$ and $\langle \cdot, \cdot \rangle$ denotes the dot product in \mathcal{H} . Equivalently, a function $k : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a kernel if and only if for every subset $\{x_1, \dots, x_q\} \subseteq \mathcal{X}$, the $q \times q$ matrix K with entries $K_{ij} = k(x_i, x_j)$ is positive semi-definite (Schölkopf et al. 2002). Given a kernel k , one can define a map $\mathcal{X} \rightarrow \mathbb{R}^{\mathcal{X}}$ as Eq. 3 where codomain of this map is the space of real valued functions on \mathcal{X} . Such a space is called a function space. Given this, it can be proven that $k(x, y) = \langle k(\cdot, x), k(\cdot, y) \rangle$. By virtue of this property, $\mathbb{R}^{\mathcal{X}}$ is called a *reproducing kernel Hilbert space* (RKHS) corresponding to the kernel k (Schölkopf et al. 2002).

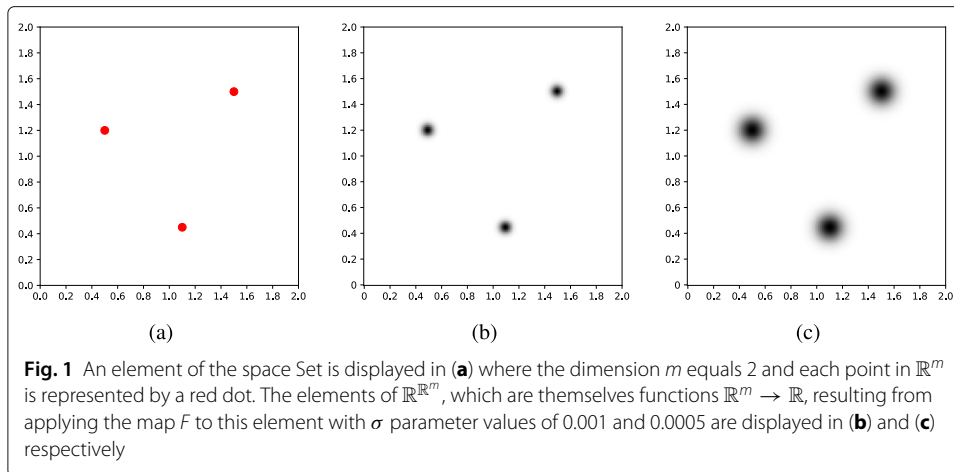
$$x \mapsto k(\cdot, x) \tag{3}$$

Let $k_\sigma^R : \mathbb{R}^m \times \mathbb{R}^m \rightarrow \mathbb{R}$ be the Gaussian kernel function defined in Eq. 4 which is parameterized by $\sigma \in \mathbb{R}_{\geq 0}$.

$$k_\sigma^R(u, v) = \exp(-\|u - v\|_2^2 / 2\sigma^2) \tag{4}$$

Given k_σ^R , we define a map $F : \text{Set} \times \mathbb{R} \rightarrow \mathbb{R}^{\mathbb{R}^m}$ in Eq. 5 where $\mathbb{R}^{\mathbb{R}^m}$ is the space of real valued functions on \mathbb{R}^m . To illustrate this map consider the element of Set displayed in Fig. 1a where the dimension m equals 2. Recall that elements in the space Set correspond to sets of points in \mathbb{R}^m . Figures 1b and c display the elements of $\mathbb{R}^{\mathbb{R}^m}$ resulting from applying the map F to this element of Set with σ parameter values of 0.001 and 0.0005 respectively.

$$F(x, \sigma) = \sum_{v \in x} k_\sigma^R(v, \cdot) \tag{5}$$



The parameter σ of the map F is a scale parameter and may be interpreted as follows. As the value of σ approaches 0, $F(x, \sigma)$ becomes a sum of a set indicator functions applied to x . In this case distinct elements of the space Set map to distinct elements of \mathbb{R}^m where the distance between these functions measured by the L^p norm is greater than zero. On the other hand, as σ approaches ∞ , differences between the functions are gradually smoothed out and in turn the distance between the functions gradually reduces. Therefore, one can view the parameter σ as controlling the discrimination power of the method.

