

A survey on graph kernels

Downloaded from: https://research.chalmers.se, 2024-09-18 20:14 UTC

Citation for the original published paper (version of record):

Kriege, N., Johansson, F., Morris, C. (2020). A survey on graph kernels. Applied Network Science, 5(1). http://dx.doi.org/10.1007/s41109-019-0195-3

N.B. When citing this work, cite the original published paper.

research.chalmers.se offers the possibility of retrieving research publications produced at Chalmers University of Technology. It covers all kind of research output: articles, dissertations, conference papers, reports etc. since 2004. research.chalmers.se is administrated and maintained by Chalmers Library

REVIEW

Open Access

A survey on graph kernels



Nils M. Kriege^{1*} ^(D), Fredrik D. Johansson² and Christopher Morris¹

*Correspondence: nils.kriege@tu-dortmund.de ¹Department of Computer Science, TU Dortmund University, Otto-Hahn-Str. 14, 44227 Dortmund, Germany Full list of author information is available at the end of the article

Abstract

Graph kernels have become an established and widely-used technique for solving classification tasks on graphs. This survey gives a comprehensive overview of techniques for kernel-based graph classification developed in the past 15 years. We describe and categorize graph kernels based on properties inherent to their design, such as the nature of their extracted graph features, their method of computation and their applicability to problems in practice. In an extensive experimental evaluation, we study the classification accuracy of a large suite of graph kernels on established benchmarks as well as new datasets. We compare the performance of popular kernels with several baseline methods and study the effect of applying a Gaussian RBF kernel to the metric induced by a graph kernel. In doing so, we find that simple baselines become competitive after this transformation on some datasets. Moreover, we study the extent to which existing graph kernels agree in their predictions (and prediction errors) and obtain a data-driven categorization of kernels as result. Finally, based on our experimental results, we derive a practitioner's guide to kernel-based graph classification.

Keywords: Supervised graph classification, Graph kernels, Machine learning

Introduction

Machine learning analysis of large, complex datasets has become an integral part of research in both the natural and social sciences. Largely, this development was driven by the empirical success of supervised learning of vector-valued data or image data. However, in many domains, such as chemo- and bioinformatics, social network analysis or computer vision, observations describe relations between objects or individuals and cannot be interpreted as vectors or fixed grids; instead, they are naturally represented by graphs. This poses a particular challenge in the application of traditional data mining and machine learning approaches. In order to learn successfully from such data, it is necessary for algorithms to exploit the rich information inherent to the graphs' structure and annotations associated with their vertices and edges.

A popular approach to learning with graph-structured data is to make use of graph kernels—functions which measure the similarity between graphs—plugged into a kernel machine, such as a support vector machine. Due to the prevalence of graph-structured data and the empirical success of kernel-based methods for classification, a large body of work in this area exists. In particular, in the past 15 years, numerous graph kernels have been proposed, motivated either by their theoretical properties or by their suitability and specialization to particular application domains. Despite this, there are no review articles aimed at comprehensive comparison between different graph kernels



© The Author(s). 2019 **Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (http://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

nor at giving practical guidelines for choosing between them. As the number of methods grow, it is becoming increasingly difficult for both non-expert practitioners and researchers new to the field to identify an appropriate set of candidate kernels for their application.

This survey is intended to give an overview of the graph kernel literature, targeted at the active researcher as well as the practitioner. First, we describe and categorize graph kernels according to their design paradigm, the used graph features and their method of computation. We discuss theoretical approaches to measure the expressivity of graph kernels and their applicability to problems in practice. Second, we perform an extensive experimental evaluation of state-of-the-art graph kernels on a wide range of benchmark datasets for *graph classification* stemming from chemo- and bioinformatics as well as social network analysis and computer vision. Finally, we provide guidelines for the practitioner for the successful application of graph kernels.

Contributions

We summarize our contributions below.

- We give a comprehensive overview of the graph kernel literature, categorizing kernels according to several properties. Primarily, we distinguish graph kernels by their mathematical definition and which graph features they use to measure similarity. Moreover, we discuss whether kernels are applicable to (i) graphs annotated with continuous attributes, or (ii) discrete labels, or (iii) unlabeled graphs only. Additionally, we describe which kernels rely on the *kernel trick* as opposed to being computed from feature vectors and what effects this has on the running time and flexibility.
- We give an overview of applications of graph kernels in different domains and review theoretical work on the expressive power of graph kernels.
- We compare state-of-the-art graph kernels in an extensive experimental study across a wide range of established and new benchmark datasets. Specifically, we show the strengths and weaknesses of the individual kernels or classes of kernels for specific datasets.
 - We compare popular kernels to simple baseline methods in order to assess the need for more sophisticated methods which are able to take more structural features into account. To this end, we analyze the ability of graph kernels to distinguish the graphs in common benchmark datasets.
 - Moreover, we investigate the effect of combining a Gaussian RBF kernel with the metric induced by a graph kernel in order to learn non-linear decision boundaries in the feature space of the graph kernel. We observe that with this approach simple baseline methods become competitive to state-of-the-art kernels for some datasets, but fail for others.
 - We study the similarity between graph kernels in terms of their classification predictions and errors on graphs from the chosen datasets. This analysis provides a qualitative, data-driven means of assessing the similarity of different kernels in terms of which graphs they deem similar.
- Finally, we provide guidelines for the practitioner and new researcher for the successful application of graph kernels.

Related work

The most recent surveys of graph kernels are the works of Ghosh et al. (2018) and Zhang et al. (2018a). Ghosh et al. (2018) place a strong emphasis on covering the fundamentals of kernel methods in general and summarizing known experimental results for graph kernels. The article does not, however, cover the most recent contributions to the literature. Most importantly, the article does not provide a detailed experimental study comparing the discussed kernels. That is, the authors do not perform (nor reproduce) original experiments on graph classification and solely report numbers found in the corresponding original paper. The survey by Zhang et al. (2018a) focuses on kernels for graphs without attributes which is a small subset of the scope of this survey. Moreover, it does not discuss the most recent developments in this area. Another survey was published in 2010 by Vishwanathan et al. (2010) but its main topic are random walk kernels and it does not include recent advances. Moreover, various PhD theses give (incomplete or dated) overviews, see, e.g., (Borgwardt 2007; Kriege 2015; Neumann 2015; Shervashidze 2012). None of the papers provides compact guidelines for choosing a kernel for a particular dataset.

Compared to the existing surveys, we provide a more complete overview covering a larger number of kernels, categorizing them according to their design, the extracted graph features and their computational properties. The validity of comparing results from different papers depends on whether these were obtained using comparable experimental setups (e.g., choices for hyperparameters, number of folds used for cross-validation, etc.), which is not the case across the entire spectrum of the graph kernel literature. Hence, we conducted an extensive experimental evaluation comparing a large number of graph kernels and datasets going beyond comparing kernels just by their classification accuracy. Another unique contribution of this article is a practitioner's guide for choosing between graph kernels.

Outline

In the "Fundamentals" section, we introduce notation and provide mathematical definitions necessary to understand the rest of the paper. The "Graph kernels" section gives an overview of the graph kernel literature. We start off by introducing kernels based on neighborhood aggregation techniques. Subsequently, we describe kernels based on assignments, substructures, walks and paths, and neural networks, as well as approaches that do not fit into any of the former categories. In the "Expressivity of graph kernels" section, we survey theoretical work on the expressivity of kernels and in the "Applications of graph kernels" section we describe applications of graph kernels in four domain areas. Finally, in the "Experimental study" section we introduce and analyze the results of a large-scale experimental study of graph kernels in classification problems, and provide guidelines for the successful application of graph kernels.

Fundamentals

In this section, we cover notation and definitions of fundamental concepts pertaining to graph-structured data, kernel methods, and graph kernels. In the "Graph kernels" section, we use these concepts to define and categorize popular graph kernels.

Graph data

A graph G is a pair (V, E) of a finite set of vertices V and a set of edges $E \subseteq \{\{u, v\} \subseteq V \mid u \neq v\}$. A vertex is typically used to represent an object (e.g., an atom) and an edge

a relation between objects (e.g., a molecular bond). We denote the set of vertices and the set of edges of *G* by V(G) and E(G), respectively. We restrict our attention to *undirected* graphs in which no two edges with identical (unordered) end points, nor any self-cycles exist. For ease of notation we denote the edge $\{u, v\}$ in E(G) by (u, v) or (v, u). A *labeled* graph is a graph *G* endowed with a *label function* $l: V(G) \rightarrow \Sigma$, where Σ is some alphabet, e.g., the set of natural or real numbers. We say that l(v) is the *label* of v. In the case $\Sigma = \mathbb{R}^d$ for some d > 0, l(v) is the *(continuous)* attribute of v. In the "Applications of graph kernels" section, we give examples of applications involving graphs with vertex labels and attributes. The edges of a graph may also be assigned labels or attributes (e.g., weights representing vertex similarity), in which case the domain of the labeling function l may be extended to the edge set.

We let N(v) denote the *neighborhood* of a vertex $v \in G$ in V(G), i.e., $N(v) = \{u \in V(G) \mid (v, u) \in E(G)\}$. The degree of a vertex is the size of its neighborhood, $\deg(u) = |N(u)|$. A *walk* ω in a graph is an ordered sequence of vertices $\omega = (u, ..., v)$ such that any two subsequent vertices are connected by an edge. A (u, v)-*path* is a walk that starts in u and ends in v with no repeated vertices. A graph G is called *connected* if there is a path between any pair of vertices in V(G) and *disconnected* otherwise. Paths, vertices, edges and neighborhoods are illustrated in Fig. 1.

