



Learning subtree pattern importance for Weisfeiler-Lehman based graph kernels

Dai Hai Nguyen^{1,2} · Canh Hao Nguyen² · Hiroshi Mamitsuka^{2,3}

Received: 22 November 2020 / Revised: 23 March 2021 / Accepted: 3 May 2021 /

Published online: 13 June 2021

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2021

Abstract

Graph is an usual representation of relational data, which are ubiquitous in many domains such as molecules, biological and social networks. A popular approach to learning with graph structured data is to make use of graph kernels, which measure the similarity between graphs and are plugged into a kernel machine such as a support vector machine. Weisfeiler-Lehman (WL) based graph kernels, which employ WL labeling scheme to extract subtree patterns and perform node embedding, are demonstrated to achieve great performance while being efficiently computable. However, one of the main drawbacks of a general kernel is the decoupling of kernel construction and learning process. For molecular graphs, usual kernels such as WL subtree, based on substructures of the molecules, consider all available substructures having the same importance, which might not be suitable in practice. In this paper, we propose a method to learn the weights of subtree patterns in the framework of WWL kernels, the state of the art method for graph classification task (Togninalli et al., in: *Advances in Neural Information Processing Systems*, pp. 6439–6449, 2019). To overcome the computational issue on large scale data sets, we present an efficient learning algorithm and also derive a generalization gap bound to show its convergence. Finally, through experiments on synthetic and real-world data sets, we demonstrate the effectiveness of our proposed method for learning the weights of subtree patterns.

Keywords Graph kernel · Optimal transport · Weisfeiler Lehman scheme

Editors: Annalisa Appice, Sergio Escalera, Jose A. Gamez, Heike Trautmann.

✉ Dai Hai Nguyen
hai@k.u-tokyo.ac.jp

Canh Hao Nguyen
canhhao@kuicr.kyoto-u.ac.jp

Hiroshi Mamitsuka
mami@kuicr.kyoto-u.ac.jp

¹ Graduate School of Frontier Sciences, The University of Tokyo, 5-1-5 Kashiwa-no-ha, Kashiwa, Chiba 277-8561, Japan

² Bioinformatics Center, Institute for Chemical Research, Kyoto University, Uji 611-0011, Japan

³ Department of Computer Science, Alato University, Espoo 02150, Finland

1 Introduction

Graphs are natural data structures, which appear in various domains such as bioinformatics (Sharan and Ideker 2006), cheminformatics (Trinajstić 2018), social network analysis (Scott 2011) and so on, where nodes (vertices) represent objects and edges represent the relations between them. A popular approach to learning with graph structured data is to make use of graph kernels. Essentially, a graph kernel is a measure of the similarity between two graphs and must satisfy two fundamental requirements of being a valid kernel: (1) symmetric and (2) positive semi-definite (PSD). Furthermore, the requirements of designing a graph kernel are: it should capture the semantic inherent in the graph structures (e.g. substructures of different levels), and it must be efficiently computable (Vishwanathan et al. 2010).

A number of graph kernels have been proposed in literature such as random walk (Kashima 2003), shortest path (Borgwardt and Kriegel 2005), Weisfeiler-Lehman (WL) subtree (Shervashidze and Borgwardt 2009) kernels, just to name a few. Most of them are based on \mathcal{R} -Convolution framework (Haussler 1999), which decomposes two graphs into substructures and adds up the similarities between their substructures to compute kernel values. Different graph kernels are defined under different ways of decomposition (types of substructures). For instance, the substructures can be random walks (Kashima 2003), shortest paths (Borgwardt and Kriegel 2005) or subtree patterns (Vishwanathan et al. 2010). Among these, WL subtree kernels have been shown to achieve great prediction performance while being efficiently computable. The key point is that it simply employs a WL based color refinement scheme to embed each node in a given graph into a vector of WL labels, which correspond to *subtree patterns* of the graphs. Then, the kernel between two graphs is defined as the sum of all pairwise similarities between any two node embeddings of the two graphs.

Following a different approach, WL based optimal assignment (WL-OA) kernel (Kriege et al. 2016) assigns one node embedding of one graph to one embedding of the other such that the total similarities between assigned node embeddings is maximized. This is also known as optimal assignment problem in combinatorial mathematics (Munkres 1957). In a similar vein, Wasserstein WL (WWL, Togninalli et al. 2019) uses optimal transport (OT), also known as Wasserstein distance (Villani 2008), for measuring the distance between two graphs based on their WL node embeddings. The distance is then converted into a similarity matrix through Laplacian kernel. Furthermore, both of these similarity matrices are shown to be valid kernels due to the hierarchy property of WL labels (see Kriege et al. 2016 and Togninalli et al. 2019 for more details).

One of the main drawbacks of these kernels is that they are predefined feature extraction without learning the importance of substructures to the problem. This results in the decoupling of data representation and learning process. In these kernels, substructures are given the same weights. However, for the problems such as molecule classification, it is known that only subparts of the molecules are responsible for their properties. Therefore, we wish to be able to give weights to their substructures to have higher classification performance and model interpretation. Based on this motivation, we propose a model to learn the weights of *subtree patterns* (extracted by WL labeling scheme). Our work extends WWL kernels (Togninalli et al. 2019) by formulating an OT based distance as a parametric function of subtree pattern weights before converting into kernels. We also propose an efficient stochastic learning algorithm to estimate the weights and derive a generalization gap bound for the algorithm. Finally, through experiments on synthetic and four real-world data

sets, we show that learning important subtree patterns by our proposed method can lead to more accurate predictive performance and extract important patterns which enhance the classification results.

The remainder of the paper is organized as follows: in Sect. 2, we review graph kernels which are based on WL labeling scheme, including WL subtree, WL-OA and WWL kernels. In Sect. 3, we present our method that parameterizes the Wasserstein distance between two graphs with WL labeling scheme as a function of subtree patterns and present the stochastic algorithm for learning parameters of the function. In Sect. 4, we derive a generalization bound for the learning algorithm. In Sect. 5, experimental results on the synthetic and real-world data sets are provided. Finally, we conclude by summarizing this work and discussing possible extensions in Sect. 6.

2 Related work

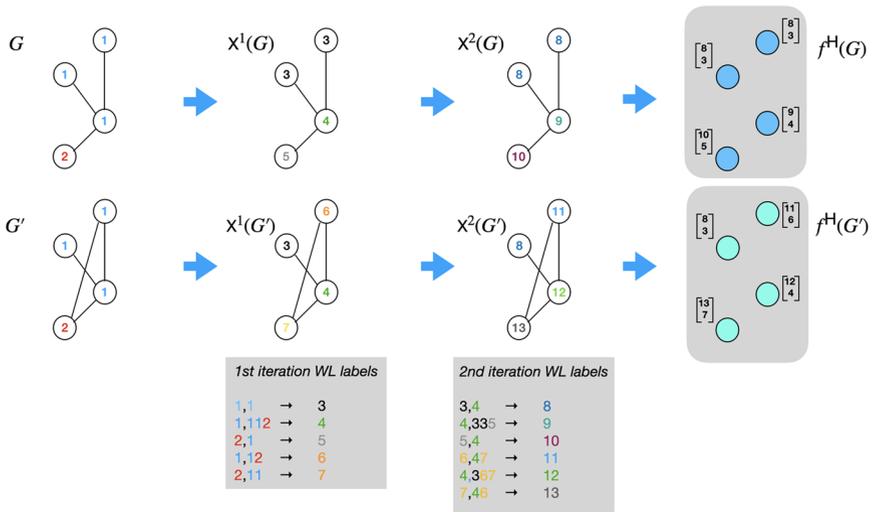
In this paper, we consider the binary classification problem for graph structured data: given a collection of labeled graphs $(g_i, y_i), i = 1, \dots, n$ (where n is the number of examples) drawn from an unknown joint distribution \mathcal{P} over $\mathcal{G} \times \{-1, 1\}$, where \mathcal{G} is a space of graphs. We wish to learn a classifier $h : \mathcal{G} \rightarrow \{-1, 1\}$, which is based on a similarity function $K : \mathcal{G} \times \mathcal{G} \rightarrow [-1, 1]$. If K is symmetric and positive semi-definite (PSD), it is called a valid kernel.

There are a number of proposed kernels on graphs, see (Kashima 2003; Borgwardt and Kriegel 2005; Shervashidze and Borgwardt 2009). In general, they are defined based on \mathcal{R} -Convolution framework (Haussler 1999), that is, each graph $g \in \mathcal{G}$ is decomposed into substructures, and a kernel value $K(g, g')$ is defined as a sum of pairwise similarities between their substructures. In fact, many graph kernels can be considered as instances of the \mathcal{R} -Convolution framework under different decomposition into substructures. The substructures can be random walks (Kashima 2003), shortest paths (Borgwardt and Kriegel 2005) or circle subtrees (Shervashidze and Borgwardt 2009). Among these, Weisfeiler-Lehman (WL) subtree kernels (Shervashidze and Borgwardt 2009) and its variants have been shown to achieve great performance for the graph classification tasks. In this work, we focus on WL based kernels and will review them in the following subsections.