Given the map F defined in Eq. 5, we define the kernel $k_\sigma^L : \mathbb{R}^m \times \mathbb{R}^m \rightarrow \mathbb{R}$ in Eq. 6. Note that, the final equality in this equation follows from the reproducing property of the RKHS related to k_σ^R and the bilinearity of the inner product (Paulsen and Raghupathi 2016). By examination of Eq. 6, we see that the kernel k_σ^L equals the dot product between elements in the codomain of the map F which is an infinite dimensional function space. That is, the kernel allows us to operate in this codomain without the computational complexity of explicitly mapping into it. In Theorem 1 we prove that k_σ^L is a valid positive-semidefinite kernel.

$$\begin{aligned}
 k_\sigma^L(F(x_i, \sigma), F(x_j, \sigma)) &= \langle F(x_i, \sigma), F(x_j, \sigma) \rangle \\
 &= \left\langle \sum_{v \in x_i} k_\sigma^R(v, \cdot), \sum_{u \in x_j} k_\sigma^R(u, \cdot) \right\rangle = \sum_{v \in x_i} \sum_{u \in x_j} k_\sigma^R(v, u)
 \end{aligned}
 \tag{6}$$

Theorem 1 The kernel k_σ^L is a positive-semidefinite kernel.

Proof The kernel k_σ^L is a positive-semidefinite kernel because it is defined in Eq. 6 to equal the dot product in the space \mathbb{R}^m . \square

The kernel k_σ^L has a specific scale which is specified by σ . In order to adopt a multi-scale approach we consider a set of s scales $\Sigma = \{\sigma_1, \dots, \sigma_s\}$ to define a corresponding set of kernels $\{k_{\sigma_1}^L, \dots, k_{\sigma_s}^L\}$. We combine these kernels using a linear combination defined in Eq. 7 where $\{\beta_1, \dots, \beta_s\} \in \mathbb{R}_{\geq 0}^s$. In Theorem 2 we prove that k_Σ^L is a valid positive-semidefinite kernel.

$$k_\Sigma^L(F(x_i), F(x_j)) = \sum_{l=1}^s \beta_l k_{\sigma_l}^L(F(x_i), F(x_j))
 \tag{7}$$

Theorem 2 The kernel k_{Σ}^L is a positive-semidefinite kernel.

Proof The kernel k_{Σ}^L is a positive-semidefinite kernel because it is the sum of positive-semidefinite kernels and the coefficients $\{\beta_1, \dots, \beta_s\}$ are all positive (see proposition 13.1 in Schölkopf et al. (2002)). \square

SVM

Recall that we consider the problem of graph classification whereby given $\mathbb{G} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_n\}$ and $\mathbb{Y} = \{\mathcal{Y}_1, \mathcal{Y}_2, \dots, \mathcal{Y}_n\}$ we wish to learn a map $\mathcal{G} \rightarrow \{0, 1\}$.

Let $f : \text{Set} \rightarrow \mathbb{R}$ be a map from which we obtain a decision function by $\text{sgn}(f)$. That is, if f returns a positive value we classify the graph in question as 1 and otherwise we classify it as 0. We determine a suitable map f lying in the RKHS \mathcal{H} corresponding to the kernel k_{Σ}^L by Eq. 8. Note that, the first term in this sum corresponds to the *soft margin loss* (Schölkopf et al. 2002) and the second term is a regularization term.

$$\hat{f} = \arg \min_{f \in \mathcal{H}} \sum_{i=1}^n \max(0, 1 - y_i f(x_i)) + \lambda \|f\|_{\mathcal{H}}^2 \quad (8)$$

By the *representer theorem* any solution to Eq. 8 can be written in the form of Eq. 9 where $\{\alpha_1, \dots, \alpha_n\} \in \mathbb{R}^n$ (Paulsen and Raghupathi 2016).

$$f(\cdot) = \sum_{j=1}^n \alpha_j k_{\Sigma}^L(\cdot, x_j) \quad (9)$$

Substituting this into Eq. 8 we obtain Eq. 10 where optimization of the function f is performed with respect to $\{\alpha_1, \dots, \alpha_n\} \in \mathbb{R}^n$. Here $K_{ij}^L = k_{\Sigma}^L(x_i, x_j)$, \odot is the elementwise multiplication operator (Hadamard product), $\vec{0}$ is a vector of zeros of size n and $\vec{1}$ is a vector of ones of size n .

$$\begin{aligned} \hat{f} &= \arg \min_{\alpha \in \mathbb{R}^n} \sum_{i=1}^n \max \left(0, 1 - y_i \sum_{j=1}^n \alpha_j k_{\Sigma}^L(x_i, x_j) \right) \\ &\quad + \lambda \left\| \sum_{j=1}^n \alpha_j k_{\Sigma}^L(\cdot, x_j) \right\|_{\mathcal{H}}^2 \\ &= \arg \min_{\alpha \in \mathbb{R}^n} \sum_{i=1}^n \max \left(0, 1 - y_i \sum_{j=1}^n \alpha_j k_{\Sigma}^L(x_i, x_j) \right) \\ &\quad + \lambda \sum_{i,j=1}^n \alpha_i \alpha_j k_{\Sigma}^L(x_i, x_j) \\ &= \arg \min_{\alpha \in \mathbb{R}^n} \left\| \max(\vec{0}, \vec{1} - y \odot K^L \alpha) \right\|_1 \\ &\quad + \lambda \alpha^T K^L \alpha \end{aligned} \quad (10)$$