We say that two unlabeled graphs *G* and *H* are *isomorphic*, denoted by $G \simeq H$, if there exists a bijection $\varphi : V(G) \rightarrow V(H)$, such that $(u, v) \in E(G)$ if and only if $(\varphi(u), \varphi(v)) \in E(H)$ for all u, v in V(G). For labeled graphs, isomorphism holds only if the bijection maps only vertices and edges with the same label. Finally, a graph G' = (V', E') is a *subgraph* of a graph G = (V, E) if $V' \subseteq V$ and $E' \subseteq E$. Let $S \subseteq V(G)$ be a subset of vertices in *G*. Then $G[S] = (S, E_S)$ denotes the *subgraph induced* by *S* with $E_S = \{(u, v) \in E(G) \mid u, v \in S\}$.

Graphs are often represented in matrix form. Perhaps most frequent is the *adjacency matrix* A with binary elements $a_{uv} = \{1 \text{ iff}(u, v) \in E\}^1$. An alternative representation is the *graph Laplacian* L, defined as L = D - A, where D is the diagonal degree matrix, such that $d_{uu} = \deg(u)$. Finally, the *incidence matrix* M of a graph is the binary $n \times n^2$ matrix with vertex-edge-pair elements $m_{ue} = \{1 \text{ iff } e = (u, v) \in E\}$ representing the event that the vertex u is incident on the edge e. It holds that $L = MM^{\top}$. The matrices A, L, and M all carry the same information.

Kernel methods

Kernel methods refer to machine learning algorithms that learn by comparing pairs of data points using particular similarity measures—*kernels*. We give an overview below; for an in-depth treatment, see (Schölkopf and Smola 2001; Shawe-Taylor and Cristianini 2004). Consider a non-empty set of data points χ , such as \mathbb{R}^d or a finite set of graphs, and let $k: \chi \times \chi \to \mathbb{R}$ be a function. Then, k is a kernel on χ if there is a Hilbert space \mathcal{H}_k and a feature map $\phi: \chi \to \mathcal{H}_k$ such that $k(x, y) = \langle \phi(x), \phi(y) \rangle$ for $x, y \in \chi$, where $\langle \cdot, \cdot \rangle$ denotes the inner product of \mathcal{H}_k . Such a feature map exists if and only if k is a *positive-semidefinite* function. A trivial example is where $\chi = \mathbb{R}^d$ and $\phi(x) = x$, in which case the kernel equals the dot product, $k(x, y) = x^\top y$.

An important concept in kernel methods is the *Gram matrix* K, defined with respect to a finite set of data points $x_1, ..., x_m \in \chi$. The Gram matrix of a kernel k has elements K_{ij} ,

¹Weighted graphs are represented by their corresponding edge weight matrix.



for $i, j \in \{0, ..., m\}$ equal to the kernel value between pairs of data points, i.e., $K_{ij} = k(x_i, x_j)$. If the Gram matrix of k is positive semidefinite for every possible set of data points, k is a kernel (Schölkopf et al. 1997). Kernel methods have the desirable property that they do not rely on explicitly characterizing the vector representation $\phi(x)$ of data points, but access data only via the Gram matrix K. The benefit of this is often illustrated using the Gaussian radial basis function (RBF) kernel on $\mathbb{R}^d, d \in \mathbb{N}$, defined as

$$k_{\text{RBF}}(\boldsymbol{x}, \boldsymbol{y}) = \exp\left(-\frac{\|\boldsymbol{x} - \boldsymbol{y}\|^2}{2\sigma^2}\right),\tag{1}$$

where σ is a bandwidth parameter. The Hilbert-space associated with the Gaussian RBF kernel has infinite dimension but the kernel may be readily computed for any pair of points (x, y) (see (Mohri et al. 2012) for further details). Kernel methods have been developed for most machine learning paradigms, e.g., support vector machines (SVM) for classification (Cortes and Vapnik 1995), Gaussian processes (GP) for regression (Rasmussen 2004), kernel PCA, k-means for unsupervised learning and clustering (Schölkopf et al. 1997), and kernel density estimation (KDE) for density estimation (Silverman 1986). In this work, we restrict our attention to *classification* of objects in a non-empty set of graphs \mathbb{G} . In this setting, a kernel $k: \mathbb{G} \times \mathbb{G} \to \mathbb{R}$ is called a *graph kernel*. Like kernels on vector spaces, graph kernels can be calculated either *explicitly* (by computing ϕ) or *implicitly* (by computing only k). Traditionally, learning with implicit kernel representations means that the value of the chosen kernel applied to every pair of graphs in the training set must be computed and stored. Explicit computation means that we compute a finite dimensional feature vector for each graph; the values of the kernel can then be computed on-the-fly during learning as the inner product of feature vectors. If explicit computation is possible, and the dimensionality of the resulting feature vectors is not too high, or the vectors are sparse, then it is usually faster and more memory efficient than implicit computation, see also (Kriege et al. 2014; Kriege et al. 2019).

Design paradigms for kernels on structured data

When working with vector-valued data, it is common practice for kernels to compare objects $x, y \in \mathbb{R}^d$ using differences between vector components (see for example the

Gaussian RBF kernel in the "Kernel methods" section). The structure of a graph, however, is invariant to permutations of its representation—the ordering by which vertices and edges are enumerated does not change the structure—and vector distances between, e.g., adjacency matrices, are typically uninformative. For this reason, it is important to compare graphs in ways that are themselves permutation invariant. As mentioned previously, two graphs with identical structure (irrespective of representation) are called *isomorphic*, a concept that could in principle be used for learning. However, not only is there no known polynomial-time algorithm for testing graph isomorphism (Johnson 2005) but isomorphism is also typically too strict for learning—it is akin to learning with the equality operator. In practice, it is often desirable to have *smoother* metrics of comparison in order to gain generalizable knowledge from the comparison of graphs.

The vast majority of graph kernels proposed in the literature are instances of so-called *convolution kernels*. Given two discrete structures, e.g., two graphs, the idea of Haussler's Convolution Framework (Haussler 1999) is to decompose these two structures into substructures, e.g., vertices or subgraphs, and then evaluate a kernel between each pair of such substructures. The convolution kernel is defined below.

Definition 1 (Convolution Kernel) Let $\mathcal{R} = \mathcal{R}_1 \times \cdots \times \mathcal{R}_d$ denote a space of components such that a composite object $X \in \mathcal{X}$ decomposes into elements of \mathcal{R} . Further, let $\mathcal{R} : \mathcal{R} \to \mathcal{X}$ denote the mapping from components to objects, such that $\mathcal{R}(x) = X$ if and only if the components $x \in \mathcal{R}$ make up the object $X \in \mathcal{X}$, and let $\mathcal{R}^{-1}(X) = \{x \in \mathcal{R} : \mathcal{R}(x) = X\}$. Then, the \mathcal{R} -convolution kernel is

$$k_{CV}(X,Y) = \sum_{x \in R^{-1}(X)} \sum_{y \in R^{-1}(Y)} \underbrace{\prod_{i=1}^{d} k_i(x_i, y_i)}_{k(x,y)},$$
(2)

where k_i is a kernel on \mathcal{R}_i for i in $\{1, \ldots, d\}$.

In our context, we may view the inverse map $R^{-1}(G)$ of the convolution kernel as the set of all components of a graph G that we wish to compare. A simple example of the R-convolution kernel is the *vertex label kernel* for which the mapping R takes the attributes $x_u \in \mathcal{R}$ of each vertex $u \in G \cup H$ and maps them to the graph that u is a member of. We expand on this notion in the "Subgraph patterns" section. A benefit of the convolution kernel framework when working with graphs is that if the kernels on substructures are invariant to orderings of vertices and edges, so is the resulting graph kernel.

A property of convolution kernels often regarded as unfavorable is that the sum in Eq. (2) applies to all pairs of components. When the considered components become more and more specific, each object becomes increasingly similar to itself, but no longer to any other objects. This phenomenon is referred to as the *diagonal dominance problem*, since the entries on the main diagonal of the Gram matrix are much higher than the others entries. This problem was observed for graph kernels, for which weights between the components were introduced to alleviate the problem (Yanardag and Vishwanathan 2015a; Aiolli et al. 2015). In addition, the fact that convolution kernels compare all pairs of components may be unsuitable in situations where each component of one object corresponds to exactly one component of the other (such as the features of two faces). Shin and Kuboyama (2008) studied *mapping kernels*, where the sum moves over a predetermined subset of pairs rather than the entire cross product. It was shown that, for general primitive kernels k, a valid mapping kernel is obtained if and only if the considered subsets of pairs are transitive on \mathcal{R} . This does not necessarily hold, when assigning the components of two objects to each other such that a correspondence of maximum total similarity w.r.t. k is obtained. As a consequence, this approach does not lead to valid kernels in general. However, graph kernels following this approach have been studied in detail and are often referred to as *optimal assignment kernels*, see in the "Assignment- and matching-based approaches" section.

Graph kernels

The first methods for graph comparison referred to as *graph kernels* were proposed in 2003 (Gärtner et al. 2003; Kashima et al. 2003). However, several approaches similar to graph kernels had been developed in the field of chemoinformatics, long before the term graph kernel was coined. The timeline in Fig. 2 shows milestones in the development of graph kernels and related learning algorithms for graphs. We postpone the discussion of the latter to "Chemoinformatics" section. Following the introduction of graph kernels, subsequent work focused for a long time on making kernels computationally tractable for large graphs with (predominantly) discrete vertex labels. Since 2012, several kernels specifically designed for graphs with continuous attributes have been proposed. It remains a current challenge in research to develop neural techniques for graphs that are able to learn feature representations that are clearly superior to the fixed feature spaces used by graph kernels.

In the following, we give an overview of the graph kernel literature in order of popular design paradigms. We begin our treatment with kernels that are based on neighborhood aggregation techniques. The subsequent subsections deal with *assignment-* and *matching-based* kernels, and kernels based on the extraction of subgraph patterns, respectively. The final subsections deal with kernels based on *walks and paths*, and kernels that do not fall into either of the previous categories. Table 1 gives an overview of the discussed graph kernels and their properties.