2.1 Weisfeiler-Lehman (WL) scheme for node embeddings

Weisfeiler-Lehman (WL) subtree kernels (Shervashidze and Borgwardt 2009) are based on an iterative colour refinement (also known as WL labeling scheme) and have been shown to achieve great performance for graph classification task. For each node of a given graph, the WL labeling scheme creates a sequence of ordered strings by the aggregation of the labels of the node and its neighbors; these strings are then hashed or indexed to produce compressed updated node labels or new indices. If the iteration of the scheme is increased, these obtained labels represent increasingly broader neighborhood of each node. More specifically, for a graph $G = (V, E)$ with initial labels $\ell_0(v)$ for $v \in V$ and let H be the number of WL iterations, we can define a sequence of refined labels $(\ell_0, \ell_1, \dots, \ell_H)$, where ℓ_{h+1} is obtained from ℓ_h by the following procedure: $\ell_{h+1}(v) = \text{hash}(\ell_h(v), \mathcal{N}_h(v))$, where $\mathcal{N}_h(v)$ denotes a lexicographically sorted sequence of labels of v 's neighbors at iteration h and the hash function is to create a updated compressed node label for v . We use perfect hashing

(a) WL node embedding scheme with $H = 2$



(b) WL based graph kernels

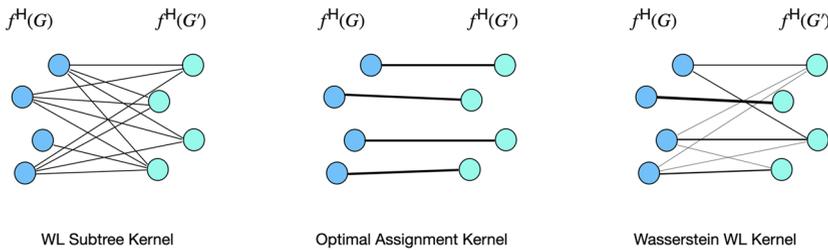


Fig. 1 **a** Illustration of Weisfeiler-Lehman (WL) node embedding scheme with two iterations ($H = 2$). **b** Illustration of different WL based graph kernels: WL Subtree kernel is computed by considering all pairwise similarities between node embeddings of two graphs; WL optimal assignment kernel takes the similarities of only aligned node embeddings into account; For Wasserstein WL kernel, one node embedding of one graph can be coupled with multiple node embeddings of the other graph for computing the kernel

for the hash function, as in Shervashidze and Borgwardt (2009), ensuring two nodes at iteration $h + 1$ have the same label if and only if their label and those of their neighbors at iteration h are the same. Throughout the rest of paper, we denote the set of WL labels at iteration h as Σ^h , for $h = 1, \dots, H$, and the set of all WL labels as $\Sigma = \cup_{h=1}^H \Sigma^h$. The WL labeling scheme is illustrated in Fig. 1a.

Based on WL labeling scheme, a WL node embedding scheme was proposed to generate node embeddings from node labels of the graphs (Shervashidze and Borgwardt 2009).

Definition 1 (WL based node embedding scheme, Shervashidze and Borgwardt 2009). Let $G = (V, E)$ and let H be the number of WL iterations. For every $h \in \{1, \dots, H\}$, we define the node embedding $x^h(v)$ of a node $v \in V$ and graph embedding $\mathbf{X}^h(G)$ of G at iteration h as follows:

$$x^h(v) = \ell^h(v), \mathbf{X}^h(G) = [x^h(v_1), \dots, x^h(v_{n_v})]^T \quad (1)$$

Then the graph embedding of G can be defined as:

$$\begin{aligned} f^H(G) : \mathcal{G} &\mapsto \Sigma^{n_v \times H} \\ G &\mapsto [\mathbf{X}^1(G)^T, \dots, \mathbf{X}^H(G)^T]^T \end{aligned}$$

where n_v is the number of nodes in G . With the WL node embedding scheme above, we are ready to introduce some notations which will be used throughout the rest of paper.

Notations Let $D_h^{\text{Ham}}(f^h(G), f^h(G'))$ be the Hamming distance matrix where each entry is the *normalized Hamming distance* between the corresponding node embeddings of G and G' at iteration h , defined as:

$$d_h^{\text{Ham}}(u, v) = \frac{1}{h} \sum_{i=1}^h d_i^{\text{disc}}(u, v), d_i^{\text{disc}}(u, v) = \begin{cases} 0 & \text{if } u_i = v_i \\ 1 & \text{otherwise} \end{cases} \quad (2)$$

where u and v denote two node embeddings, $d_h^{\text{Ham}}(u, v)$ denotes the normalized Hamming distance between u and v at iteration h and $d_i^{\text{disc}}(u, v)$ denotes the *discrete distance* between u and v at iteration i . Similarly, let $D_h^{\text{Disc}}(\mathbf{X}^h(G), \mathbf{X}^h(G'))$ be the discrete distance matrix where each entry is the discrete distance between the corresponding node embeddings of G and G' at iteration h . It is easy to see that $[D_h^{\text{Ham}}]_{ij} \in [0, 1]$ and $[D_h^{\text{Disc}}]_{ij} \in \{0, 1\}$. We also define a base kernel which corresponds to the averaged number of feature shared by two node embeddings as:

$$k_h(u, v) = \frac{1}{h} \sum_{i=1}^h \mathbb{1}(u_i = v_i) \quad (3)$$

It is easy to see that $k_h(u, v) = 1 - d_h^{\text{Ham}}(u, v)$. Thanks to WL node embedding scheme, a graph can be represented as a point cloud or a set of node embeddings. Measuring the similarity (or dissimilarity) between two graphs boils down to measuring the similarity (or dissimilarity) between two sets of embeddings. Each node embedding captures information about the neighborhood of the corresponding node (or a rooted *subtree pattern*). In the following subsections, we will review different WL based graph kernels derived from different ways of comparing their sets of WL node embeddings.

2.2 WL subtree kernels

WL subtree kernels (Shervashidze and Borgwardt 2009) simply employ the aforementioned WL labeling scheme to extract subtree patterns, which represent the neighborhood of each node in the graph up to a given distance (or number of hops H). Essentially WL subtree kernel counts the number of common WL labels. In the context of WL node embedding scheme, it can be computed by summing all pairwise similarities between node embeddings of two graphs. More formally, for two graphs G and G' with two sets of node embeddings $f^H(G)$ and $f^H(G')$, respectively, the WL subtree kernel value between them can be defined as:

$$\mathbf{K}^{\text{WL}}(G, G') = \sum_{\mathbf{x} \in f^H(G)} \sum_{\mathbf{y} \in f^H(G')} k_H(\mathbf{x}, \mathbf{y}) \quad (4)$$

It is obvious that as the base kernel $k_H(\mathbf{x}, \mathbf{y})$ (defined in Eq. (3)) is equal to the number of WL labels (subtree patterns) shared by two node embeddings, the kernel $\mathbf{K}^{\text{WL}}(G, G')$ is equal to the total number of WL subtree patterns shared by two graphs.

2.3 WL-based optimal assignment kernels

Optimal assignments are natural measures of similarity between two sets of points. In particular, for two sets of points, the goal is to assign one point of one set to another point of the other set (one-to-one correspondence) such that the sum of similarities between assigned points is maximized. Finding such an optimal alignment or bijection is also known as the well-studied assignment problem in combinatorial optimization (Munkres 1957). However, a challenge is how to design a valid kernel based on optimal assignments.

Kriege et al. (2016) introduced a restricted class of kernels, called *strong kernels*, that guarantees the construction of valid optimal assignment based kernels. An important result is that the strong kernels give rise to hierarchies defined on the domain of kernels. Based on this, the authors proposed WL-based optimal assignment (WL-OA) kernels with the WL node embedding scheme. WL-OA kernels employ the base kernel, defined in Eq. (3), which satisfies the requirement of being a strong kernel as the sequence of refined WL labels $(\ell_0, \ell_1, \dots, \ell_H)$ gives rise to a family of nested subsets, which can be represented by a hierarchy. Consequently, WL-OA kernels are valid. Formally, given two graphs G and G' with two sets of node embeddings $f^H(G)$ and $f^H(G')$, respectively, the WL-OA kernel value between them is defined as:

$$\mathbf{K}^{\text{OA}}(G, G') = \max_{B \in \mathcal{B}(f^H(G), f^H(G'))} \sum_{(x,y) \in B} k_H(\mathbf{x}, \mathbf{y}) \quad (5)$$

where $\mathcal{B}(f^H(G), f^H(G'))$ is the set of all bijections between two sets $f^H(G)$ and $f^H(G')$. To apply this kernel to graphs of different number of nodes, we can fill up the graph with smaller number of nodes, says $f^H(G')$, by new node embeddings z with $k(\mathbf{x}, \mathbf{z}) = 0$ for all $\mathbf{x} \in f^H(G)$ without changing the result. It is worth noting that the WL-OA kernels take the similarities of aligned node embeddings into account, while the WL subtree kernels consider all pairwise similarities.

2.4 Wasserstein WL kernels

Optimal transport (OT), also known as Wasserstein distance function (Villani 2008), has gained much attraction in machine learning community as a powerful tool for the comparison of two probability distributions. The naive computation of this distance between two discrete measures, e.g. point clouds, involves solving transport problem. Formally, let $X = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ and $Y = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$ be two sets of points, where m and n denote the size of two sets X and Y , respectively; $\mathbf{p} \in \mathbb{R}^m$ and $\mathbf{q} \in \mathbb{R}^n$ are two discrete probability distributions over X and Y , respectively. We use d_{ij} to denote the distance between \mathbf{x}_i and \mathbf{y}_j , e.g. the squared Euclidean distance $d_{ij} = \|\mathbf{x}_i - \mathbf{y}_j\|_2^2$. The Wasserstein distance is formulated as a linear program over the *transportation matrix* (or joint probability) $\mathbf{P} \in \mathbb{R}^{m \times n}$.

$$\begin{aligned}
\mathcal{W}_1(X, Y, d) &= \min_{\mathbf{P}_{ij}} \sum_{i=1}^m \sum_{j=1}^n \mathbf{P}_{ij} d_{ij} \\
\text{subject to } \sum_{i=1}^m \mathbf{P}_{ij} &= \mathbf{q}_j, \forall j \in [1, n] \\
\sum_{j=1}^n \mathbf{P}_{ij} &= \mathbf{p}_i, \forall i \in [1, m]
\end{aligned} \tag{6}$$

With WL node embedding scheme to generate node embeddings for graphs, Togninalli et al. (2019) evaluated the pairwise Wasserstein distance between graphs with the normalised Hamming (2) as the ground distance. Then, Wasserstein WL (WWL) kernel is defined as an instance of Laplacian kernels, see Eq. (7).

$$\begin{aligned}
\mathbf{D}^{\text{WWL}}(G, G') &= \mathcal{W}_1(f^H(G), f^H(G'), d_H^{\text{Ham}}) \\
\mathbf{K}^{\text{WWL}}(G, G') &= e^{-\gamma \mathbf{D}^{\text{WWL}}(G, G')}
\end{aligned} \tag{7}$$

where γ is a hyperparameter.