End-to-end optimization

As described in the previous subsections, the proposed classification model contains three steps with each having corresponding parameters which require optimization with respect to the objective function defined in Eq. 10. The parameters in question are the sets of convolutional layer parameters \mathbf{W}^l and \mathbf{b}^l defined in Eq. 2, the sets of kernel parameters σ_l and β_l defined in Eq. 7, and the set of SVM parameters α_j defined in Eq. 9. All of

these parameters are unconstrained real values apart from the sets of kernel parameters σ_l and β_l which are constrained to be positive real values. As such, the optimization problem in question is a constrained optimization problem. In this work we wish to optimize all the above model parameters jointly in an end-to-end manner. We refer to this as the end-to-end optimization problem. Note that, if only the SVM parameters were optimized and all other parameters were fixed, the optimization problem could be formulated as a quadratic program by taking the dual and solved in closed-form (Schölkopf et al. 2002). This is the most commonly used method for optimizing the parameters of an SVM.

In order to solve the end-to-end optimization problem we use a gradient based optimization method. Such methods are the most commonly used methods for optimizing neural network parameters (Goodfellow et al. 2016). There are two main approaches that can be used to apply a gradient based optimization method to a constrained optimization problem. The first approach is to project the result of each gradient step back into the feasible region. The second approach is to transform the constrained optimization problem into an unconstrained optimization problem and solve this problem. Such a transformation can be achieved using the Karush-Kuhn-Tucker (KKT) method (Nocedal and Wright 2006). In this work we use the former approach. In practice this reduces to passing the parameters σ_l and β_l through the function $\max(\cdot, 0)$ after each gradient step. The above optimization can be used in conjunction with any gradient based optimization method such as stochastic gradient descent. In this work the Adam method was used (Kingma and Ba 2014).

Evaluation

In this section we present an evaluation of the proposed end-to-end graph classification model with respect to current state-of-the-art models. This section is structured as follows. “Implementation details” section provides implementation details for the proposed model. “Baseline methods” section describes the baseline models used to compare the proposed model against. Finally, “Datasets and results” section describes the datasets used in this evaluation and compares the performance of all models on these datasets.

Implementation details

The parameters of the proposed model were initialized as follows. The convolutional layer weights \mathbf{W}^t and biases \mathbf{b}^t in Eq. 2 were initialized using Kaiming initialization (He et al. 2015) and to a value of 0 respectively. The kernel parameters $\{\sigma_1, \dots, \sigma_s\}$ and $\{\beta_1, \dots, \beta_s\}$ and s in Eq. 7 were all initialized to a value of 1.

The model hyper-parameters were set as follows. The dimension of the convolutional hidden layers was set equal to 25. The Adam optimizer learning rate was set to its default value of 0.001 and training was performed for 300 epochs. The hyper-parameters λ in Eq. 10 and s in Eq. 7 were selected from the sets $\{0.0, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0\}$ and $\{1, 2\}$ respectively by considering classification accuracy on a validation set. Larger values for the hyper-parameter s were not considered to ensure scalability of the model to medium sized datasets.

The time and space complexity of classifying a given graph is $O(n)$ where n is the number of graphs in the training dataset. This is a consequence of the summation in Eq. 9 over all training examples. The time and space complexity of performing an update of the method parameters using backprop is $O(n^2)$ because this step computes the complete

kernel matrix K in Eq. 10. Note that, the above time complexity analysis assumes that each element in the kernel matrix K can be computed in constant time. In reality, if we assume that each graph contains m vertices the time complexity of computing each element in K is m^2 (see Eq. 6). In this case the time complexity of classifying a given graph is $O(nm^2)$ while the time complexity of performing an update of the method parameters using backprop is $O(n^2m^2)$.

Baseline methods

As described in the related work section of this paper, existing models for graph classification belong to two main categories of feature engineered kernel and end-to-end deep learning models. For the purposes of this evaluation, we compared the model proposed in this work to baseline models in each of these categories.