Neighborhood aggregation approaches

One of the dominating paradigms in the design of graph kernels is representation and comparison of *local* structure. Two vertices are considered similar if they have identical labels—even more so if their neighborhoods are labeled similarly. Expanding on this notion, two graphs are considered similar if they are composed of vertices with similar neighborhoods, i.e., that they have similar local structure. The different ways by which local structure is defined, represented and compared form the basis for several influential graph kernels. We describe a first example next.

Neighborhood aggregation approaches work by assigning an attribute to each vertex based on a summary of the local structure around them. Iteratively, for each vertex, the attributes of its immediate neighbors are aggregated to compute a new attribute for the target vertex, eventually representing the structure of its extended neighborhood. Shervashidze et al. (2011) introduced a highly influential class of neighborhood aggregation kernels for graphs with discrete labels based on the 1-dimensional *Weisfeiler-Lehman* (1-WL) or *color refinement* algorithm—a well-known heuristic for the graph isomorphism



cheminformatics are shown in *italics*, kernels for attributed graphs in **bold**

problem, see, e.g., (Babai and Kucera 1979). We illustrate an application of the 1-WL algorithm in Fig. 3.

Let *G* and *H* be graphs, and let $l: V(G) \cup V(H) \rightarrow \Sigma$ be the observed vertex label function of *G* and *H*.² In a series of iterations i = 0, 1, ..., the 1-WL algorithm computes new label functions $l^i: V(G) \cup V(H) \rightarrow \Sigma$, each of which can be used to compare *G* and *H*. In iteration 0 we set $l^0 = l$ and in subsequent iterations i > 0, we set

²If the graph is unlabeled, let l map to a constant.

Table 1 Summary of selected graph kernels: Computation by explicit (EX) and implicit (IM) feature mapping and support for attributed graphs

Graph kernel	Computation	Labels	Attributes
Shortest-Path (Borgwardt and Kriegel 2005)	IM	$+^{\dagger}$	$+^{\dagger}$
Generalized Shortest-Path (Hermansson et al. 2015)	IM	+	$+^{\dagger}$
Graphlet (Shervashidze et al. 2009)	EX	-	-
Cycles and Trees (Horváth et al. 2004)	EX	$+^{\star}$	-
Tree Pattern Kernel (Ramon and Gärtner 2003; Mahé and Vert 2009)	IM	+	$+^{\star}$
Ordered Directed Acyclic Graphs (Da San Martino et al. 2012a; 2012b)	EX	+	-
GraphHopper (Feragen et al. 2013)	IM	$+^{\dagger}$	+
Graph Invariant (Orsini et al. 2015)	IM	+	+
Subgraph Matching (Kriege and Mutzel 2012)	IM	+	+
Weisfeiler-Lehman Subtree (Shervashidze et al. 2011)	EX	+	-
Weisfeiler-Lehman Edge (Shervashidze et al. 2011)	EX	+	-
Weisfeiler-Lehman Shortest-Path (Shervashidze et al. 2011)	EX	+	-
k-dim. Local Weisfeiler-Lehman Subtree (Morris et al. 2017)	EX	+	_
Neighborhood Hash Kernel (Hido and Kashima 2009)	EX	+	_
Propagation Kernel (Neumann et al. 2016)	EX	+	+
Neighborhood Subgraph Pairwise Distance Kernel (Costa and De Grave 2010)	EX	+	-
Random Walk (Gärtner et al. 2003; Kashima et al. 2003; Mahé et al. 2004; Vishwanathan et al. 2010; Sugiyama and Borgwardt 2015; Kang et al. 2012)	IM	+	+
Optimal Assignment Kernel (Fröhlich et al. 2005)	IM	+	+
Weisfeiler-Lehman Optimal Assignment (Kriege et al. 2016)	IM	+	-
Pyramid Match (Nikolentzos et al. 2017b)	IM	+	-
Matchings of Geometric Embeddings (Johansson and Dubhashi 2015)	IM	+	$+^{\star}$
Descriptor Matching Kernel (Su et al. 2016)	IM	+	$+^{\dagger}$
Graphlet Spectrum (Kondor et al. 2009)	EX	+	-
Multiscale Laplacian Graph Kernel (Kondor and Pan 2016)	IM	+	$+^{\star\dagger}$
Global Graph Kernel (Johansson et al. 2014)	EX	-	-
Deep Graph Kernels (Yanardag and Vish- wanathan 2015a)	IM	+	-
Smoothed Graph Kernels (Yanardag and Vishwanathan 2015b)	IM	$+^{\star}$	-
Hash Graph Kernel (Morris et al. 2016)	EX	+	+
Depth-based Representation Kernel (Bai et al. 2014)	IM	_	_
Aligned Subtree Kernel (Bai et al. 2015)	IM	+	-

The column 'Labels' refers to whether the kernels support comparison of graphs with discrete vertex and edge labels in a way that depends on the interplay between structure and labels. The column 'Attributes' refer to the same capability but for continuous or more general vertex attributes. \star — not considered in publication, but method can be extended; \dagger — vertex annotations only

$$l^{i}(v) = \text{relabel}\left(\left(l^{i-1}(v), \text{sort}\left(\left\{\!\!\left\{l^{i-1}(u) \mid u \in N(v)\,\!\right\}\!\!\right\}\right)\!\!\right),\tag{3}$$

for $\nu \in V(G) \cup V(H)$, where sort(*S*) returns a sorted tuple of the multiset *S* and the injection relabel(*p*) maps the pair *p* to a unique value in Σ which has not been used in previous iterations. Now if *G* and *H* have an unequal number of vertices with label $\sigma \in \Sigma$, we can conclude that the graphs are not isomorphic. Moreover, if the cardinality of the image of l^{i-1} equals the cardinality of the image of l^i , the algorithm terminates.

The idea of the *Weisfeiler-Lehman subtree kernel* is to compute the above algorithm for $h \ge 0$ iterations, and after each iteration *i* compute a feature vector $\phi^i(G) \in \mathbb{R}^{|\Sigma_i|}$ for each graph *G*, where $\Sigma_i \subseteq \Sigma$ denotes the image of l^i . Each component $\phi^i(G)_{\sigma_j^i}$ counts the number of occurrences of vertices labeled with $\sigma_j^i \in \Sigma_i$. The overall feature vector $\phi_{WL}(G)$ is defined as the concatenation of the feature vectors of all *h* iterations, i.e.,

$$\left(\phi^0(G)_{\sigma_1^0},\ldots,\phi^0(G)_{\sigma_{|\Sigma_0|}^0},\ldots,\phi^h(G)_{\sigma_1^h},\ldots,\phi^h(G)_{\sigma_{|\Sigma_h|}^h}\right).$$

Then the Weisfeiler-Lehman subtree kernel for *h* iterations is $k_{WL}(G, H) = \langle \phi_{WL}(G), \phi_{WL}(H) \rangle$. The running time for a single feature vector computation is in $\mathcal{O}(hm)$ and $\mathcal{O}(Nhm + N^2hn)$ for the computation of the Gram matrix for a set of *N* graphs (Shervashidze et al. 2011), where *n* and *m* denote the maximum number of vertices and edges over all *N* graphs, respectively.

The WL subtree kernel suggests a general paradigm for comparing graphs at different levels of resolution: iteratively relabel graphs using the WL algorithm and construct a graph kernel based on a base kernel applied at each level. Indeed, in addition to the subtree kernel, Shervashidze et al. (2011) introduced two other variants, the *Weisfeiler-Lehman edge* and the *Weisfeiler-Lehman shortest-path kernel*. Instead of counting the labels of vertices after each iteration the Weisfeiler-Lehman edge kernel counts the colors of the two endpoints for all edges. The Weisfeiler-Lehman shortest-path kernel is the sum of shortest-path kernels applied to the graphs with refined labels l^i for $i \in \{0, ..., h\}$.

Morris et al. (2017) introduced a graph kernel based on higher dimensional variants of the Weisfeiler-Lehman algorithm. Here, instead of iteratively labeling vertices, the algorithm labels k-tuples or sets of cardinality k. Morris et al. (2017) also provide efficient approximation algorithm to scale the algorithm up to large datasets. In (Hido and Kashima 2009), a graph kernel similar to the 1-WL was introduced which replaces the neighborhood aggregation function Eq. (3) by a function based on binary arithmetic. Similarly, in Neumann et al. (2016) the propagation kernel is defined which propagates labels, and real-valued attributes for several iterations while tracking their distribution for every vertex. A randomized approach based on *p*-stable locality-sensitive hashing is used to obtain unique features after each iteration. In recent years, graph neural networks (GNNs) have emerged as an alternative to graph kernels. Standard GNNs can be viewed as a feedforward neural network version of the 1-WL algorithm, where colors (labels) are replaced by continuous feature vectors and network layers are used to aggregate over vertex neighborhoods (Hamilton et al. 2017; Kipf and Welling 2017). Recently, a connection between the 1-WL and GNNs has been established (Morris et al. 2019), showing that any possible GNN architecture cannot be more powerful than the 1-WL in terms of distinguishing non-isomorphic graphs.

Bai et al. (2014; 2015) proposed graph kernels based on *depth-based representations*, which can be seen as a different form of neighborhood aggregation. For a vertex v the



m-layer expansion subgraph is the subgraph induced by the vertices of shortest-path distance at most *m* from the vertex *v*. In order to obtain a vertex embedding for *v* the Shannon entropy of these subgraphs is computed for all $m \le h$, where *h* is a given parameter (Bai et al. 2014). A similar concept is applied in (Bai et al. 2015), where depth-based representations are used to compute strengthened vertex labels. Both methods are combined with matching-based techniques to obtain a graph kernel.