In general cases, it is not necessarily possible to derive a valid kernel from the Wasserstein distance. However, thanks to the special property of WL labels and normalized Hamming distance (2), \mathbf{D}^{WWL} was shown to be conditionally negative definite (CND), resulting in the validity of \mathbf{K}^{WWL} , by proving the following lemma (see Togninalli et al. 2019 for its proof):

Lemma 1 *If a transportation matrix P^H is optimal solution of (6) with the ground distance d_H^{Ham} (2) between node embeddings at iteration H , then we have the two following claims:*

1. P^H is also optimal solution of (6) with the discrete distance d_H^{Disc} between H -iteration values.
2. P^H is also optimal solution of (6) with the normalised Hamming distance d_{H-1}^{Ham} between node embeddings at iteration $H - 1$.

Let P^* be the optimal solution of (6) for D_H^{Ham} . From the above lemma, it is also the optimal solution for D_h^{Disc} , $h = 1, \dots, H$. The Wasserstein distance between two graphs G and G' in (7) can be simplified as follows:

$$\mathbf{D}^{\text{WWL}}(G, G') = \frac{1}{H} \sum_{h=1}^H \mathcal{W}_1(\mathbf{X}^h(G), \mathbf{X}^h(G'), d_h^{\text{Disc}}) \tag{8}$$

The Eq. (8) is a sum of OT distances with the discrete distances as ground metrics, which are CND. Therefore, the sum is also CND, leading to the validity of the similarity matrix \mathbf{K}^{WWL} .

3 Incorporating subtree pattern importance into WL based graph kernels

One of the main limitations of kernels is the decoupling of data representation and learning process, that is, the kernel must be predefined prior to learning, leading to limited predictive performance. Furthermore, in prediction tasks for molecular data, the output might be determined by the presence of a few important substructures, while these kernels contain all substructures with equal weights. Motivated by this drawback, in this paper we address the problem of incorporating subtree pattern weights for WWL kernel (Togninalli et al. 2019). To this end, we aim to learn new kernels from a parametric form of Wasserstein distance taking into account subtree pattern weights (8), and learn these weights from data optimally for the task.

3.1 Parametric form of Wasserstein distance with subtree pattern weights

To derive a parametric form of the distance function (8), we rely on the following simple observation:

Lemma 2 *Let S be a set of elements, $X = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ and $Y = \{\mathbf{y}_1, \dots, \mathbf{y}_n\}$ ($X, Y \subseteq S$) be two multiset of S of m and n samples, respectively, the Wasserstein distance between them with the discrete distance as the ground metric is determined by:*

$$W_1(X, Y) = 1 - \sum_{v \in S} \min(\mu_X(v), \mu_Y(v)) \tag{9}$$

where $\mu_X(v)$ denotes the mass density function of the multiset X with $v \in S$.

Applying this lemma to Eq. (8), we have:

$$D^{WWL}(G, G') = 1 - \frac{1}{H} \sum_{h=1}^H \sum_{v \in \Sigma^h} \min(\mu_{\mathbf{X}_h(G)}(v), \mu_{\mathbf{X}_h(G')}(v)) \tag{10}$$

Our idea is to give each substructure or WL label v a nonnegative weight $w_v \in \mathbb{R}_{\geq 0}$ for its importance to the problem, so the parametric form of Eq. (10) is defined as follows:

$$b - \frac{1}{H} \sum_{h=1}^H \langle \mathbf{w}_h, \mathbf{z}_h(G, G') \rangle \tag{11}$$

where $\mathbf{w}_h, \mathbf{z}_h(G, G') \in \mathbb{R}^{|\Sigma^h|}$ are the vectors of entries w_v and $\min(\mu_{\mathbf{X}_h(G)}(v), \mu_{\mathbf{X}_h(G')}(v))$, respectively, for $v \in \Sigma^h$; b is a constant to ensure that the value of parametric function is nonnegative. In vector form, this can be expressed as:

$$d_{\mathbf{W}}(G, G') = b - \langle \mathbf{W}, \mathbf{Z}(G, G') \rangle \tag{12}$$

where $\begin{cases} \mathbf{Z}(G, G') = \frac{1}{H} [\mathbf{z}_1(G, G')^T, \dots, \mathbf{z}_H(G, G')^T]^T \\ \mathbf{W} = [\mathbf{w}_1^T, \dots, \mathbf{w}_H^T]^T \end{cases}$

The parametric form (12) is a linear function with respect to the parameter vector $\mathbf{W} \in \mathbb{R}^d$ ($d = |\Sigma^1| + \dots + |\Sigma^H|$) and $\mathbf{Z}(G, G')$ is considered as a feature vector of a pair of

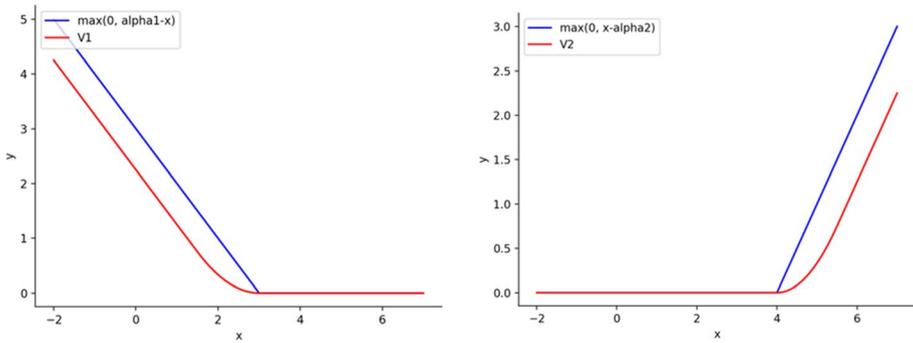


Fig. 2 Illustration of Hinge loss $\max(0, \alpha_1 - x)$ and $\max(0, x - \alpha_2)$; and their smooth versions: V_1 (left) and V_2 (right), respectively, with $\alpha_1 = 3, \alpha_2 = 4$ and $\sigma = 1.5$

graphs G and G' . Once the parameters are estimated, we can derive a similarity matrix through the Laplacian kernel as in Eq. (7). More importantly, as \mathbf{W} is nonnegative, it is easy to see that $d_{\mathbf{W}}$ is a CND function, and thus the derived similarity matrix is valid.

3.2 Formulation of learning subtree pattern weights \mathbf{W}

We aim to learn the parameters W in Eq. (12) using the notions of metric learning (Kulis 2012). That is two input graphs with the same labels are encouraged to be closer while the two with different labels become far away from each other. In other words, within class distances should be small, while between class distances should be large. For this purpose, as a loss function for a graph pair g and g' , we can use the following two hinge loss function: $\max(0, \alpha_1 - d_{\mathbf{W}}(g, g'))$ if g and g' are with different labels and $\max(0, d_{\mathbf{W}}(g, g') - \alpha_2)$ otherwise, for learning subtree pattern weights, where α_1 and α_2 are constants ($\alpha_1 \geq \alpha_2$). The former yields a penalty if g and g' of different labels are closer than α_1 while the latter yields a penalty when g and g' of the same label are more distant than α_2 . Instead of using these functions in the optimization problem, we use their smooth versions: V_1 and V_2 (see Fig. 2), which offer useful properties for deriving a generalization bound for the problem in Sect. 4. The derivation of these functions is based on the connection between the strong convexity of a function and Lipschitz continuous gradient of its Fenchel dual (see more details in Nesterov (2005)).

More concretely, let $D_n = \{z_1 = (g_1, y_1), \dots, z_n = (g_n, y_n)\}$ where $g_i \in \mathcal{G}$ and $y_i \in \mathcal{Y} = \{-1, 1\}$, for $i = 1, \dots, n$, we formulate a constrained minimization problem for learning subtree pattern weights as follows:

$$\begin{aligned} & \underset{\mathbf{W}}{\text{minimize}} && \frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \ell(\mathbf{W}, z_i, z_j) \\ & \text{subject to} && \mathbf{W} \in \mathcal{C} \end{aligned} \tag{13}$$

where $\mathcal{C} = \{\mathbf{W} \in \mathbb{R}^d : \mathbf{W} = [\mathbf{w}_1^T, \dots, \mathbf{w}_H^T]^T, \|\mathbf{w}_h - \mathbf{c}_h\|_2 \leq \epsilon_h, 1 \leq h \leq H\}$ (\mathbf{c}_h and ϵ_h are constant vectors and scalars); ℓ is a continuously differentiable function and defined as follows: $\ell(\mathbf{W}, z_i, z_j) = \begin{cases} V_1(\mathbf{W}, g_i, g_j) & \text{if } y_i \neq y_j \\ V_2(\mathbf{W}, g_i, g_j) & \text{otherwise} \end{cases}$

$$\begin{aligned}
 &V_1(\mathbf{W}, g_i, g_j) \\
 &= \begin{cases} 0 & \text{if } d_{\mathbf{W}}(g_i, g_j) \geq \alpha_1 \\ \alpha_1 - \frac{\sigma}{2} - d_{\mathbf{W}}(g_i, g_j) & \text{if } d_{\mathbf{W}}(g_i, g_j) \leq \alpha_1 - \sigma \\ \frac{1}{2\sigma} (d_{\mathbf{W}}(g_i, g_j) - \alpha_1)^2 & \text{if } \alpha_1 - \sigma < d_{\mathbf{W}}(g_i, g_j) < \alpha_1 \end{cases} \quad (14)
 \end{aligned}$$

$$\begin{aligned}
 &V_2(\mathbf{W}, g_i, g_j) \\
 &= \begin{cases} 0 & \text{if } d_{\mathbf{W}}(g_i, g_j) \leq \alpha_2 \\ d_{\mathbf{W}}(g_i, g_j) - \alpha_2 - \frac{\sigma}{2} & \text{if } d_{\mathbf{W}}(g_i, g_j) \geq \alpha_2 + \sigma \\ \frac{1}{2\sigma} (d_{\mathbf{W}}(g_i, g_j) - \alpha_2)^2 & \text{if } \alpha_2 < d_{\mathbf{W}}(g_i, g_j) < \alpha_2 + \sigma \end{cases} \quad (15)
 \end{aligned}$$

and $\alpha_1, \alpha_2 (\alpha_1 \geq \alpha_2)$ and σ are constants; $d_{\mathbf{W}}(g_i, g_j)$ is calculated from $\mathbf{Z}(g_i, g_j)$ as in Eq. (12). To reduce notation, we use \mathbf{Z}_{ij} rather than $\mathbf{Z}(g_i, g_j)$ in the rest of the paper.