A set of 17 baseline models were considered where this set contains 5 feature engineered kernel models and 12 end-to-end deep learning models. We considered so many baseline models to ensure we were comparing to state of the art; many existing models claim to outperform each other so it is difficult to determine which models are in fact state of the art. The end-to-end deep learning baseline models considered in the evaluation are end-to-end models but not are kernel-based models. The proposed model is the first end-to-end kernel-based model for graphs.

In order to cover the breadth of different feature engineered kernel models, we considered three kernel functions in the \mathcal{R} -convolution framework, one kernel function in the Weisfeiler-Lehman framework and one kernel function which uses unsupervised learning. The kernel functions in question are entitled Graphlet by Shervashidze et al. (2009), Shortest Path by Borgwardt and Kriegel (2005), Vertex Histogram by Sugiyama and Borgwardt (2015), Weisfeiler Lehman by Shervashidze et al. (2011) and Pyramid Match by Nikolentzos et al. (2017). These kernel functions are described in the Appendix section of this article. For each kernel function, classification was performed using a Support Vector Machine (SVM) with a kernel matrix precomputed using the kernel function in question. Implementations for the kernel functions were obtained from the *GraKeL* Python library (Siglidis et al. 2020).

The end-to-end deep learning baseline models considered are entitled GCN, GCN-WithJK, GIN, GIN0, GINWithJK, GIN0WithJK, GraphSAGE, GraphSAGEWithJK, Diff-Pool, GlobalAttentionNet, Set2SetNet and SortPool. The architectures of these models are described in the Appendix section of this article. Implementations for these models were obtained from the *PyTorch Geometric* Python library (Fey and Lenssen 2019); these can be downloaded directly from the benchmark section of the PyTorch Geometric website¹. For each end-to-end deep learning baseline model the corresponding model parameters were optimized using the Adam optimizer with the default learning rate of 0.001 and run for 300 epochs. In all cases a negative log likelihood loss function was used. Model hyper-parameters corresponding to the number and dimension of hidden layers were selected from the sets $\{1, 2, 3, 4, 5\}$ and $\{16, 32, 64, 128\}$ respectively by considering the loss on a validation set.

¹https://github.com/rusty1s/pytorch_geometric/tree/master/benchmark

Table 1 Summary statistics for each dataset used in the evaluation

	MUTAG	PTC_MR	BZR_MD	PTC_FM	COX2
Number graphs	188	344	306	349	467
Mean number vertices	17.93	14.29	35.75	14.11	41.22
Mean number edges	19.79	14.69	38.36	14.48	43.45

Datasets and results

To evaluate the proposed graph classification model we considered five commonly used graph classification datasets obtained from the TU Dortmund University graph dataset repository (Kersting et al. 2016)². Summary statistics for each of these datasets are displayed in Table 1. The MUTAG dataset contains graphs corresponding to chemical compounds and the binary classification problem concerns predicting a particular characteristic of the chemical (Debnath et al. 1991). The PTC_MR dataset contains graphs corresponding to chemical compounds and the binary classification problem concerns predicting a carcinogenicity property. The BZR_MD dataset contains graphs corresponding to chemical compounds and the binary classification problem concerns predicting a particular characteristic of the chemical (Sutherland et al. 2003). The PTC_FM dataset contains graphs corresponding to chemical compounds and the binary classification problem and concerns predicting a carcinogenicity property. The COX2 dataset contains graphs corresponding to molecules and the binary classification problem concerns predicting if a given molecule is active or inactive (Sutherland et al. 2003).

Stratified k -folds cross-validation with a k value of 10 was used to split the data into training, validation and testing sets. During each of the k training steps, one of the $k - 1$ folds in the training set was randomly selected to be a validation set and classification accuracy on this set was used to select model hyper-parameters. It has been shown that different training, validation and testing set splits of the data can lead to quite different rankings of graph classification models (Shchur et al. 2018). However averaging the performance of k different splits, as done in this work, helps to reduce this instability. In our analysis the same training, testing and validation splits were used for all graph classification models considered. This is an important point because the performance of a given model may vary as a function of the split used. For each dataset we computed the mean accuracy on the test sets for each method. The results of this analysis are displayed in Table 2. For two of the five datasets, the proposed graph classification model achieved the best mean performance and outperformed most other models by a significant margin. For the remaining three datasets, the proposed method achieved a better mean performance than many but not all baseline methods. These positive results demonstrate the utility of the proposed model. In most cases the proposed model achieved best performance on the validation set with the hyper-parameter s having a value of 2. Recall, from Eq. 7, that this hyper-parameter equals the number of individual kernels integrated by the model. This demonstrates the utility of integrating multiple kernels.