Assignment- and matching-based approaches

A common approach to comparing two composite or structured objects is to identify the best possible matching of the components making up the two objects. For example, when comparing two chemical molecules it is instructive to map each atom in one graph to the atom in the other graph that is most similar in terms of, for example, neighborhood structure and attached chemical and physical measurements. This idea has been used also in graph kernels, an early example of which was proposed by Fröhlich et al. (2005) in the *optimal assignment* (OA) kernel. In the OA kernel, each vertex is endowed with a representation (e.g., a label) that is compared using a base kernel. Then, a similarity value for a pair of graphs is computed based on a mapping between their vertices such that the total similarity between the matched vertices with respect to a base kernel is maximized. An illustration of the optimal assignment kernel can be seen in Fig. 4. The OA kernel can be defined as follows.

Definition 2 (Optimal assignment kernel) Let $X = \{x_1, ..., x_n\}$ and $Y = \{y_1, ..., y_n\}$ be sets of components from \mathcal{R} and $k : \mathcal{R} \times \mathcal{R} \to \mathbb{R}$ a base kernel on components. The optimal assignment kernel is

$$K_A(X,Y) = \max_{\pi \in \Pi_n} \sum_{i=1}^n k(x_i, y_{\pi(i)}),$$
(4)

where Π_n is the set of all possible permutations of $\{1, ..., n\}$. In order to apply the assignment kernel to sets of different cardinality, we fill the smaller set with objects z and define k(z, x) = 0 for all $x \in \mathbb{R}$.



The careful reader may have noticed a superficial similarity between the OA kernel and the *R*-convolution and mapping kernels (see in the "Design paradigms for kernels on structured data" section). However, instead of summing the base kernel over a fixed ordering of component pairs, the OA kernel searches for the optimal mapping between components of two objects X, Y. Unfortunately, this means that Eq. 4 is not a positivesemidefinite kernel in general (Vert 2008; Vishwanathan et al. 2010). This fact complicates the use of assignment similarities in kernel methods, although generalizations of SVMs for arbitrary similarity measures have been developed, see, e.g., (Loosli et al. 2015) and references therein. Moreover, kernel methods, such as SVMs, have been found to work well empirically also with indefinite kernels (Johansson and Dubhashi 2015), without enjoying the guarantees that apply to positive definite kernels.

Several different approaches to obtain positive definite graph kernels from indefinite assignment similarities have been proposed. Woźnica et al. (2010) derived graph kernels from set distances and employed a matching-based distance to compare graphs, which was shown to be a metric (Ramon and Bruynooghe 2001). In order to obtain a valid kernel, the authors use so-called prototypes, an idea prevalent also in the theory of learning with (non-kernel) similarity functions under the name landmarks (Balcan et al. 2008). Prototypes are a selected set of instances (e.g., graphs) to which all other instances are compared. Each graph is then represented by a feature vector in which each component is the distance to a different prototype. Prototypes were used also by Johansson and Dubhashi (2015) who proposed to embed the vertices of a graph into the *d*-dimensional real vector space in order to compute a matching between the vertices of two graphs with respect to the Euclidean distance. Several methods for the embedding were proposed; in particular, the authors used Cholesky decompositions of matrix representations of graphs including the graph Laplacian and its pseudo-inverse. The authors found empirically that the indefinite graph similarity matrix from the matching worked as well as prototypes. In the "Experimental study" section, we use this, indefinite version.

Instead of generating feature vectors from prototypes, Kriege et al. (2016) showed that Eq. 4 is a valid kernel for a restricted class of base kernels *k*. These, so-called *strong base kernels*, give rise to hierarchies from which the optimal assignment kernels are computed in linear time by histogram intersection. For graph classification, a base

kernel was obtained from Weisfeiler-Lehman refinement. The derived *Weisfeiler-Lehman* optimal assignment kernel often provides better classification accuracy on real-world benchmark datasets than the Weisfeiler-Lehman subtree kernel (see in the "Experimental study" section). The weights of the hierarchy associated with a strong base kernel can be optimized via multiple kernel learning (Kriege 2019).

Pachauri et al. (2013) studied a generalization of the assignment problem to more than two sets, which was used to define *transitive assignment kernels* for graphs (Schiavinato et al. 2015). The method is based on finding a single assignment between the vertices of all graphs of the dataset instead of finding an optimal assignment for each pairs of graphs. This approach satisfies the transitivity constraint of mapping kernels and therefore leads to positive-semidefinite kernels. However, non-optimal assignments between individual pairs of graphs are possible. Nikolentzos et al. (2017b) proposed a matchingbased approach based on the Earth Mover's Distance, which results in an indefinite kernel function. In order to deal with this they employ a variation of the SVM algorithm, specialized for learning with indefinite kernels. Additionally, they propose an alternative solution based on the *pyramid match kernel*, a generic kernel for comparing sets of features (Grauman and Darrell 2007b). The pyramid match kernel avoids the indefiniteness of other assignment kernels by comparing features through a multi-resolution histograms (with bins determined globally, rather than for each pair of graphs).

Subgraph patterns

In many applications, a strong baseline for representations of composite objects such as documents, images or graphs is one that ignores the structure altogether and represents objects as *bags of components*. A well-known example is the so-called bag-of-words representation of text—statistics of word occurrences without context—which remains a staple in natural language processing. For additional specificity, it is common to compare statistics also of bigrams (sequences of two words), trigrams, etc. A similar idea may be used to compare graphs by ignoring large-scale structure and viewing graphs as bags of vertices or edges. The *vertex label kernel* does precisely this by comparing graphs only at the level of similarity between all pairs of vertex labels from two different graphs,

$$k_{\mathrm{VL}}(G,H) = \sum_{u \in V(G)} \sum_{v \in V(G)} k(l(u), l(v))$$

With the base kernel k the equality indicator function, k_{VL} is a linear kernel on the (unnormalized) distributions of vertex labels in G and H. Similar in spirit, the *edge label kernel* is defined as the sum of base kernel evaluations on all pairs of edge labels (or triplets of the edge label and incident vertex labels). Note that such kernels are a paramount example for instances of the convolution kernel framework, see in the "Design paradigms for kernels on structured data" section.

A downside of vertex and edge label kernels is that they ignore the interplay between structure and labels and are almost completely uninformative for unlabeled graphs. Instead of viewing graphs as bags of vertices or edges, we may view them as bags of *subgraph patterns*. To this end, Shervashidze et al. (2009) introduced a kernel based on counting occurrences of subgraph patterns of a fixed size—so called *graphlets* (see Fig. 5). Every graphlet is an instance of an isomorphism type—a set of graphs that are all isomorphic—such as a graph on three vertices with two edges. While there are three

graphs that connect three vertices with two edges, they are all isomorphic and considered equivalent as graphlets.

Graphlet kernels count the isomorphism types of all induced (possibly disconnected) subgraphs on k > 0 vertices of a graph G. Let $\phi(G)_{\sigma_i}$ for $1 \le i \le N$ denote the number of instances of isomorphism type σ_i where N denotes the number of different types. The kernel computes a feature map $\phi_{GR}(G)$ for G,

$$\phi_{\mathrm{GR}}(G) = \left(\phi(G)_{\sigma_1}, \dots, \phi(G)_{\sigma_N}\right) \,.$$

The graphlet kernel is finally defined as $k_{GR}(G, H) = \langle \phi_{GR}(G), \phi_{GR}(H) \rangle$ for two graphs *G* and *H*.

The time required to compute the graphlet kernel scales exponentially with the size of the considered graphlets. To remedy this, Shervashidze et al. (2009) proposed two algorithms for speeding up the computation time of the feature map for k in {3, 4}. In particular, it is common to restrict the kernel to connected graphlets (isomorphism types). Additionally, the statistics used by the graphlet kernel may be estimated approximately by subgraph sampling, see, e.g., (Johansson et al. 2015; Ahmed et al. 2016; Chen and Lui 2016; Bressan et al. 2017). Please note that the graphlet kernel as proposed by Shervashidze et al. (2009) does not consider any labels or attributes. However, the concept (but not all speed-up tricks) can be extended to labeled graphs by using labeled isomorphism types as features, see, e.g., (Wale et al. 2008). Mapping (sub)graphs to their isomorphism type is known as graph canonization problem, for which no polynomial time algorithm is known (Johnson 2005). However, this is not a severe restriction for small graphs such as graphlets and, in addition, well-engineered algorithms solving most practical instances in a short time exist (McKay and Piperno 2014). Horváth et al. (2004) proposed a kernel which decomposes graphs into cycles and tree patterns, for which the canonization problem can be solved in polynomial time and simple practical algorithms for this are known.

Costa and De Grave (2010) introduced the *neighborhood subgraph pairwise distance kernel* which associates a string with every vertex representing its neighborhood up to a certain depth. In order to avoid solving the graph canonization problem, they proposed using a graph invariant that may, in rare cases, map non-isomorphic neighborhood subgraphs to the same string. Then, pairs of these neighborhood graphs together with



the shortest-path distance between their central vertices are counted as features. The approach is similar to the Weisfeiler-Lehman shortest-path kernel (see in the "Neighborhood aggregation approaches" section).

An alternative to subgraph patterns, tree patterns may contain repeated vertices just like random walks and were initially proposed for use in graph comparison by Ramon and Gärtner (2003) and later refined by Mahé and Vert (2009). Tree pattern kernels are similar to the Weisfeiler-Lehman subtree kernel, but do not consider all neighbors in each step, but also all possible subsets (Shervashidze et al. 2011), and hence do not scale to larger datasets. Da San Martino et al. (2012b) proposed decomposing a graph into trees and applying a kernel defined on trees. In (Da San Martino et al. 2012a), a fast hashing-based computation scheme for the aforementioned graph kernel is proposed.

Walks and paths

A downside of the subgraph pattern kernels described in the previous section is that they require the specification of a set of patterns, or subgraph size, in advance. To ensure efficient computation, this often restricts the patterns to a fairly small scale, emphasizing local structure. A popular alternative is to compare the sequences of vertex or edge attributes that are encountered through traversals through graphs. In this section, we describe two families of traversal algorithms which yield different attribute sequences and thus different kernels—shortest paths and random walks.