Lemma 3 Let $\beta = \max_{1 \leq i < j \leq n} \|\mathbf{Z}_{ij}\|_2, L = \frac{\beta^2}{\sigma}$, and $M = \max(b - \frac{\sigma}{2} - \alpha_2, \alpha_1 - \frac{\sigma}{2})$, the loss function ℓ defined in problem (13) is L -lipschitz, β -smooth and upper bounded by M .

Proof First, it is obvious to see that as $0 \leq d_{\mathbf{W}} \leq b$, we have the following bounds: $0 \leq V_1 \leq \alpha_1 - \frac{\sigma}{2}$ and $0 \leq V_2 \leq b - \frac{\sigma}{2} - \alpha_2$. Therefore, ℓ is upper bounded by M .

We derive the first and second order derivatives of the smooth function V_1 as follows:

$$\begin{aligned}
 &\nabla V_1(\mathbf{W}, g_i, g_j) \\
 &= \begin{cases} 0 & \text{if } d_{\mathbf{W}}(g_i, g_j) \geq \alpha_1 \\ \mathbf{Z}_{ij} & \text{if } d_{\mathbf{W}}(g_i, g_j) \leq \alpha_1 - \sigma \\ \frac{1}{\sigma} (\alpha_1 - d_{\mathbf{W}}(g_i, g_j)) \mathbf{Z}_{ij} & \text{if } \alpha_1 - \sigma < d_{\mathbf{W}}(g_i, g_j) < \alpha_1 \end{cases} \quad (16)
 \end{aligned}$$

$$\begin{aligned}
 &\nabla^2 V_1(\mathbf{W}, g_i, g_j) \\
 &= \begin{cases} 0 & \text{if } d_{\mathbf{W}}(g_i, g_j) \geq \alpha_1 \\ 0 & \text{if } d_{\mathbf{W}}(g_i, g_j) \leq \alpha_1 - \sigma \\ \frac{1}{\sigma} \mathbf{Z}_{ij} \mathbf{Z}_{ij}^T & \text{if } \alpha_1 - \sigma < d_{\mathbf{W}}(g_i, g_j) < \alpha_1 \end{cases} \quad (17)
 \end{aligned}$$

In order to prove the function V_1 is L -Lipschitz and β -smooth, it is sufficient to show that the norm of its derivative is always less than L : $\|\nabla V_1(\mathbf{W}, g_i, g_j)\|_2 \leq L$ and the spectral norm (or the maximum eigen value) of its second order derivative is always less than β : $\|\nabla^2 V_1(\mathbf{W}, g_i, g_j)\| \leq \beta, \forall g_i, g_j \in \mathcal{G}$. Indeed, from Eqs. (16) and (17), we have: $\|V_1(\mathbf{W}, g_i, g_j)\|_2 \leq \|\mathbf{Z}_{ij}\|_2$ and $\|\nabla^2 V_1(\mathbf{W}, g_i, g_j)\| \leq \frac{1}{\sigma} \|\mathbf{Z}_{ij} \mathbf{Z}_{ij}^T\| = \frac{1}{\sigma} \|\mathbf{Z}_{ij}\|_2^2$. Also, we can bound $\|\mathbf{Z}_{ij}\|_2$ by the inequality $\|\mathbf{Z}_{ij}\|_2 \leq \beta$. Thus V_1 is β -smooth and L -Lipschitz. Similarly, we can also show that V_2 is β -smooth and L -Lipschitz. The lemma is proven. \square

3.3 A stochastic learning algorithm for constrained optimization

The constrained optimization problem (13) is convex and thus guarantees to find its global optimum. Standard methods such as projected gradient descent can be used to solve the problem (13). However, for large scale data sets, solving the problem (13), involving n^2 terms with d parameters, might be computationally expensive. For instance, the data set PROTEIN (see Table 1) has $n^2 > 10^6$ pairs of examples and the weight vector size of $d > 10^5$ with the number of WL iterations $H = 5$. In this subsection, we present an efficient stochastic learning algorithm for dealing with this issue.

Algorithm 1: A stochastic algorithm for learning \mathbf{W}

Input: $D_n = \{z_1, \dots, z_n\}$, c , μ : learning rate and T :#maxIters

Output: solution \mathbf{W}^*

$\mathbf{W}^{(0)} = \mathbf{1}_d$, $t = 0$;

while $t \leq T$ **do**

 randomly pick two examples $z = (g, y)$ and $z' = (g', y')$ from D_n ;

if $y = y'$ **then**

 | $\text{grad}^{(t)} = \nabla V_1(\mathbf{W}^{(t)}, g, g')$;

else

 | $\text{grad}^{(t)} = \nabla V_2(\mathbf{W}^{(t)}, g, g')$;

end

$\mathbf{W}^{(t+1)} = \text{proj}_{\mathcal{C}}[\mathbf{W}^{(t)} - \mu \text{grad}^{(t)}]$;

end

$\mathbf{W}^* = \mathbf{W}^{(T)}$;

Let $\mathbf{W}^{(t)}$ denote the weight at iteration t . The weight is initialized by a vector of ones: $\mathbf{W}^{(0)} = \mathbf{1}_d$, which is also the case of WWL kernel without learning subtree pattern importance. At each iteration t , we randomly pick up a pair of examples ($z = (g, y), z' = (g', y')$) from the training data set D_n and compute the gradient $\text{grad}^{(t)}$ corresponding to this pair. In fact this step can be done efficiently due to the sparsity of the feature vector $\mathbf{Z}(g, g')$ in Eq. (12). Then, we update the current solution $\mathbf{W}^{(t)}$ to $\mathbf{W}^{(t+1)}$ by the following rule:

$$\mathbf{W}^{(t+1)} = \text{proj}_{\mathcal{C}}[\mathbf{W}^{(t)} - \mu \text{grad}^{(t)}]$$

where $\text{proj}_{\mathcal{C}}[\mathbf{W}]_h = \begin{cases} \mathbf{w}_h & \text{if } \|\mathbf{w}_h - \mathbf{c}_h\|_2 \leq \epsilon_h \\ \frac{\epsilon_h}{\|\mathbf{w}_h - \mathbf{c}_h\|_2}(\mathbf{w}_h - \mathbf{c}_h) + \mathbf{c}_h & \text{otherwise} \end{cases}$ maps a point \mathbf{w}_h ($1 \leq h \leq H$) back to the bounded feasible region. The procedure is illustrated in Algorithm 1.

4 Theoretical guarantees: a bound on generalization gap

In this section, we provide a bound on the generalization gap of the proposed stochastic learning algorithm for solving the problem (13). The gap is defined as the expected difference between the generalization error $R(\cdot)$ and empirical error $R_{D_n}(\cdot)$. In order to derive the generalization bound, we first provide basic setup and notations; then prove that our learning algorithm has a uniform stability, which is established in Theorem 5 using Lemma 3

and Lemma 4; finally derive our generalization bound, which is established in Theorem 9 using the McDiarmid inequality (see Theorem 6).

4.1 Basic setup and notations

Generalization Error. Let \mathbf{W}_n be the parameters of the parametric function (12) obtained by training on the data set D_n using Algorithm 1. Then the generalization error (or risk) $R(\mathbf{W}_n)$ with a loss function ℓ is defined as:

$$R(\mathbf{W}_n) = \mathbf{E}_{z,z'}[\ell(\mathbf{W}_n, z, z')]$$

where $\mathbf{E}_{z,z'}[\ell(\cdot, \cdot, \cdot)]$ denotes the expectation of function ℓ when z and z' are sampled according the distribution \mathcal{P} .

Empirical Error. The empirical error $R_{D_n}(\mathbf{W}_n)$ is defined on the training data set D_n as :

$$R_{D_n}(\mathbf{W}_n) = \frac{1}{n^2} \sum_{z_i \in D_n} \sum_{z_j \in D_n} \ell(\mathbf{W}_n, z_i, z_j)$$

Expected Generalization Gap. As Algorithm 1 is based on a randomized procedure, we use the definition of the expected generalization gap as follows:

$$\mathbb{K}_n = \mathbf{E}_{\text{SGD}}[R(\mathbf{W}_n) - R_{D_n}(\mathbf{W}_n)]$$

where \mathbf{E}_{SGD} denotes the expectation taken over the inherent randomness of the stochastic algorithm.