It is important to note that the proposed method was compared against a large number of benchmark methods (17). This makes it challenging for any single method to perform best on all datasets. It is difficult to interpret exactly why one deep learning architecture performs better or worse than another on a particular dataset. However, one limitation

²<https://chrsmrrs.github.io/datasets/>

Table 2 For each dataset, the mean classification accuracy plus standard deviation of 10-fold cross validation for each graph classification model are displayed

Model	MUTAG	PTC_MR	BZR_MD	PTC_FM	COX2
Graphlet	0.86 ±0.05	0.54 ±0.08	0.61 ±0.12	0.57 ±0.07	0.75 ±0.06
Shortest path	0.82 ±0.06	0.55 ±0.08	0.68 ±0.07	0.55 ±0.07	0.77 ±0.06
Vertex histogram	0.85 ±0.05	0.58 ±0.07	0.70 ±0.08	0.58 ±0.06	0.75 ±0.06
Weisfeiler Lehman	0.71 ±0.06	0.59 ±0.07	0.60 ±0.07	0.63 ±0.08	0.65 ±0.08
Pyramid match	0.86 ±0.06	0.54 ±0.06	0.62 ±0.07	0.57 ±0.09	0.73 ±0.06
GCN	0.73 ±0.06	0.57 ±0.03	0.68 ±0.06	0.61 ±0.08	0.79 ±0.03
GCNWithJK	0.73 ±0.07	0.57 ±0.05	0.68 ±0.07	0.62 ±0.08	0.78 ±0.02
GIN	0.82 ±0.07	0.54 ±0.05	0.62 ±0.09	0.57 ±0.07	0.80 ±0.04
GIN0	0.85 ±0.04	0.57 ±0.08	0.63 ±0.13	0.59 ±0.05	0.77 ±0.04
GINWithJK	0.83 ±0.07	0.55 ±0.07	0.61 ±0.15	0.60 ±0.04	0.80 ±0.06
GINOWithJK	0.83 ±0.06	0.54 ±0.07	0.63 ±0.10	0.59 ±0.06	0.82 ±0.05
GraphSAGE	0.72 ±0.06	0.57 ±0.08	0.68 ±0.09	0.61 ±0.06	0.78 ±0.01
GraphSAGEWithJK	0.71 ±0.09	0.56 ±0.04	0.68 ±0.09	0.60 ±0.07	0.77 ±0.01
DiffPool	0.84 ±0.12	0.57 ±0.03	0.69 ±0.07	0.61 ±0.06	0.77 ±0.02
GlobalAttentionNet	0.74 ±0.07	0.56 ±0.05	0.67 ±0.06	0.63 ±0.06	0.77 ±0.01
Set2SetNet	0.73 ±0.07	0.56 ±0.03	0.68 ±0.10	0.62 ±0.06	0.78 ±0.04
SortPool	0.75 ±0.11	0.59 ±0.08	0.66 ±0.09	0.60 ±0.08	0.77 ±0.01
Proposed model	0.87 ±0.06	0.60 ±0.08	0.62 ±0.10	0.62 ±0.06	0.79 ±0.05

Boldface indicates best performing model

of the proposed method that may limit its ability to accurately discriminate is that it only models the distribution of node embeddings and not the position of these nodes in the graph. The recent work by You et al. (2019) suggests position information is important. The DiffPool method which performed best on the BZR_MD dataset actually uses node position information when performing clustering in the pooling step (this is illustrated in Figure 1 of the original paper by Ying et al. (2018)). We hypothesize that position information may not be important for some graph classification tasks while being important for others. This may explain why the proposed method does not uniformly outperform all others. It is also worth noting that the proposed method achieved similar performance to the GIN method on the BZR_MD dataset. In a recent paper by Errica et al. (2019), the authors found the GIN method to achieve best results on a number of datasets. Finally, it is interesting to note that the end-to-end deep learning models did not uniformly outperform the feature engineered kernel models. In fact, the best mean performance on the BZR_MD dataset was achieved by a feature engineered kernel model.

Conclusions and future work

This article proposes a novel kernel-based support vector machine (SVM) for graph classification. Unlike existing kernel-based models, the proposed model is trained in a supervised end-to-end manner whereby the convolutional layers, the kernel function and SVM parameters are jointly optimized. The proposed model outperforms existing deep learning models on a number of datasets which demonstrates the utility of the model.

Despite these positive results, the proposed model is not a suitable candidate solution for all graph classification problems. Like all kernel-based models, the proposed model does not natively scale to large datasets. This is a consequence of the fact that training the model requires computation and storing of the kernel matrix whose size is quadratic

in the number of training examples. This limitation may potentially be overcome by performing an approximation of the kernel function (Rahimi and Recht 2008). The authors plan to investigate this research direction in future work.