Shortest-path kernels

One of the very first, and most influential, graph kernels is the *shortest-path* (SP) kernel (Borgwardt and Kriegel 2005). The idea of the SP kernel is to compare the attributes and lengths of the shortest paths between all pairs of vertices in two graphs. The shortest path between two vertices is illustrated in Fig. 1. Formally, let *G* and *H* be graphs with label function $l: V(G) \cup V(H) \rightarrow \Sigma$ and let d(u, v) denote the shortest-path distance between the vertices *u* and *v* in the same graph. Then, the kernel is defined as

$$k_{\rm SP}(G,H) = \sum_{\substack{(u,v) \in V(G)^2 \ (w,z) \in V(H)^2 \\ u \neq v}} \sum_{\substack{(w,z) \in V(H)^2 \\ w \neq z}} k((u,v),(w,z)),\tag{5}$$

where

$$k((u, v), (w, z)) = k_{\mathrm{L}}(l(u), l(w)) \cdot k_{\mathrm{L}}(l(v), l(z)) \cdot k_{\mathrm{D}}(d(u, v), d(w, z)).$$

Here, $k_{\rm L}$ is a kernel for comparing vertex labels and $k_{\rm D}$ is a kernel to compare shortest-path distances, such that $k_{\rm D}(d(u, v), k(w, z)) = 0$ if $d(u, v) = \infty$ or $d(w, z) = \infty$.

The running time for evaluating the general form of the SP kernel for a pair of graphs is in $\mathcal{O}(n^4)$. This is prohibitively large for most practical applications. However, in the case of discrete vertices and edge labels, e.g., a finite subset of the natural numbers, and k the indicator function, we can compute the feature map $\phi_{\text{SP}}(G)$ corresponding to the kernel explicitly. In this case, each component of the feature map counts the number of triples (l(u), l(v), d(u, v)) for u and v in V(G) and $u \neq v$. Using this approach, the time complexity of the SP kernel is reduced to the time complexity of the Floyd-Warshall algorithm, which is in $O(n^3)$. In (Hermansson et al. 2015) the shortest-path is generalized by considering all shortest paths between two vertices.

Random walk kernels

Gärtner et al. (2003) and Kashima et al. (2003) simultaneously proposed graph kernels based on random walks, which count the number of (label sequences along) walks that two graphs have in common. The description of the random walk kernel by Kashima et al. (2003) is motivated by a probabilistic view of kernels and based on the idea of so-called *marginalized kernels*. The feature space of the kernel comprises all possible label sequences produced by random walks; since the length of the walks is unbounded, the space is of infinite dimension. A method of computation is proposed based on a recursive reformulation of the kernel, which at the end boils down to finding the stationary state of a discrete-time linear system. Since this kernel was later generalized by (Vishwanathan et al. 2010) we do not go into the mathematical details of the original publication. The approach fully supports attributed graphs, since vertex and edge labels encountered on walks are compared by user-specified kernels.

Mahé et al. (2004) extended the original formulation of random walk kernels with a focus on application in cheminformatics (Mahé et al. 2005) to improve the scalability and relevance as similarity measure. A mostly unfavorable characteristic of random walks is that they may visit the same vertex several times. Walks are even allowed to traverse an edge from *u* to *v* and instantly return to *u* via the same edge, a problem referred to as *tottering*. These repeated consecutive vertices do not provide useful information and may even harm the validity as similarity measure. Hence, the marginalized graph kernel was extended to avoid tottering by replacing the underlying first-order Markov random walk model by a second-order Markov random walk model. This technique to prevent tottering only eliminates walks (v_1, \ldots, v_n) with $v_i = v_{i+2}$ for some *i*, but it does not require the considered walks to be paths, i.e., repeated vertices still occur.

Like other random walk kernels, Gärtner et al. (2003) define the feature space of their kernel as the label sequences derived from walks, but propose a different method of computation based on the *direct product graph* of two labeled input graphs.

Definition 3 (Direct Product Graph) For two labeled graphs G = (V, E) and H = (V', E') the direct product graph is denoted by $G \times H = (V, \mathcal{E})$ and defined as

$$\mathcal{V} = \left\{ (v, v') \in V \times V' \mid l(v) = l(v') \right\}$$

$$\mathcal{E} = \left\{ ((u, u'), (v, v')) \in \mathcal{V} \mid (u, v) \in E \land (u', v') \in E' \land l((u, v)) = l((u', v')) \right\}.$$

A vertex (edge) in $G \times H$ has the same label as the corresponding vertices (edges) in G and H.

The concept is illustrated in Fig. 6. There is a one-to-one correspondence between walks in $G \times H$ and walks in the graphs G and H with the same label sequence. The *direct product kernel* is then defined as

$$K_{\text{RW}}(G,H) = \sum_{i,j=1}^{|\mathcal{V}|} \left[\sum_{l=0}^{\infty} \lambda_l A_{\times}^l \right]_{ij},\tag{6}$$

where A_{\times} is the adjacency matrix of $G \times H$ and $\lambda = (\lambda_0, \lambda_1, ...)$ a sequence of weights such that the above sum converges. This is the case for $\lambda_i = \gamma^i$, $i \in \mathbb{N}$, and $\gamma < \frac{1}{a}$ with



 $a \ge \Delta$, where Δ is the maximum degree of $G \times H$. For this choice of weights and with I the identity matrix, there exists a closed-form expression,

$$K_{\rm GRW}(G,H) = \sum_{i,j=1}^{|\mathcal{V}|} \left[(I - \gamma A_{\times})^{-1} \right]_{ij}$$
⁽⁷⁾

which can be computed by matrix inversion. Since the expression reminds of the geometric series transferred to matrices, Eq. 7 is referred to as *geometric random walk kernel*. The running time to compute the geometric random walk kernel between two graphs is dominated by the inversion of the adjacency matrix associated with the direct product graph. The running time is given as roughly $O(n^6)$ (Vishwanathan et al. 2010).

Vishwanathan et al. (2010) propose a generalizing framework for random walk based graph kernels and argue that the approach by Kashima et al. (2003) and Gärtner et al. (2003) can be considered special cases of this kernel. The paper does not address vertex labels and makes extensive use of the Kronecker product between matrices denoted by \otimes and lifts it to the feature space associated with an (edge) kernel. Given an edge kernel $\kappa_{\rm E}$ on attributes from the set \mathcal{A} , let $\phi : \mathcal{A} \to \mathcal{H}$ be a feature map. For an attributed graph G, the feature matrix $\Phi(G)$ is then defined as $\Phi_{ij}(G) = \phi((v_i, v_j))$ if $(v_i, v_j) \in E(G)$ and **0** otherwise. Then, $W_{\times} = \Phi(G) \otimes \Phi(H)$ yields a weight matrix of the direct product graph $G \times H^3$. The proposed kernel is defined as

$$K_{\rm RW}(G,H) = \sum_{l=0}^{\infty} \mu_l \boldsymbol{q}_{\times}^T \boldsymbol{W}_{\times}^l \boldsymbol{p}_{\times}, \qquad (8)$$

where p_{\times} and q_{\times} are initial and stopping probability distributions and μ_l coefficients such that the sum converges. Several methods of computation are proposed, which yield different running times depending on a parameter l, specific to that approach. The parameter leither denotes the number of fixed-point iterations, power iterations or the effective rank of W_{\times} . The running times to compare graphs of order n also depend on the edge labels of the input graphs and the desired edge kernel: For unlabeled graphs the running time $\mathcal{O}(n^3)$ is achieved and $\mathcal{O}(dln^3)$ for labeled graphs, where $d = |\mathcal{L}|$ is the size of the label alphabet. The same running time is attained by edge kernels with a d-dimensional feature space, while $\mathcal{O}(ln^4)$ time is required in the infinite case. For sparse graphs, $\mathcal{O}(ln^2)$ is

³Here vertex labels are ignored, i.e., $V(G \times H) = V(G) \times V(H)$.

achieved in all cases, where a graph *G* is said to be sparse if |E(G)| = O(|V(G)|). Further improvements of the running time were subsequently achieved by non-exact algorithms based on low rank approximations (Kang et al. 2012). Recently, the phenomenon of *halt-ing* in random walk kernels has been studied Sugiyama and Borgwardt (2015), which refers to the fact that walk-based graph kernels may down-weight longer walks so much that their value is dominated by walks of length 1.

The classical random walk kernels described above in theory take all walks without a limitation in length into account, which leads to a high-dimensional feature space. Several application-related papers used walks up to a certain length only, e.g., for the prediction of protein functions (Borgwardt et al. 2005) or image classification (Harchaoui and Bach 2007). These walk based kernels are not susceptible to the phenomenon of halting. Kriege et al. (2014); Kriege et al. (2019) systematically studied kernels based on all the walks of a predetermined fixed length ℓ , referred to as ℓ -walk kernel, and all the walks with length at most ℓ , called *Max*- ℓ -walk kernel, respectively. For these, computation schemes based on implicit and explicit feature maps were proposed and compared experimentally. Computation by explicit feature maps provides a better performance for graphs with discrete labels with a low label diversity and small walk lengths. Conceptually different, Zhang et al. (2018b) derived graph kernels based on return probabilities of random walks.

Kernels for graphs with continuous labels

Most real-world graphs have attributes, mostly real-valued vectors, associated with their vertices and edges. For example, atoms of chemical molecules have physical and chemical properties; individuals in social networks have demographic information; and words in documents carry semantic meaning. Kernels based on pattern counting or neighborhood aggregation are of a discrete nature, i.e., two vertices are regarded as similar if and only if they exactly match, structure-wise as well as attribute-wise. However, in most applications it is desirable to compare real-valued attributes with more nuanced similarity measures such as the Gaussian RBF kernel defined in the "Kernel methods" section.