4.2 Uniform stability of the stochastic learning algorithm

Intuitively, a learning algorithm is said to have a uniform stability if its output is stable under a small modification of the training data set. For a randomized learning algorithm, the uniform stability property is defined as follows:

Definition 2 (Uniform Stability of the randomized algorithm) A randomized algorithm \mathbb{A} is β_n -uniformly stable with respect to a loss function ℓ , if the following inequality holds:

$$\forall(D_n, k), \sup_{z,z'} |\mathbf{E}_{\text{SGD}}[\ell(\mathbf{W}_n, z, z')] - \mathbf{E}_{\text{SGD}}[\ell(\mathbf{W}_{n,k}, z, z')]| \leq \beta_n$$

where $D_{n,k}$ is the new data set obtained from D_n by replacing $z_k \in D_n$ with a new example \hat{z}_k sampled from \mathcal{P} ; \mathbf{W}_n and $\mathbf{W}_{n,k}$ are the outputs of \mathbb{A} trained on two data sets D_n and $D_{n,k}$, respectively.

In order to prove that Algorithm 1 has the uniform stability property, we need the following lemma (its proof is placed in the "appendix" section):

Lemma 4 *Let the loss function ℓ defined in the problem (13) be β -smooth and L -Lipschitz; $\mathbf{W}_n^{(T)}$ and $\mathbf{W}_{n,k}^{(T)}$ be the parameters of the parametric form (12) trained on D_n and $D_{n,k}$, respectively, using Algorithm 1 with the number of iterations T and learning rate μ . Then, the expected difference in the model parameters is upper bounded by:*

$$\mathbf{E}_{\text{SGD}} \left[\left\| \mathbf{W}_n^{(T)} - \mathbf{W}_{n,k}^{(T)} \right\|_2 \right] \leq \frac{4}{n} \mu TL \quad (18)$$

Using Lemma 4 and L -Lipschitz property of function ℓ (see Lemma 3), we can now prove the stability of Algorithm 1.

Theorem 5 [Uniform Stability of Algorithm 1] *Let the loss function ℓ defined in the problem (13) be β -smooth and L -Lipschitz. Then Algorithm 1 with the fixed learning rate μ is k_n -uniformly stable where $k_n = \frac{4}{n} \mu TL^2$.*

Proof We have the following inequalities:

$$|\mathbf{E}_{\text{SGD}}[\ell(\mathbf{W}_n, z, z')] - \mathbf{E}_{\text{SGD}}[\ell(\mathbf{W}_{n,k}, z, z')]| \leq L \mathbf{E}_{\text{SGD}} \|\mathbf{W}_n - \mathbf{W}_{n,k}\|_2 \leq \frac{4}{n} \mu TL^2 \quad (19)$$

where the first and second inequalities are obtained by the L -Lipschitz property of ℓ and Lemma 4, respectively. This completes the proof. \square

4.3 Bound on generalization gap

Using the property of uniform stability in the previous subsection, we can derive the generalization bound which is done by the McDiarmid inequality (McDiarmid 1989).

Theorem 6 [McDiarmid inequality (McDiarmid 1989)] *Let X_1, X_2, \dots, X_n be n independent random variables taking values in \mathcal{X} and let $Z = f(X_1, \dots, X_n)$. If, for each $1 \leq i \leq n$, there exists a constant c_i such that*

$$\sup_{x_1, \dots, x_n, x'_i} |f(x_1, \dots, x_n) - f(x_1, \dots, x'_i, \dots, x_n)| \leq c_i, \forall 1 \leq i \leq n,$$

$$\text{then for any } \epsilon > 0, \Pr[|Z - \mathbf{E}[Z]| \geq \epsilon] \leq 2 \exp\left(\frac{-2\epsilon^2}{\sum_{i=1}^n c_i^2}\right)$$

To derive the bound on $R(W_n)$, we replace Z by K_n in Theorem 6 and bound $\mathbf{E}_{\text{SGD}}[K_n]$ and $|K_n - \mathbb{K}_{n,k}|$ which are established by the following lemmas (see their proofs in the "appendix" section).

Lemma 7 *For the loss function satisfying a uniform stability in k_n , we have the following inequality:*

$$\mathbf{E}_{D_n}[K_n] \leq 2k_n \quad (20)$$

Lemma 8 *For the loss function satisfying a uniform stability in k_n and upper bounded by M , we have the following inequality:*

$$|K_n - \mathbb{K}_{n,k}| \leq 2k_n + \frac{2M}{n} \quad (21)$$

Now we can derive the generalization bound for $R(W_n)$ in the following theorem:

Theorem 9 Let D_n be a training data set with n samples, \mathbf{W}_n be the solution obtained by minimizing the optimization problem (13) using Algorithm 1 with uniform stability k_n . Then the following inequality holds for probability of at least $1 - \delta$ ($0 \leq \delta \leq 1$):

$$\mathbf{E}_{\text{SGD}}[R(\mathbf{W}_n) - R_{D_n}(\mathbf{W}_n)] \leq 2k_n + (nk_n + M)\sqrt{\frac{2}{n}\log\frac{2}{\delta}} \tag{22}$$

Proof Applying McDiarmid’s concentration inequality (6) by replacing Z with \mathbb{K}_n , we have:

$$\Pr(\mathbb{K}_n - \mathbf{E}_{D_n}[\mathbb{K}_n] \geq \epsilon) \leq 2\exp\left(\frac{-2\epsilon^2}{n\left(2k_n + \frac{2M}{n}\right)^2}\right)$$

By fixing $\delta = 2\exp\left(\frac{-2\epsilon^2}{n\left(2k_n + \frac{2M}{n}\right)^2}\right)$, we get $\epsilon = (nk_n + M)\sqrt{\frac{2}{n}\log\frac{1}{\delta}}$ which completes the proof of Theorem 9. □

The generalization bound is meaningful if the bound converge to 0 as $n \rightarrow \infty$. Our derived generalization bound converges to 0 as k_n decays with $O(\frac{1}{n})$. This confirms Algorithm 1 converges.

5 Experiments

In this section, we demonstrate the benefit of learning subtree pattern weights by experiments on both synthetic and real-world data. We performed classification experiments using the C-SVM implementation LIBSVM (Chang and Lin 2011). The necessary parameters of SVM were selected by cross-validation on the training set. These are the regularization parameter $C \in \{10^{-3}, 10^{-2}, \dots, 10^2, 10^3\}$ and kernel parameter $\gamma \in \{0.0001, 0.001, 0.01\}$. For learning the weights \mathbf{W} of subtree patterns (or WL labels) in Algorithm 1, the learning rate μ and maximum number of iterations T were set as 0.0001 and 500, respectively; $\epsilon_h \in \{0.1, 0.5, 1.0\}$ were selected by cross validation based on the training set and \mathbf{c}_h was fixed as a vector of ones, i.e. $\mathbf{c}_h = \mathbf{1}_{|\Sigma^h|}$ for $h = 1, \dots, H$; the hyperparameters α_1, α_2 and σ were empirically determined as 1.0, 0.5 and 0.1, respectively. All kernels were implemented in Python 3.0 and experiments were conducted on an Intel Core i9 at 2.3 Ghz with 64GB of RAM using a single processor only. The source code can be accessed through <https://github.com/haidnguyen0909/weightedWWL>.

5.1 Synthetic data

We designed eight substructures, shown in Fig. 3, in which substructures indexed by 1, 2, 5 and 6 are assumed to be indicative to positive class (+ 1) as they have a pattern '1-0(-2)-0' in common. The others are indicative to negative class (- 1). Our synthetic data set consists of *eight groups of graphs, each corresponds to one of these eight substructures* by randomly adding noisy nodes and edges. We used groups corresponding substructures 1, 2, 3 and 4 as training data and the others as testing data. We constructed two kernels for graphs: WWL and the proposed method with number of WL iterations $H = 2$, then

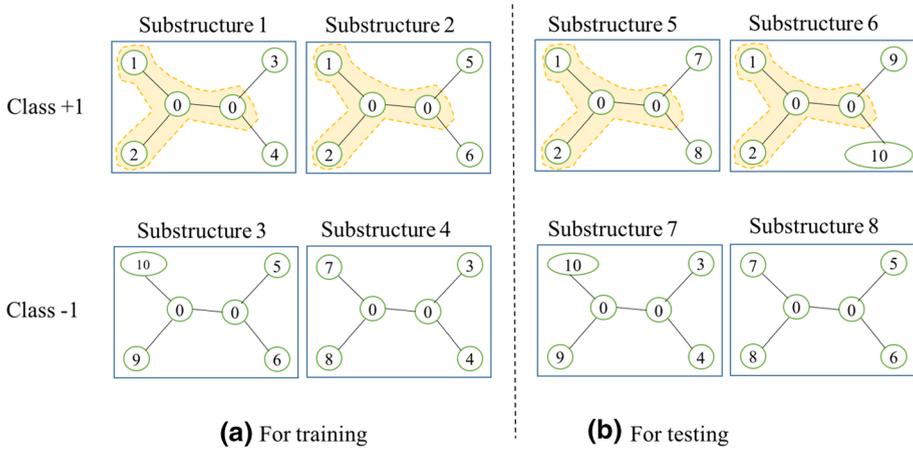


Fig. 3 Designed substructures 1–8: substructures 1, 2, 5 and 6 are indicative to positive class as they contain pattern ‘1–0(– 2)–0’ (emphasized in yellow). The rest are indicative to negative class. Graphs are generated from substructures by adding random nodes and noisy edges (20 examples for each substructure). Graphs generated from the substructures 1, 2, 3 and 4 are used for training. Graph generated from the substructures 5, 6, 7 and 8 are used for testing

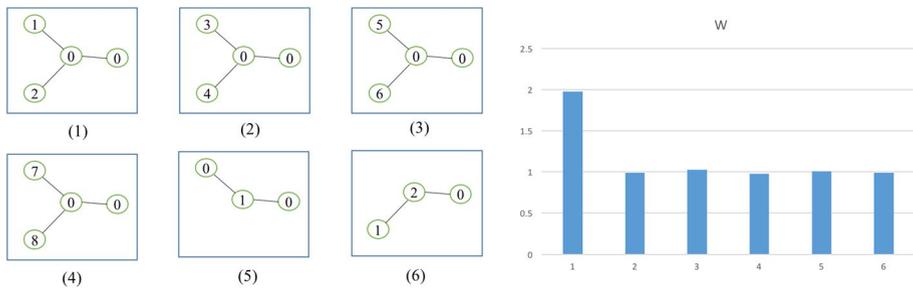


Fig. 4 Examples of selected subtree patterns of the first level $h = 1$ (left) and their weights learned by the proposed learning algorithm (right)

used SVM for classification. We reported mean accuracy obtained by ten synthetic data sets generated in this way.