Appendix

We briefly describe the kernel functions corresponding to the feature engineered kernel baseline models considered in the work.

Graphlet - This is a kernel in the \mathcal{R} -convolution framework which uses substructures based on Graphlets and was proposed by Shervashidze et al. (2009).

Shortest Path - This is a kernel in the \mathcal{R} -convolution framework which uses substructures based on shortest paths and was proposed by Borgwardt and Kriegel (2005).

Vertex Histogram - This is a kernel in the \mathcal{R} -convolution which uses substructures based on random walks and was proposed by Sugiyama and Borgwardt (2015).

Weisfeiler Lehman - This is a kernel in the Weisfeiler-Lehman framework which uses the Vertex Histogram Kernel as the base kernel (Sugiyama and Borgwardt 2015) and was proposed by Shervashidze et al. (2011).

Pyramid Match - This kernel uses unsupervised learning and was proposed by Nikolentzos et al. (2017).

We briefly describe the architectures corresponding to the end-to-end deep learning baseline models considered in the work. More specific implementation details can be found at the benchmark section of the PyTorch Geometric website.

GCN - This model consists of graph convolutional layers proposed by Kipf and Welling (2017), followed by mean pooling, followed by a non-linear layer, followed by a dropout layer, followed by a linear layer, followed by a softmax layer.

GCNWithJK - This model is equal to GCN but with the addition of jump or skip connections before mean pooling as proposed by Xu et al. (2018).

GIN - This model consists of the graph convolutional layers proposed by Xu et al. (2019), followed by mean pooling, followed by a non-linear layer, followed by a dropout layer, followed by a linear layer, followed by a softmax layer. The convolution layer in question has a parameter ϵ which is learned.

GIN0 - This model is equal to GIN with the exception that the parameter ϵ is not learned and instead is set to a value of 0.

GINWithJK - This model is equal to GIN but with the addition of jump or skip connections before mean pooling as proposed by Xu et al. (2018).

GIN0WithJK - This model is equal to GIN0 but with the addition of jump or skip connections before mean pooling as proposed by Xu et al. (2018).

GraphSAGE - This model consists of the graph convolutional layers proposed by Hamilton et al. (Hamilton et al. 2017), followed by a mean pooling layer, followed by a non-linear layer, followed by a dropout layer, followed by a linear layer, followed by a softmax layer.

GraphSAGEWithJK - This model is equal to GraphSAGE but with the addition of jump or skip connections before mean pooling as proposed by Xu et al. (2018).

DiffPool - This model consists of the graph convolutional layers proposed by Hamilton et al. (2017), followed by the pooling method proposed by Ying et al. (2018), followed by a non-linear layer, followed by a dropout layer, followed by a linear layer, followed by a softmax layer.

GlobalAttentionNet - This model consists of the graph convolutional layers proposed by Hamilton et al. (2017), followed by the pooling layer proposed by Li et al. (2016), followed by a dropout layer, followed by a non-linear layer, followed by a linear layer, followed by a softmax layer.

Set2SetNet - This model consists of the graph convolutional layers proposed by Hamilton et al. (2017), followed by the pooling layer proposed by Vinyals et al. (2016), followed by a non-linear layer, followed by a dropout layer, followed by a linear layer, followed by a softmax layer.

SortPool - This model consists of the graph convolutional layers proposed by Hamilton et al. (2017), followed by the pooling layer proposed by Zhang et al. (2018a), followed by a non-linear layer, followed by a dropout layer, followed by linear layer, followed by a softmax layer.

Abbreviations

SVM: Support vector machine; RKHS: Reproducing kernel Hilbert space; KKT: Karush-Kuhn-Tucker

Acknowledgements

The author would like to acknowledge the many useful discussions he had with Bertrand Gauthier concerning kernel methods.

Author's contributions

Padraig Corcoran was solely responsible for all work presented in this paper. The author(s) read and approved the final manuscript.

Funding

Nothing to declare.

Availability of data and materials

All data used in this work is publicly available from the TU Dortmund University graph dataset repository.

Competing interests

The author declares that they have no competing interests.