Kernels suitable for attributed graphs typically rely on user-defined kernels for the comparison of vertex and edge labels. These kernels are then combined with kernels on structure through operations that yield a valid kernel on graphs, such as addition or multiplication. Two examples of this, the recently proposed kernels for attributed graphs, *GraphHopper* (Feragen et al. 2013) and *GraphInvariant* (Orsini et al. 2015), can be expressed as

$$k_{WV}(G,H) = \sum_{\nu \in V(G)} \sum_{\nu' \in V(H)} k_W(\nu,\nu') \cdot k_V(\nu,\nu').$$
(9)

Here, k_V is a user-specified kernel comparing vertex attributes and k_W is a kernel that determines a weight for a vertex pair based on the individual graph structures. Kernels belonging to this family are easily identifiable as instances of *R*-convolution kernels, cf. Definition 1.

For graphs with real-valued attributes, one could set k_V to the Gaussian RBF kernel. The selection of the kernel k_W is essential to take the graph structure into account and allows to obtain different instances of weighted vertex kernels. One implementation of k_W motivated along the lines of *GraphInvariant* (Orsini et al. 2015) is

$$k_W(\nu,\nu') = \sum_{i=0}^h k_\delta(\tau_i(\nu),\tau_i(\nu'))$$

where $\tau_i(\nu)$ denotes the discrete label of the vertex ν after the *i*-th iteration of Weisfeiler-Lehman label refinement of the underlying unlabeled graph. Intuitively, this kernel reflects to what extent the two vertices have a structurally similar neighborhood.

Another graph kernel, which fits into the framework of weighted vertex kernels, is the GraphHopper kernel (Feragen et al. 2013) with

$$k_W(\nu,\nu') = \langle \boldsymbol{M}(\nu), \boldsymbol{M}(\nu') \rangle_F.$$

Here M(v) and M(v') are $\delta \times \delta$ matrices, where the entry $M(v)_{ij}$ for v in V(G) counts the number of times the vertex v appears as the *i*-th vertex on a shortest path of discrete length *j* in *G*, where δ denotes the maximum diameter over all graphs, and $\langle \cdot, \cdot \rangle_F$ is the Frobenius inner product.

Kriege and Mutzel (2012) proposed the *subgraph matching kernel* which is computed by considering all bijections between all subgraphs on at most *k* vertices, and allows to compare vertex attributes using a custom kernel. Moreover, in (Su et al. 2016) the *Descriptor Matching kernel* is defined, which captures the graph structure by a propagation mechanism between neighbors, and uses a variant of the pyramid match kernel (Grauman and Darrell 2007a) to compare attributes between vertices. The kernel can be computed in time linear in the number of edges.

Morris et al. (2016) introduced a scalable framework to compare attributed graphs. The idea is to iteratively turn the continuous attributes of a graph into discrete labels using randomized hash functions. This allows to apply fast explicit graph feature maps, which are limited to graphs with discrete annotations such as the one associated with the Weisfeiler-Lehman subtree kernel (Shervashidze et al. 2011). For special hash functions, the authors obtain approximation results for several state-of-the-art kernels which can handle continuous information. Moreover, they derived a variant of the Weisfeiler-Lehman subtree kernel which can handle continuous attributes.

Other approaches

Kondor et al. (2009) derived a graph kernel using graph invariants based on group representation theory. In (Kondor and Pan 2016), a graph kernel is proposed which is able to capture the graph structure at multiple scales, i.e., neighborhoods around vertices of increasing depth, by using ideas from spectral graph theory. Moreover, the authors provide a low-rank approximation algorithm to scale the kernel computation to large graphs. Johansson et al. (2014) define a graph kernel based on the the Lovász number (Lovász 2006) and provide algorithms to approximate this kernel.

In (Li et al. 2015), a kernel for dynamic graphs is proposed, where vertices and edges are added or deleted over time. The kernel is based on eigen decompositions. Kriege et al. (2014); Kriege et al. (2019) investigated under which conditions it is possible and more efficient to compute the feature map corresponding to a graph kernel explicitly. They provide theoretical as well as empirical results for walk-based kernels. Li et al. (2012) proposed a streaming version of the Weisfeiler-Lehman algorithm using a hashing technique. Aiolli et al. (2015) and Massimo et al. (2016) applied multiple kernel learning to the

graph kernel domain. Nikolentzos et al. (2018) proposed to first build the k-core decomposition of graphs to obtain a hierarchy of nested subgraphs, which are then individually compared by a graph similarity measure. The approach has been combined with several graph kernels such as the Weisfeiler-Lehman subtree kernel and was shown to improve the accuracy on some datasets.

Yanardag and Vishwanathan (2015a) uses recent neural techniques from neural language modeling, such as *skip-gram* (Mikolov et al. 2013). The authors build on known state-of-the-art kernels, but allow to respect relationships between their features. This is demonstrated by hand-designed matrices encoding the similarities between features for selected graph kernels such as the graphlet and Weisfeiler-Lehman subtree kernel. Similar ideas were used in (Yanardag and Vishwanathan 2015b) where smoothing methods for multinomial distributions were applied to the graph domain.

Expressivity of graph kernels

While a large literature has studied the empirical performance of various graph kernels, there exists comparatively few works that deal with graph kernels exclusively from a theoretical point of view. Most works that provide learning guarantees for graph kernels attempt to formalize their *expressivity*.

The *expressivity* of a graph kernel refers broadly to the kernel's ability to distinguish certain patterns and properties of graphs. In an early attempt to formalize this notion, Gärtner et al. (2003) introduced the concept a *complete graph kernel*—kernels for which the corresponding feature map is an injection. If a kernel is not complete, there are non-isomorphic graphs *G* and *H* with $\phi(G) = \phi(H)$ that cannot be distinguished by the kernel. In this case there is no way any classifier based on this kernel can separate these two graphs. However, computing a complete graph kernel is GI-hard, i.e., at least as hard as deciding whether two graphs are isomorphic (Gärtner et al. 2003). For this problem no polynomial time algorithm for general graphs is known (Johnson 2005). Therefore, none of the graph kernels used in practice are complete. Note however, that a kernel may be injective with respect to a finite or restricted family of graphs.

As no practical kernels are complete, attempts have been made to characterize expressivity in terms of which graph properties can be distinguished by existing graph kernels. In (Kriege et al. 2018), a framework to measure the expressivity of graph kernels based on ideas from property testing was introduced. The authors show that graph kernels such as the Weisfeiler-Lehman subtree, the shortest-path and the graphlet kernel are not able to distinguish basic graph properties such as planarity or connectedness. Based on these results they propose a graph kernel based on frequency counts of the isomorphism type of subgraphs around each vertex up to a certain depth. This kernel is able to distinguish the above properties and computable in polynomial time for graphs of bounded degree. Finally, the authors provide learning guarantees for 1-nearest neighborhood classifiers. Similarly, (Johansson and Dubhashi 2015) gave bounds on the classification margin obtained when using the optimal assignment kernel, with Laplacian embeddings, to classify graphs with different densities or random graphs with and without planted cliques. In Johansson et al. (2014), the authors studied global properties of graphs such as girth, density and clique number and proposed kernels based on vertex embeddings associated with the Lovász- ϑ and SVM- ϑ numbers which have been shown to capture these properties.

The expressivity of graph kernels has been studied also from statistical perspectives. In particular, Oneto et al. (2017) use well-known results from statistical learning theory to give results which bound measures of expressivity in terms of Rademacher complexity and stability theory. Moreover, they apply their theoretical findings in an experimental study comparing the estimated expressivity of popular graph kernels, confirming some of their known properties. Finally, Johansson et al. (2015) studied the statistical tradeoff between expressivity and differential privacy (Dwork et al. 2014).

Applications of graph kernels

The following section outlines a *non-exhaustive* list of applications of the kernels described in the "Graph kernels" section, categorized by scientific area.

Chemoinformatics Chemoinformatics is the study of chemistry and chemical compounds using statistical and computational resources (Brown 2009). An important application is drug development in which new, untested medical compounds are modeled in silico before being tested in vitro or in animal tests. The primary object of study-the molecule—is well represented by a graph in which vertices take the places of atoms and edges that of bonds. The chemical properties of these atoms and bonds may be represented as vertex and edge attributes, and the properties of the molecule itself through features of the structure and attributes. The graphs derived from small molecules have specific characteristics. They typically have less than 50 vertices, their degree is bounded by a small constant (<4 with few exceptions), and the distribution of vertex labels representing atom types is specific (e.g., most of the atoms are carbon). Almost all molecular graphs are planar, most of them even outerplanar (Horváth et al. 2010), and they have a tree-like structure (Yamaguchi et al. 2003). Molecular graphs are not only a common benchmark dataset for graph kernels, but several kernels were specifically proposed for this domain, e.g., (Horváth et al. 2004; Swamidass et al. 2005; Ceroni et al. 2007; Mahé and Vert 2009; Fröhlich et al. 2005). The pharmacophore kernel was introduced by Mahé et al. (2006) to compare chemical compounds based on characteristic features of vertices together with their relative spatial arrangement. As a result, the kernel is designed to handle with continuous distances. The pharmacophore kernel was shown to be an instance of the more general subgraph matching kernel (Kriege and Mutzel 2012). Mahé and Vert (2009) developed new tree pattern kernels for molecular graphs, which were then applied in toxicity and anti-cancer activity prediction tasks. Kernels for chemical compounds such as this have been successfully employed for various tasks in cheminformatics including the prediction of mutagenicity, toxicity and anti-cancer activity (Swamidass et al. 2005).