We observed that WWL obtained mean accuracies of 82.4%, while the proposed method achieved significantly higher accuracy of 95%. It is noted that the testing examples were confusing the classifier. For instance, the substructure 5 has the same similarity with substructures 2 (indicative to positive class) and 4 (indicative to negative class) according to the WWL kernel, making it hard for the classifier to distinguish graphs containing the substructure 5. In contrast, this confusion can be alleviated by learning subtree pattern weights. In particular, the pattern ‘1–0(– 2)–0’ present in the substructure 5 is assigned a high weight by the proposed method (see Fig. 4). Therefore, graphs generated from the substructure 5 are more likely to be classified as positive.

Table 1 Statistics of datasets used in experiments

Datasets	#Graphs	#Classes	Avg. card (V)	Avg. card (E)	#labels
MUTAG	188	2 (125 vs. 63)	17.9	39.6	7
PTC-MR	344	2 (192 vs. 152)	25.6	51.9	19
PROTEIN	1113	2 (663 vs. 450)	39.1	145.63	3
NCI1	4110	2 (2053 vs. 2057)	N/A	N/A	N/A

5.2 Real-world data

In this subsection we present an experimental evaluation of the proposed method on real-world data. We report experimental results on four benchmark bioinformatics data sets, involving node-labeled graphs, particularly, MUTAG, PTC-MR, PROTEIN and NCI1. The MUTAG dataset consists of graph structures of 188 chemical compounds which are either mutagenic aromatic or heteroaromatic nitro compounds and nodes can have 7 discrete labels. The PTC-MR dataset consists of 344 chemical compounds which are known to cause or not cause cancer in rats and mice, and nodes can have 19 discrete labels. The PROTEIN dataset consists of relations between secondary structure elements represented by nodes and neighborhood in the amino-acid sequence or in 3D space by edges, and nodes can have 3 discrete labels. The NCI1 dataset is a balanced subset of chemical compounds screened for their ability to inhibit the growth of a panel of human tumor cell lines, and nodes can have 37 discrete labels. Some statistics of these data sets are shown in Table 1.

We compared the proposed method to several state-of-the-art graph kernels. Due to the large number of graph kernels in the literature, we selected representatives of the major families of graph kernels. In particular, for the family of walk based kernels, we compared the proposed method to the fast random walk kernel (Kashima 2003) that essentially counts the common labeled walks. For the family of path based graph kernels, we compared to the shortest path kernel (Borgwardt and Kriegel 2005). For the family of WL based graph kernels, we compared to WL subtree (Vishwanathan et al. 2010), WL-OA (Kriege et al. 2016) and WWL kernels (Togninalli et al. 2019). The WL based kernels have been shown to be superior to previous approaches.

We report mean predictive accuracies and standard deviations obtained by 10-fold cross-validation repeated 10 times with random fold assignments. Within each fold, the number of hops $H \in \{1, 2, \dots, 6\}$ was selected by cross validation based on the training set. The results evaluated by classification accuracy are summarised in Table 2. We used one-sided paired *t-test* to verify if the accuracy differences between two methods on data sets are statistically significant. We empirically observed that random walk and shortest path kernels were less competitive to WL-based kernels on four data sets. On three datasets MUTAG, PROTEIN and NCI1, the proposed method was comparable with WL-OA while it improved its unweighted version WWL by 1.4%, 1.5% and 0.8%, respectively (the calculated *p* values were 0.061, 0.0087 and 0.055, respectively, smaller than the significance level of $\alpha = 0.1$). On PTC-MR, the proposed method improved WWL by 0.6% while outperforming WL-OA by nearly 5% (the calculated *p* values were 0.0035 and < 0.001 , respectively). In all these data sets, random walk, shortest path and WL subtree kernels were dominated by the rest in large margins. Furthermore, we also investigated some selected subtree patterns at the first and second levels ($h = 1, 2$) along with their weights learned by the proposed algorithm from two data sets: MUTAG and PTC-MR (see Figs. 5

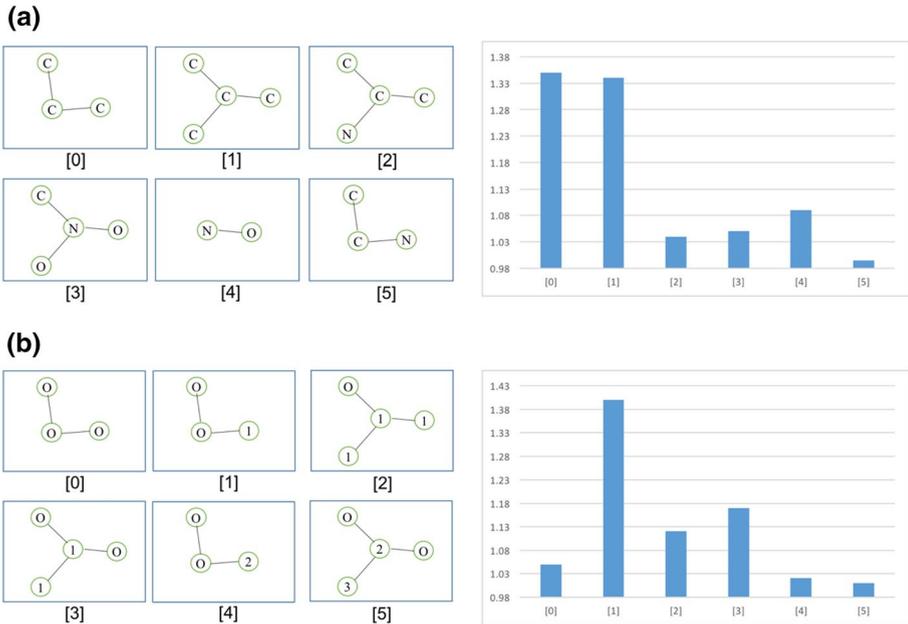


Fig. 5 Examples of selected subtree patterns and their weights of the first level $h = 1$ **a** and the second level $h = 2$ **b** extracted from MUTAG by Algorithm 1

and 6, respectively). Interestingly, the weights were found different over substructures, indicating their different degrees of importance in the prediction task. These experimental results confirmed the effectiveness of learning important subtree patterns for WWL kernels.

5.3 Computational efficiency of the proposed stochastic algorithm

In this subsection, we evaluate the computational efficiency of the proposed Algorithm 1. We empirically compared two variants: batch and stochastic (Algorithm 1), for solving the minimization problem (13) in terms of running time. The first variant considers all pairs of graphs for every step of projected gradient descent. The second variant considers one pair of graphs at a time to take a single step.

First we assessed the running time of two variants on randomly generated graphs (as described in Subsect. 5.1) with respect to two parameters: number of graphs N and number of WL iterations H . We varied N in range $\{50, 100, 200, 400, 600, 800, 1000\}$ and H in range $\{1, 2, 3, 4, 5, 6, 7\}$. For each individual experiment, we fixed one parameter at its default value and varied the other. The default values were 100 for N and 2 for H . We report CPU running times in seconds in Fig. 7. Empirically, we observed that the running time of full batch variant increased quickly when increasing the number of graphs N and the number of WL iterations H . In contrast, the stochastic variant scaled much better with much lower running times, indicating that Algorithm 1 has high scalability for large scale data sets. The computational efficiency of Algorithm 1 can be explained by the fact that

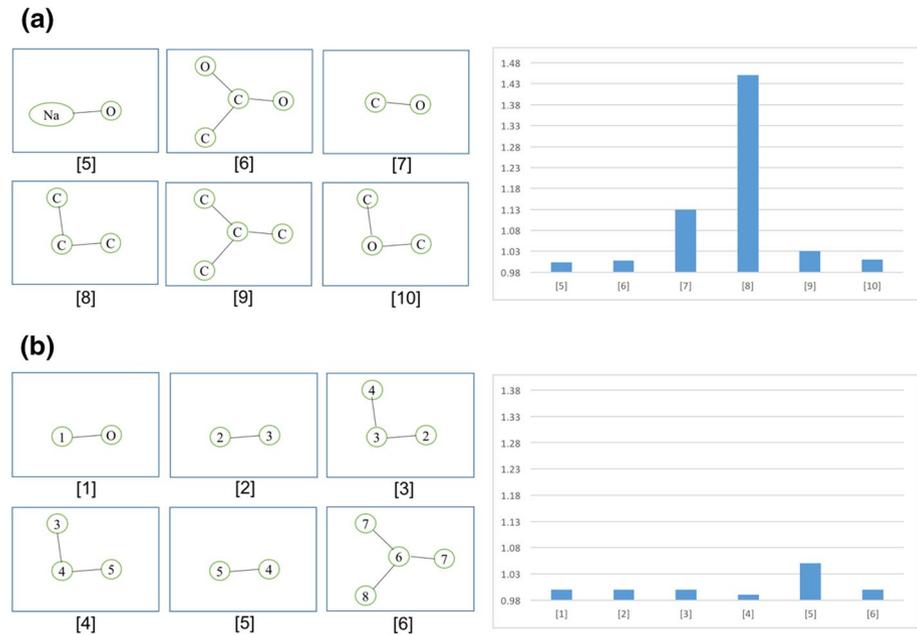


Fig. 6 Examples of selected subtree patterns and their weights of the first level $h = 1$ **a** and the second level $h = 2$ **b** extracted from PTC-MR by Algorithm 1

Table 2 Classification accuracies and standard deviation on real-world graph data sets: MUTAG, PTC-MR, PROTEIN and NCI1