Received: 27 February 2020 Accepted: 7 July 2020

Published online: 22 July 2020

References

- Borgwardt KM, Kriegel H-P (2005) Shortest-path kernels on graphs. In: IEEE International Conference on Data Mining, Houston. p 8
- Chen V, Varma P, Krishna R, Bernstein M, Re C, Fei-Fei L (2019) Scene graph prediction with limited labels. In: International Conference on Computer Vision. <https://doi.org/10.1109/iccv.2019.00267>
- Debnath AK, Lopez de Compadre RL, Debnath G, Shusterman AJ, Hansch C (1991) Structure-activity relationship of mutagenic aromatic and heteroaromatic nitro compounds. correlation with molecular orbital energies and hydrophobicity. *J Med Chem* 34(2):786–797
- Duvenaud D, Maclaurin D, Iparraguirre J, Bombarell R, Hirzel T, Aspuru-Guzik A, Adams R (2015) Convolutional networks on graphs for learning molecular fingerprints. In: Advances in Neural Information Processing Systems, Montreal. pp 2224–2232
- Errica F, Podda M, Bacciu D, Micheli A (2019) A fair comparison of graph neural networks for graph classification. arXiv preprint arXiv:1912.09893
- Fey M, Lenssen JE (2019) Fast graph representation learning with PyTorch Geometric. In: ICLR Workshop on Representation Learning on Graphs and Manifolds, New Orleans
- Franceschi L, Niepert M, Pontil M, He X (2019) Learning discrete structures for graph neural networks. In: International Conference on Machine Learning, California. pp 1972–1982
- Gagarin A, Corcoran P (2018) Multiple domination models for placement of electric vehicle charging stations in road networks. *Comput Oper Res* 96:69–79
- Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE (2017) Neural message passing for quantum chemistry. In: International Conference on Machine Learning, Sydney. pp 1263–1272
- Goodfellow I, Bengio Y, Courville A (2016) Deep Learning. MIT press, Massachusetts
- Hamilton W, Ying Z, Leskovec J (2017) Inductive representation learning on large graphs. In: Advances in Neural Information Processing Systems, California. pp 1024–1034
- Hausler D (1999) Convolution kernels on discrete structures. Technical report
- He K, Zhang X, Ren S, Sun J (2015) Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In: Proceedings of the IEEE International Conference on Computer Vision, Las Condes. pp 1026–1034
- Ivanov S, Burnaev E (2018) Anonymous walk embeddings. In: International Conference on Machine Learning, vol. 80, Stockholmssmassan. pp 2191–2200
- Kearnes S, McCloskey K, Berndl M, Pande V, Riley P (2016) Molecular graph convolutions: moving beyond fingerprints. *J Computer-aided Mol Des* 30(8):595–608