However, such tasks have been addressed by computational methods long before the advent of graph kernels, cf. Fig. 2. So-called *fingerprints* are a well-established classical technique in cheminformatics to represent molecules by feature vectors (Brown 2009). Commonly features are obtained by (i) enumeration of all substructures of a certain class contained in the molecular graphs, (ii) taken from a predefined dictionary of relevant substructures or (iii) generated in a preceding data-mining phase. Fingerprints are then used to encode the number of occurrences of a feature or only its presence or absence by a single bit per feature. Often hashing is used to reduce the fingerprint length to a fixed size at the cost of information loss (see, e.g., (Daylight 2008)). Such fingerprints are typically compared using similarity measures such as the *Tanimoto coefficient*, which

are closely related to kernels (Ralaivola et al. 2005). Approaches of the first category are, e.g., based on all paths contained in a graph (Daylight 2008) or all subgraphs up to a certain size (Wale et al. 2008), similar to graphlets. Ralaivola et al. (2005) experimentally compared random walk kernels to kernels derived from path-based fingerprints and has shown that these reach similar classification performance on molecular graph datasets. *Extended connectivity fingerprints* encode the neighborhood of atoms iteratively similar to the graph kernels discussed in the "Neighborhood aggregation approaches" section and can be considered a standard tool in cheminformatics for decades (Rogers and Hahn 2010). Predefined dictionaries compiled by experts with domain-specific knowledge exist, e.g., MACCS/MDL Keys for drug discovery (Durant et al. 2002).

Bioinformatics Understanding proteins, one of the fundamental building blocks of life, is a central goal in bioinformatics. Proteins are complex molecules which are often represented in terms of larger components such as helices, sheets and turns. Borgwardt et al. (2005) model protein data as graphs where each vertex represents such a component, and each edge indicates proximity in space or in amino acid sequence. Both vertices and edges are annotated by categorical and real-valued attributes. The authors used a modified random walk kernel to classify proteins as enzymes or non-enzymes. In related work, Borgwardt et al. (2007) predict disease outcomes from protein-protein interaction networks. Here, each vertex is a protein and each edge the physical interaction between a protein-protein pair. In order to take missing edges into account, which is crucial for studying protein-protein-interaction networks, the kernel

$$K_{\rm CP}(G,H) = K_{\rm RW}(G,H) + K_{\rm RW}(G,\overline{H})$$

was proposed, which is the sum of a random walk kernel K_{RW} applied to the original graphs G and H as well as to their complement graphs \overline{G} and \overline{H} . Studying pairs of complement graphs may be useful also in other applications.

Neuroscience The connectivity and functional activity between neurons in the human brain are indicative of diseases such as Alzheimer's disease as well as subjects' reactions to sensory stimuli. For this reason, researchers in neuroscience have studied the similarities of brain networks among human subjects to find patterns that correlate with known differences between them. Representing parts of the brain as vertices and the strength of connection between them as edges, several authors have applied graph kernels for this purpose (Vega-Pons et al. 2014; Takerkart et al. 2014; Vega-Pons and Avesani 2013; Wang et al. 2016; Jie et al. 2016). Unlike many other applications, the vertices in brain networks often have an identity, representing a specific part of the brain. Jie et al. (2016) exploited this fact in learning to classify mild cognitive impairments (MCI). They find that their proposed kernel, based on iterative neighborhood expansion (similar to the Weisfeiler-Lehman kernel), which exploits the one-to-one mapping of vertices (brain regions) between different graphs consistently outperforms baseline kernels in this task.

Natural language processing Natural language processing is ripe with relational data: words in a document relate through their location in text, documents relate through their

publication venue and authors, named entities relate through the contexts in which they are mentioned. Graph kernels have been used to measure similarity between all of these concepts. For example, Nikolentzos et al. (2017a) use the shortest-path kernel to compute document similarity by converting each document to a graph in which vertices represent terms and two vertices are connected by an edge if the corresponding terms appear together in a fixed-size window. Hermansson et al. (2013) used the co-occurrence network of person names in a large news corpus to classify which names belong to multiple individuals in the database. Each name was represented by the subgraph corresponding to the neighborhood of co-occuring names and labeled by domain experts. The output of the system was intended for use as preprocessing to an entity disambiguation system. In (Li et al. 2016) the Weisfeiler-Lehman subtree kernel was used to define a similarity function for call graphs of Java programs to identify similar call graphs. de Vries (2013) extended the Weisfeiler-Lehman subtree kernel so that it can handle RDF data.

Computer vision Harchaoui and Bach (2007) applied kernels based on walks of a fixed length to image classification and developed a dynamic programming approach for their computation. The also modified tree pattern kernels for image classification, where graphs typically have a fixed embedding in the plane. Wu et al. (2014) proposed graph kernels for human action recognition in video sequences. To this end, they encode the features of each frame as well as the dynamic changes between successive frames by separate graphs. These graphs are compared by a linear combination of random walk kernels using multiple kernel learning, which leads to an accurate classification of human actions. The propagation kernel was applied to predict object categories in order to facilitate robot grasping (Neumann et al. 2013). To this end, 3D point cloud data was represented by *k*-nearest neighbor graphs.

Experimental study

In our experimental study, we investigate various kernels considered to be state-of-theart in detail and compare them to simple baseline methods using vertex and edge label histograms. We would like to answer the following research questions.

- Q1 **Expressivity.** Are the proposed graph kernels sufficiently expressive to distinguish the graphs of common benchmark datasets from each other according to their labels and structure?
- Q2 **Non-linear decision boundaries.** Can the classification accuracy of graph kernels be improved by finding non-linear decision boundaries in their feature space?
- Q3 Accuracy. Is there a graph kernel that is superior over the other graph kernels in terms of classification accuracy? Does the answer to Q1 explain the differences in prediction accuracy?
- Q4 **Agreement.** Which graph kernels predict similarly? Do different graph kernels succeed and fail for the same graphs?
- Q5 **Continuous attributes.** Is there a kernel for graphs with continuous attributes that is superior over the other graph kernels in terms of classification accuracy?

Methods

We describe the methods we used to answer the research questions and summarize our experimental setup.

Classification accuracy

In order to answer several of our research questions, it is necessary to determine the prediction accuracy achieved by the different graph kernels. We performed classification experiments using the *C*-SVM implementation LIBSVM (Chang and Lin 2011). We used nested cross-validation with 10 folds in the inner and outer loop. In the inner loop the kernel parameters and the regularization parameter *C* were chosen by cross-validation based on the training set for the current fold. In the same way it was determined whether the kernel matrix should be normalized. The parameter *C* was chosen from $\{10^{-3}, 10^{-2}, \ldots, 10^3\}$. We repeated the outer cross-validation ten times with different random folds, and report average accuracies and standard deviations.

Complete graph kernels

The theoretical concept of complete graph kernels has little practical relevance and is not suitable for answering Q1. Therefore we generalize the concept of complete graph kernels. For a given dataset $\mathcal{D} = \{(G_1, y_1), \ldots, (G_n, y_n)\}$ of graphs G_i with class labels $y_i \in \mathcal{Y}$ for all $1 \leq i \leq n$, we say a graph kernel K with a feature map ϕ is *complete for* \mathcal{D} if for all graphs G_i, G_j the implication $\phi(G_i) = \phi(G_j) \Longrightarrow i = j$ holds; it is *label complete for* \mathcal{D} if for all graphs G_i, G_j the implication $\phi(G_i) = \phi(G_j) \Longrightarrow y_i = y_j$ holds. Note that we may test whether $\phi(G_i) = \phi(G_j)$ holds using the kernel trick without constructing the feature vectors. For a kernel K on \mathcal{X} with a feature map $\phi : \mathcal{X} \to \mathcal{H}$ the *kernel metric* is

$$d_K(x, y) = \|\phi(x) - \phi(y)\|$$
(10)

$$= \sqrt{K(x,x) + K(y,y) - 2K(x,y)}.$$
(11)

Therefore, $\phi(G) = \phi(H)$ if and only if K(G, G) + K(H, H) - 2K(G, H) = 0. We define the *(label) completeness ratio* of a graph kernel w.r.t. a dataset as the fraction of graphs in the dataset that can be distinguished from all other graphs (with different class labels) in the dataset.

We investigate how these measures align with the observed prediction accuracy. Note that the label completeness ratio limits the accuracy of a kernel on a specific dataset. Vice versa, classifiers based on complete kernels not necessarily achieve a high accuracy. A kernel that is one for two isomorphic graphs and zero otherwise, for example, would achieve the highest possible completeness ratio, but is too strict for learning, cf. "Design paradigms for kernels on structured data" section. Moreover, a complete graph kernel not necessarily maps graphs in different classes to feature vectors that are linearly separable. In this case (an additional) mapping in a high-dimensional feature space might improve the accuracy.

Non-linear decision boundaries in the feature space of graph kernels

Many graph kernels explicitly compute feature vectors and thus essentially transform graph data to vector data, cf. "Graph kernels" section. Typically, these kernels then just apply the linear kernel to these vectors to obtain a graph kernel. This is surprising since it is well-known that for vector data often better results can be obtained by a polynomial or Gaussian RBF kernel. These, however, are usually not used in combination with graph kernels. Sugiyama and Borgwardt (2015) observed that applying a Gaussian RBF kernel to vertex and edge label histograms leads to a clear improvement over linear kernels. Moreover, for some datasets the approach was observed to be competitive with random walk

kernels. Going beyond the application of standard kernels to graph feature vectors, Kriege (2015) proposed to obtain modified graph kernels also from those based on implicit computation schemes by employing the kernel trick, e.g., by substituting the Euclidean distance in the Gaussian RBF kernel by the metric associated with a graph kernel. Since the kernel metric can be computed without explicit feature maps, any graph kernel can thereby be modified to operate in a different (high-dimensional) feature space. However, the approach was generally not employed in experimental evaluations of graph kernels. Only recently, Nikolentzos and Vazirgiannis (2018) presented first experimental results of the approach for the shortest-path, Weisfeiler-Lehman and pyramid match graph kernel using a polynomial and Gaussian RBF kernel for successive embedding. Promising experimental results were presented, in particular, for the Gaussian RBF kernel. We present an in detail evaluation of the approach on a wide range of graph kernels and datasets.