Kernels	MUTAG	PTC-MR	PROTEIN	NCI1
Random Walk (Kashima 2003)	85.06 ± 0.18	55.74 ± 3.64	71.11 ± 0.83	62.88 ± 0.22
Shortest path (Borgwardt and Kriegel 2005)	85.49 ± 0.59	53.29 ± 0.92	73.03 ± 1.13	61.36 ± 0.19
WL Subtree (Shervashidze and Borgwardt 2009)	85.61 ± 0.85	61.89 ± 1.97	72.5 ± 0.32	85.61 ± 0.13
WL-OA (Kriege et al. 2016)	88.17 ± 1.98	60.49 ± 1.39	75.89 ± 0.41	86.17 ± 0.35
WWL (Togninalli et al. 2019)	86.95 ± 1.35	64.86 ± 1.57	74.25 ± 0.74	85.69 ± 0.28
Proposed	88.37 ± 1.82	65.44 ± 0.97	75.73 ± 0.57	86.45 ± 0.11

Best results are shown in bold

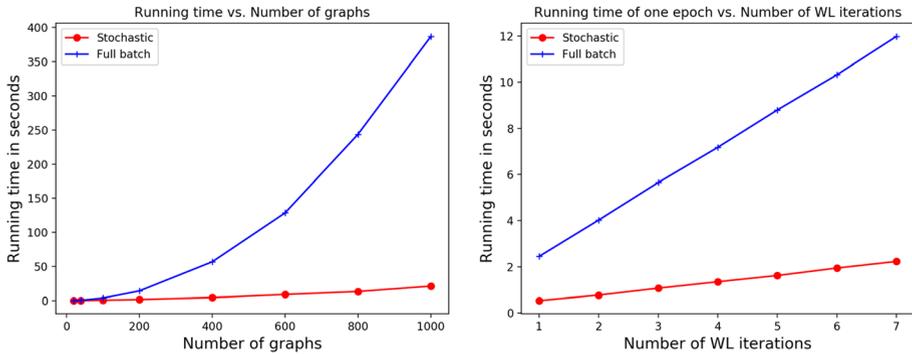


Fig. 7 Running time in seconds on synthetic data sets of two variants: full batch and stochastic algorithms, for learning the subtree pattern importance W (in the optimization problem (13)) (Default values: dataset size $N = 100$, WL iteration $H = 2$)

Table 3 Running time in seconds of two variants: full batch and stochastic on real-world data sets: MUTAG, PTC-MR, PROTEIN and NCI1

Variants/data sets	MUTAG	PTC-MR	PROTEIN	NCI1
Full batch	1' 32"	3' 48"	1 h 20'40"	5 h 5'
Stochastic	4h 29'	8h 50'	> 3 days	> 1 week

computing gradient of the loss function for a graph pair g and g' involves a few substructures (or WL labels) v shared by g and g' in Eq. (12), i.e., sparsity of the feature vector $\mathbf{Z}(g, g')$.

Second we assessed the running time of two variants on real-world data sets: MUTAG, PTC-MR, PROTEIN and NCI1. We reported the running time of two variants to finish the entire classification tasks, including learning subtree pattern weights, computing kernels and doing classification, in Table 3. The running time were taken average by 10-fold cross validation. We empirically observed that the full batch variant was slow when running on even small data sets MUTAG and PTC-MR, taking in approximately 4 h and 9 h, respectively. But the stochastic variant was much faster, taking only less than 4 min on the two data sets. We also observed that the stochastic variant could easily scale up to data sets with thousands of graphs. Particularly, on data sets PROTEIN and NCI1, the tasks were performed in nearly 1h 30' and 5 h, respectively. However, it was impossible for the full batch variant to finish the tasks in less than 3 days. These evidence showed that the proposed stochastic variant is highly scalable.

6 Conclusion and discussion

In this work, we proposed to learn the weights of substructures of graphs, particularly, *subtree patterns* (extracted by WL labeling scheme), to overcome the limitations of current graph kernels. We considered the problem of incorporating subtree pattern weights for WWL kernel (Togninalli et al. 2019) by formulating the parametric form of Wasserstein

distance taking into account subtree pattern weights, and learning these weights from data optimally for the tasks.

Our proposed method has several advantages. First, it can learn the importance of subtree patterns specifically for the tasks through their weights in the parametric distance function. Second, the kernels converted from the learned parametric function of subtree pattern weights are valid. Third, the efficient stochastic algorithm for learning the weights has high scalability for large scale data sets, and its theoretical guarantees are provided.

Although we considered WWL kernel for extracting important subtree patterns, an interesting and worthwhile extension of our work would be to apply this idea to other WL based graph kernels such as WL subtree and WL-OA kernels. The improvements of the optimization algorithm for learning subtree pattern weights in terms of both convergence and efficiency would also be our future work.

A Appendices

A.1 Proof of Lemma 4

Proof We prove the lemma by following the notion of using the same randomness for two data set D_n and $D_{n,k}$ as in Hardt et al. (2016). Particularly, we supply the sample sequences $S = \{p_1 = (z_{i_1}, z_{j_1}), \dots, p_T = (z_{i_T}, z_{j_T})\}$ to two identical learning algorithms except that for some t ($1 \leq t \leq T$), if p_t contains z_k ($p_t = (z_k, z_{j_t})$ or (z_{i_t}, z_k)), we replace it with $\hat{p}_t = (\hat{z}_k, z_{j_t})$ or (z_{i_t}, \hat{z}_k) . So, there are two cases to consider:

Case 1: At step t , Algorithm 1 picks a pair of samples (z, z') that contain no z_k ($z \neq z_k$ and $z' \neq z_k$) and this case occurs with probability $(1 - \frac{1}{n})^2$. Then, we have:

$$\begin{aligned} \left\| \mathbf{W}_n^{(t+1)} - \mathbf{W}_{n,k}^{(t+1)} \right\|_2^2 &= \left\| \text{proj}_C[\mathbf{W}_n^{(t)} - \mu \nabla \ell(\mathbf{W}_n^{(t)}, z, z')] - \text{proj}_C[\mathbf{W}_{n,k}^{(t)} - \mu \nabla \ell(\mathbf{W}_{n,k}^{(t)}, z, z')] \right\|_2^2 \\ &\leq \left\| \mathbf{W}_n^{(t)} - \mu \nabla \ell(\mathbf{W}_n^{(t)}, z, z') - \mathbf{W}_{n,k}^{(t)} + \mu \nabla \ell(\mathbf{W}_{n,k}^{(t)}, z, z') \right\|_2^2 \\ &= \left\| \mathbf{W}_n^{(t)} - \mathbf{W}_{n,k}^{(t)} \right\|_2^2 - 2\mu \left(\mathbf{W}_n^{(t)} - \mathbf{W}_{n,k}^{(t)} \right) \left(\nabla \ell(\mathbf{W}_n^{(t)}, z, z') - \nabla \ell(\mathbf{W}_{n,k}^{(t)}, z, z') \right) \\ &\quad + \mu^2 \left\| \nabla \ell(\mathbf{W}_n^{(t)}, z, z') - \nabla \ell(\mathbf{W}_{n,k}^{(t)}, z, z') \right\|_2^2 \\ &\leq \left\| \mathbf{W}_n^{(t)} - \mathbf{W}_{n,k}^{(t)} \right\|_2^2 - \mu \left(\frac{2}{\beta} - \mu \right) \left\| \nabla \ell(\mathbf{W}_n^{(t)}, z, z') - \nabla \ell(\mathbf{W}_{n,k}^{(t)}, z, z') \right\|_2^2 \end{aligned}$$

The second line is obtained by the fact that $\left\| \text{proj}_C[\mathbf{u}] - \text{proj}_C[\mathbf{v}] \right\|_2 \leq \left\| \mathbf{u} - \mathbf{v} \right\|_2$ for \mathbf{u}, \mathbf{v} in the domain. The last line is obtained by the β -smoothness of function ℓ (see Lemma 3). So we have the following inequality (by selecting $\mu \leq \frac{2}{\beta}$):

$$\left\| \mathbf{W}_n^{(t+1)} - \mathbf{W}_{n,k}^{(t+1)} \right\|_2 \leq \left\| \mathbf{W}_n^{(t)} - \mathbf{W}_{n,k}^{(t)} \right\|_2 \tag{23}$$

Case 2: At step t , Algorithm 1 picks a pair of samples (z, z') that contain z_k ($z = z_k$ and $z' = z_k$) and this case occurs with probability $1 - (1 - \frac{1}{n})^2$. Then we have:

$$\left\| \mathbf{W}_n^{(t+1)} - \mathbf{W}_{n,k}^{(t+1)} \right\|_2 \leq \left\| \mathbf{W}_n^{(t)} - \mathbf{W}_{n,k}^{(t)} \right\|_2 + 2\mu L \tag{24}$$

The above inequality holds as the norm of gradient of loss function ℓ is upper bounded by L . From two inequalities (23) and (24), and considering the probabilities of two cases, we have the following inequality:

$$\mathbf{E}_{\text{SGD}} \left[\left\| \mathbf{W}_n^{(t+1)} - \mathbf{W}_{n,k}^{(t+1)} \right\|_2 \right] \leq \left(1 - \frac{1}{n} \right)^2 \mathbf{E}_{\text{SGD}} \left[\left\| \mathbf{W}_n^{(t)} - \mathbf{W}_{n,k}^{(t)} \right\|_2 \right] \tag{25}$$

$$+ \left(1 - \left(1 - \frac{1}{n} \right)^2 \right) \left(\mathbf{E}_{\text{SGD}} \left[\left\| \mathbf{W}_n^{(t)} - \mathbf{W}_{n,k}^{(t)} \right\|_2 \right] + 2\mu L \right) \tag{26}$$

$$= \mathbf{E}_{\text{SGD}} \left[\left\| \mathbf{W}_n^{(t)} - \mathbf{W}_{n,k}^{(t)} \right\|_2 \right] + 2\mu L \left(1 - \left(1 - \frac{1}{n} \right)^2 \right) \tag{27}$$