- Kersting K, Kriege NM, Morris C, Mutzel P, Neumann M (2016) Benchmark Data Sets for Graph Kernels. <http://graphkernels.cs.tu-dortmund.de>. Accessed 07 Jul 2020
- Kingma DP, Ba J (2014) Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980
- Kipf TN, Welling M (2017) Semi-supervised classification with graph convolutional networks. In: International Conference on Learning Representations, Toulon
- Kondor R, Pan H (2016) The multiscale laplacian graph kernel. In: Advances in Neural Information Processing Systems, Barcelona. pp 2990–2998
- Kriege NM, Giscard P-L, Wilson R (2016) On valid optimal assignment kernels and applications to graph classification. In: Advances in Neural Information Processing Systems, Barcelona. pp 1623–1631
- Krishna R, Zhu Y, Groth O, Johnson J, Hata K, Kravitz J, Chen S, Kalantidis Y, Li L-J, Shamma DA, et al (2017) Visual genome: Connecting language and vision using crowdsourced dense image annotations. *Int J Comput Vis* 123(1):32–73
- Landrieu L, Simonovsky M (2018) Large-scale point cloud semantic segmentation with superpoint graphs. In: The IEEE Conference on Computer Vision and Pattern Recognition. <https://doi.org/10.1109/cvpr.2018.00479>
- Li Y, Tarlow D, Brockschmidt M, Zemel R (2016) Gated graph sequence neural networks. In: International Conference on Learning Representations, San Juan
- Luzhnica E, Day B, Liò P (2019) On graph classification networks, datasets and baselines. In: ICML Workshop on Learning and Reasoning with Graph-Structured Representations, California
- Nikolentzos G, Meladianos P, Tixier AJ-P, Skianis K, Vazirgiannis M (2018) Kernel graph convolutional neural networks. In: International Conference on Artificial Neural Networks. Springer. pp 22–32
- Nikolentzos G, Meladianos P, Vazirgiannis M (2017) Matching node embeddings for graph similarity. In: AAAI Conference on Artificial Intelligence, California
- Noce J, Wright S (2006) Numerical Optimization. Springer
- Paulsen VI, Raghupathi M (2016) An Introduction to the Theory of Reproducing Kernel Hilbert Spaces, vol. 152. Cambridge University Press
- Rahimi A, Recht B (2008) Random features for large-scale kernel machines. In: Advances in Neural Information Processing Systems, British Columbia. pp 1177–1184
- Rieck B, Bock C, Borgwardt K (2019) A persistent weisfeiler-lehman procedure for graph classification. In: International Conference on Machine Learning, California. pp 5448–5458
- Schölkopf B, Smola AJ, Bach F, et al (2002) Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond. MIT press, Massachusetts
- Shchur O, Mumme M, Bojchevski A, Günnemann S (2018) Pitfalls of graph neural network evaluation. arXiv preprint arXiv:1811.05868
- Shervashidze N, Schweitzer P, Leeuwen E. J. v., Mehlhorn K, Borgwardt KM (2011) Weisfeiler-lehman graph kernels. *J Mach Learn Res* 12(Sep):2539–2561
- Shervashidze N, Vishwanathan S, Petri T, Mehlhorn K, Borgwardt K (2009) Efficient graphlet kernels for large graph comparison. In: Artificial Intelligence and Statistics, Florida. pp 488–495
- Siglidis G, Nikolentzos G, Linnios S, Giatsidis S, Skianis K, Vazirgiannis M (2020) Grakel: A graph kernel library in python. *J Mach Learn Res* 21(54):1–5
- Sugiyama M, Borgwardt K (2015) Halting in random walk kernels. In: Advances in Neural Information Processing Systems, Montreal. pp 1639–1647
- Sutherland JJ, O'Brien LA, Weaver DF (2003) Spline-fitting with a genetic algorithm: A method for developing classification structure-activity relationships. *J Chem Inf Comput Sci* 43(6):1906–1915
- Vinyals O, Bengio S, Kudlur M (2016) Order matters: Sequence to sequence for sets. In: International Conference on Learning Representations, San Juan, Puerto Rico
- Wu Z, Pan S, Chen F, Long G, Zhang C, Yu PS (2019) A comprehensive survey on graph neural networks. arXiv preprint arXiv:1901.00596
- Wu L, Yen IE-H, Zhang Z, Xu K, Zhao L, Peng X, Xia Y, Aggarwal C (2019) Scalable global alignment graph kernel using random features: From node embedding to graph embedding. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. ACM, Alaska. pp 1418–1428
- Xu K, Hu W, Leskovec J, Jegelka S (2019) How powerful are graph neural networks?. In: International Conference on Learning Representations, New Orleans
- Xu K, Li C, Tian Y, Sonobe T, Kawarabayashi K.-i., Jegelka S (2018) Representation learning on graphs with jumping knowledge networks. In: International Conference on Machine Learning, vol. 80, Stockholm. pp 5453–5462
- Xu D, Zhu Y, Choy C, Fei-Fei L (2017) Scene graph generation by iterative message passing. In: The IEEE Conference on Computer Vision and Pattern Recognition. <https://doi.org/10.1109/cvpr.2017.330>
- Yan S, Xiong Y, Lin D (2018) Spatial temporal graph convolutional networks for skeleton-based action recognition. In: AAAI Conference on Artificial Intelligence, New Orleans
- Yanardag P, Vishwanathan S (2015) Deep graph kernels. In: ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp 1365–1374. <https://doi.org/10.1145/2783258.2783417>
- Ying Z, You J, Morris C, Ren X, Hamilton W, Leskovec J (2018) Hierarchical graph representation learning with differentiable pooling. In: Advances in Neural Information Processing Systems, Montreal. pp 4800–4810
- You J, Ying R, Leskovec J (2019) Position-aware graph neural networks. In: Chaudhuri K, Salakhutdinov R (eds). Proceedings of the 36th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 97, Long Beach, California. pp 7134–7143
- You J, Ying R, Ren X, Hamilton W, Leskovec J (2018) Graphrnn: Generating realistic graphs with deep auto-regressive models. In: International Conference on Machine Learning, Stockholm. pp 5694–5703
- Zhang M, Cui Z, Neumann M, Chen Y (2018a) An end-to-end deep learning architecture for graph classification. In: AAAI Conference on Artificial Intelligence, New Orleans
- Zhang Z, Cui P, Zhu W (2018b) Deep learning on graphs: A survey. arXiv preprint arXiv:1812.04202
- Zhang Z, Wang M, Xiang Y, Huang Y, Nehorai A (2018c) Retgk: Graph kernels based on return probabilities of random walks. In: Advances in Neural Information Processing Systems, Montreal. pp 3964–3974

Publisher's Note

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