By the above recursive formula, we obtain the following:

$$\mathbf{E}_{\text{SGD}} \left[\left\| \mathbf{W}_n^{(T)} - \mathbf{W}_{n,k}^{(T)} \right\|_2 \right] \leq 2 \left(1 - \left(1 - \frac{1}{n} \right)^2 \right) \mu TL \leq \frac{4}{n} \mu TL \tag{28}$$

which completes the proof. □

A.2 Proof of Lemma 7

Proof By the definition of \mathbb{K}_n as in (4.1), we have the following:

$$\begin{aligned} \mathbf{E}_{D_n} [\mathbb{K}_n] &= \mathbf{E}_{D_n} \mathbf{E}_{\text{SGD}} [R(\mathbf{W}_n) - R_{D_n}(\mathbf{W}_n)] \\ &= \mathbf{E}_{D_n} \mathbf{E}_{\text{SGD}} \mathbf{E}_{z_i, z'_i} \ell(\mathbf{W}_n, z_i, z'_i) - \mathbf{E}_{D_n} \mathbf{E}_{\text{SGD}} \frac{1}{n^2} \sum_{i,j=1}^n \ell(\mathbf{W}_n, z_i, z_j) \\ &= \mathbf{E}_{D_n, z_i, z'_i} \mathbf{E}_{\text{SGD}} \left[\frac{1}{n^2} \sum_{k=1}^n \sum_{j=1}^n [\ell(\mathbf{W}_n, z, z') - \ell(\mathbf{W}_n, z_k, z') + \ell(\mathbf{W}_n, z_k, z') - \ell(\mathbf{W}_n, z_k, z_j)] \right] \\ &= \mathbf{E}_{D_n, z_i, z'_i} \mathbf{E}_{\text{SGD}} \left[\frac{1}{n^2} \sum_{k=1}^n \sum_{j=1}^n [\ell(\mathbf{W}_n, z, z') - \ell(\mathbf{W}_n, z_k, z') + \ell(\mathbf{W}_n, z_k, z') - \ell(\mathbf{W}_n, z_k, z_j)] \right] \\ &= \underbrace{\mathbf{E}_{D_n, z_i, z'_i} \mathbf{E}_{\text{SGD}} \left[\frac{1}{n^2} \sum_{k=1}^n \sum_{j=1}^n [\ell(\mathbf{W}_n, z, z') - \ell(\mathbf{W}_n, z_k, z')] \right]}_{(a)} \\ &\quad + \underbrace{\mathbf{E}_{D_n, z_i, z'_i} \mathbf{E}_{\text{SGD}} \left[\frac{1}{n^2} \sum_{k=1}^n \sum_{j=1}^n [\ell(\mathbf{W}_n, z_k, z') - \ell(\mathbf{W}_n, z_k, z_j)] \right]}_{(b)} \end{aligned}$$

We first process the part (a) which is equivalent to the following:

$$\begin{aligned}
 \text{(a)} &= \frac{1}{n} \sum_{k=1}^n \mathbf{E}_{D_n, z, z'} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z, z')] - \frac{1}{n} \sum_{k=1}^n \mathbf{E}_{D_n, z, z'} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z_k, z')] \\
 &= \frac{1}{n} \sum_{k=1}^n \mathbf{E}_{D_n, z, z'} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z, z')] - \frac{1}{n} \sum_{k=1}^n \mathbf{E}_{D_n, \hat{z}_k, z'} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z_k, z')] \\
 &= \frac{1}{n} \sum_{k=1}^n \mathbf{E}_{D_n, z, z'} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z, z')] - \frac{1}{n} \sum_{k=1}^n \mathbf{E}_{D_n, z, z'} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_{n,k}, z, z')] \\
 &\leq \frac{1}{n} \sum_{k=1}^n \mathbf{E}_{D_n, z, z'} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z, z') - \ell(\mathbf{W}_{n,k}, z, z')] \leq k_n
 \end{aligned}$$

Similarly, we also prove that $(b) \leq k_n$ which completes the proof. □

A.3 Proof of Lemma 8

Proof By the definition of \mathbb{K}_n as in (4.1), we have:

$$\begin{aligned}
 |\mathbb{K}_n - \mathbb{K}_{n,k}| &= |\mathbf{E}_{\text{SGD}}[R(\mathbf{W}_n) - R_{D_n}(\mathbf{W}_n)] - \mathbf{E}_{\text{SGD}}[R(\mathbf{W}_{n,k}) - R_{D_{n,k}}(\mathbf{W}_{n,k})]| \\
 &\leq \underbrace{|\mathbf{E}_{\text{SGD}}R(\mathbf{W}_n) - \mathbf{E}_{\text{SGD}}R(\mathbf{W}_{n,k})|}_{(c)} + \underbrace{|\mathbf{E}_{\text{SGD}}R_{D_n}(\mathbf{W}_n) - \mathbf{E}_{\text{SGD}}R_{D_{n,k}}(\mathbf{W}_{n,k})|}_{(d)}
 \end{aligned}$$

We bound the two terms (c) and (d) as follows:

$$\begin{aligned}
 (c) &= |\mathbf{E}_{\text{SGD}} \mathbf{E}_{z, z'} [\ell(\mathbf{W}_n, z, z') - \ell(\mathbf{W}_{n,k}, z, z')]| \\
 &\leq \mathbf{E}_{z, z'} |\mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z, z') - \ell(\mathbf{W}_{n,k}, z, z')]| \leq k_n \\
 (d) &= |\mathbf{E}_{\text{SGD}} \frac{1}{n^2} \sum_{z_i \in D_n} \sum_{z_j \in D_n} \ell(\mathbf{W}_n, z_i, z_j) - \mathbf{E}_{\text{SGD}} \frac{1}{n^2} \sum_{z_i \in D_{n,k}} \sum_{z_j \in D_{n,k}} \ell(\mathbf{W}_{n,k}, z_i, z_j)| \\
 &= |\frac{1}{n^2} \sum_{i \neq k, j \neq k} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z_i, z_j) - \ell(\mathbf{W}_{n,k}, z_i, z_j)] \\
 &\quad + \frac{1}{n^2} \sum_{i \neq k} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z_i, z_k) - \ell(\mathbf{W}_{n,k}, z_i, z_k)] \\
 &\quad + \frac{1}{n^2} \sum_{j \neq k} \mathbf{E}_{\text{SGD}} [\ell(\mathbf{W}_n, z_k, z_j) - \ell(\mathbf{W}_{n,k}, z_k, z_j)]| \\
 &\leq \frac{(n-1)^2}{n^2} k_n + \frac{2(n-1)}{n^2} M k_n + \frac{2M}{n}
 \end{aligned}$$

The two above inequalities complete the proof. □

Acknowledgements D. H. N. has been supported in part by Otsuka Toshimi scholarship and JSPS Research Fellowship for Young Scientists (DC2) with KAKENHI [grant number 19J14714]. C. H. N. has been supported in part by MEXT KAKENHI [grant number 18K11434]. H. M. has been supported in part by JST ACCEL [grant number JPMJAC1503], MEXT KAKENHI [grant numbers 16H02868, 19H04169], FiDiPro by Tekes (currently Business Finland) and AIPSE program by Academy of Finland.

References

- Borgwardt, K.M., & Kriegel, H.P. (2005). Shortest-path kernels on graphs. In: Fifth IEEE International Conference on Data Mining (ICDM'05), pp. 8–pp. IEEE
- Chang, C. C., & Lin, C. J. (2011). Libsvm: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3), 1–27.
- Hardt, M., Recht, B., & Singer, Y. (2016). Train faster, generalize better: Stability of stochastic gradient descent. In: International Conference on Machine Learning, pp. 1225–1234. PMLR
- Haussler, D. (1999). Convolution kernels on discrete structures. Tech. rep., Technical report, Department of Computer Science, University of California.
- Kashima, H., Tsuda, K., & Inokuchi, A. (2003). Marginalized kernels between labeled graphs. In: Proceedings of the 20th International Conference on Machine Learning (ICML-03), pp. 321–328
- Kriege, N.M., Giscard, P.L., Wilson, R. (2016). On valid optimal assignment kernels and applications to graph classification. In: Advances in Neural Information Processing Systems, pp. 1623–1631
- Kulis, B., et al. (2012). Metric learning: A survey. *Foundations and Trends in Machine Learning*, 5(4), 287–364.
- McDiarmid, C. (1989). On the method of bounded differences. *Surveys in Combinatorics*, 141(1), 148–188.
- Munkres, J. (1957). Algorithms for the assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1), 32–38.
- Nesterov, Y. (2005). Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1), 127–152.
- Scott, J. (2011). Social network analysis: Developments, advances, and prospects. *Social Network Analysis and Mining*, 1(1), 21–26.
- Sharan, R., & Ideker, T. (2006). Modeling cellular machinery through biological network comparison. *Nature Biotechnology*, 24(4), 427–433.
- Shervashidze, N., & Borgwardt, K. (2009). Fast subtree kernels on graphs. In: Advances in neural information processing systems, pp. 1660–1668
- Togninalli, M., Ghisu, E., Llinares-López, F., & Rieck, B. (2019). Borgwardt, K.: Wasserstein weisfeiler-lehman graph kernels. In: Advances in Neural Information Processing Systems, pp. 6439–6449
- Trinajstić, N. (2018). *Chemical graph theory*. London: Routledge.
- Villani, C. (2008). *Optimal transport: Old and new* (Vol. 338). New York: Springer Science & Business Media.
- Vishwanathan, S. V. N., Schraudolph, N. N., Kondor, R., & Borgwardt, K. M. (2010). Graph kernels. *The Journal of Machine Learning Research*, 11, 1201–1242.

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