We apply the Gaussian RBF kernel to the feature vectors associated with graph kernels by substituting the Euclidean distance in Eq. (1) by the metric associated with graph kernels. Note that the kernel metric can be computed from feature vectors according to Eq. (10) or by employing the kernel trick according to Eq. (11). In order to study the effect of this modification experimentally, we have modified the computed kernel matrices as described above. The parameter σ was selected from $\{2^{-7}, 2^{-6}, \ldots, 2^7\}$ by cross-validation in the inner cross-validation loop based on the training data sets.

Datasets

In our experimental evaluation, we have considered graph data from various domains, most of which has been used previously to compare graph kernels. Moreover, we derived new large datasets from the data published by the National Center for Advancing Translational Sciences in the context of the *Tox21 Data Challenge 2014*⁴ initiated with the goal to develop better toxicity assessment methods for small molecules. These datasets each contain more than 7000 graphs and thus exceed the size of the datasets typically used to evaluate graph kernels. We have made all datasets publicly available (Kersting et al. 2016). Their statistics are summarized in Table 2.

The datasets AIDS, BZR, COX2, DHFR, Mutagenicity, MUTAG, NCI1, NCI109, PTC and Tox21 are graphs derived from small molecules, where class labels encode a certain biological property such as toxicity and activity against cancer cells. The vertices and edges of the graphs represent the atoms and their chemical bonds, respectively, and are annotated by their atom and bond type. The datasets DD, ENZYMES and PRO-TEINS represent macromolecules using different graph models. Here, the vertices either represent protein tertiary structures or amino acids and the edges encode spatial proximity. The class labels are the 6 EC top-level classes or encode whether a protein is an enzyme. The datasets REDDIT-BINARY, IMDB-BINARY and IMDB-MULTI are derived from social networks. The MSRC datasets are associated with computer vision tasks. Images are encoded by graphs, where vertices represent superpixels with a semantic label and edges their adjacency. Finally, SYNTHETICnew and Synthie are synthetically generated graphs with continuous attributes. FRANKENSTEIN contains graphs derived from small molecules, where atom types are represented by high dimensional vectors of pixel intensities of associated images.

Dataset	Properties				Labels		Attributes		Dof
	Graphs	Clas.	Avg. V	Avg. <i> E</i>	Vertex	Edge	Vertex	Edge	1161.
AIDS	2000	2	15.69	16.20	+	+	+ (4)	I	(Riesen and Bunke 2008)
BZR	405	2	35.75	38.36	+	I	+ (3)	I	(Sutherland et al. 2003)
COX2	467	2	41.22	43.45	+	I	+ (3)	I	(Sutherland et al. 2003)
DHFR	467	2	42.43	44.54	+	I	+ (3)	I	(Sutherland et al. 2003)
DD	1178	2	284.32	715.66	+	I	I	I	(Dobson and Doig 2003; Shervashidze et al. 2011)
ENZYMES	600	9	32.63	62.14	+	I	+ (18)	I	(Borgwardt et al. 2005; Schomburg et al. 2004)
FRANKENSTEIN	4337	2	16.90	17.88	I	I	+ (780)	I	(Orsini et al. 2015)
IMDB-BINARY	1000	2	19.77	96.53	I	I	I	I	(Yanardag and Vishwanathan 2015a)
IMDB-MULTI	1500	ŝ	13.00	65.94	I	I	I	I	(Yanardag and Vishwanathan 2015a)
Mutagenicity	4337	2	30.32	30.77	+	+	Ι	I	(Riesen and Bunke 2008; Kazius et al. 2005)
MSRC-9	221	8	40.58	97.94	+	I	I	I	(Neumann et al. 2016)
MSRC-21	563	20	77.52	198.32	+	I	I	I	(Neumann et al. 2016)
MSRC-21C	209	20	40.28	96.60	+	I	I	I	(Neumann et al. 2016)
MUTAG	188	2	17.93	19.79	+	+	I	I	(Debnath et al. 1991; Kriege and Mutzel 2012)
NC11	4110	2	29.87	32.30	+	I	Ι	I	(Shervashidze et al. 2011)
NCI109	4127	2	29.68	32.13	+	I	I	I	(Shervashidze et al. 2011)
PTC-FM	349	2	14.11	14.48	+	+	Ι	I	(Helma et al. 2001; Kriege and Mutzel 2012)
PTC-FR	351	2	14.56	15.00	+	+	Ι	I	(Helma et al. 2001; Kriege and Mutzel 2012)
PTC-MM	336	2	13.97	14.32	+	+	I	Ι	(Helma et al. 2001; Kriege and Mutzel 2012)
PTC-MR	344	2	14.29	14.69	+	+	Ι	I	(Helma et al. 2001; Kriege and Mutzel 2012)
PROTEINS	1113	2	39.06	72.82	+	I	+ (1)	I	(Borgwardt et al. 2005; Dobson and Doig 2003)
REDDIT-BINARY	2000	2	429.63	497.75	I	I	I	I	(Yanardag and Vishwanathan 2015b)
SYNTHETICnew	300	2	100.00	196.25	I	I	+ (1)	I	(Feragen et al. 2013)
Synthie	400	4	95.00	173.92	I	I	+ (15)	I	(Morris et al. 2016)
Tox21-AR	9362	2	18.39	18.84	+	+	Ι	I	(Tox21 Data Challenge 2014)
Tox21-MMP	7320	2	17.49	17.83	+	+	I	I	(Tox21 Data Challenge 2014)
Tox21-AHR	8169	2	18.09	18.50	+	+	I	I	(Tox21 Data Challenge 2014)

Table 2 Dataset statistics and properties

Graph kernels

As a baseline we included the *vertex label kernel* (VL) and *edge label kernel* (EL), which are the dot products on vertex and edge label histograms, respectively. An edge label is a triplet consisting of the labels of the edge and the label of its two endpoints. We used the Weisfeiler-Lehman subtree (WL) and Weisfeiler-Lehman optimal assignment kernel (WL-OA), see in the "Neighborhood aggregation approaches" section. For both the number of refinement operations was chosen from $\{0, 1, \ldots, 8\}$ by cross-validation. In addition we implemented a *graphlet kernel* (GL3) and the shortest-path kernel (SP) (Borgwardt and Kriegel 2005). GL3 is based on connected subgraphs with three vertices taking labels into account similar to the approach used by Shervashidze et al. (2011). For SP we used the indicator function to compare path lengths and computed the kernel by explicit feature maps in case of discrete vertex labels, cf. (Shervashidze et al. 2011). These kernels were implemented in Java based on the same common data structures and support both vertex labels and—with exception of VL and SP—edge labels.

We compare three kernels based on matching of vertex embeddings, the matching kernel of Johansson and Dubhashi (2015) with inverse Laplacian (MK-IL) and Laplacian (MK-L) embeddings and the Pyramid Match (PM) kernel of (Nikolentzos et al. 2017b). The MK kernels lack hyperparameters and for the PM-kernel, we used the default settings—vertex embedding dimension (d = 6) and matching levels (L = 3)—in the implementation by Nikolentzos (2016). Finally, we include the shortest-path variant of the Deep Graph Kernel (DeepGK) (Yanardag and Vishwanathan 2015a) with parameters as suggested in Yanardag (2015) (SP feature type, MLE kernel type, window size 5, 10 dimensions)⁵, the DBR kernel of Bai et al. (2014) (no parameters, code obtained through correspondence) and the propagation kernel (Prop) (Neumann et al. 2016; Neumann 2016) for which we select the number of diffusion iterations by cross-validation and use the settings recommended by the authors for other hyperparameters.

In a comparison of kernels for graphs with continuous vertex attributes we use the shortest-path kernel (Borgwardt and Kriegel 2005) with a Gaussian RBF base kernel to compare vertex attributes, see also "Shortest-path kernels" section, the GraphHopper kernel (Feragen et al. 2013), the GraphInvariant kernel (Orsini et al. 2015), the Propagation kernel (P2K) (Neumann et al. 2016), and the Hash Graph kernel (Morris et al. 2016). We set the parameter σ of the Gaussian RBF kernel to $\sqrt{D/2}$ for the GraphHopper and the GraphInvariant kernel, as reported in (Feragen et al. 2013; Orsini et al. 2015), where D denotes the number of components of the vertex attributes. For datasets that do not have vertex labels, we either used the vertex degree instead or uniform labels (selected by (double) cross-validation). Following (Morris et al. 2016), we set the number of iteration for the Hash Graph kernel to 20 for all datasets, excluding the Sythnie datasets where we used 100.

Results and discussion

We present our experimental results and discuss the research questions.

Q1 Expressivity. For these experiments we only considered kernels that are permutation-invariant and guarantee that two isomorphic graphs are represented by the

⁵We did not perform a parameter search for the parameters of the Deep Graph kernel and the accuracy of the kernel may improve with a more tailored choice.

same feature vector. This is not the case for the MK-* and PM kernels because of the vertex embedding techniques applied.

Figure 7 shows the completeness ratio of various permutation invariant graph kernels with different parameters on the datasets as a heatmap. The WL-OA kernels achieved the same results as the WL kernels and are therefore not depicted. As expected, VL achieves only a weak completeness ratio, since it ignores the graph structure completely. To a lesser extent, this also applies to EL and GL₃. The SP and the WL_h kernels with $h \ge 2$ provide a high completeness ratio close to one for most datasets. However, for the IMDB-BINARY dataset shortest-paths appear to be less powerful features than small local graphlets. This indicates structural differences between this dataset and the molecular graph datasets, where SP consistently achieves better results than GL_3 . As expected DeepGK performs similar to the SP kernel. WL and Prop are both based on a neighborhood aggregation mechanism, but WL achieves a higher completeness ratio on several datasets. This is explained by the fact that Prop does not support edge labels and does not employ a relabeling function after each propagation step. DBR does not take labels into account and consequently fails to distinguish many graphs of the datasets, for which vertex labels are informative. The difficulty of distinguishing the graphs in a dataset varies strongly based on the type of graphs. The computer vision graphs are almost perfectly distinguished by just considering the vertex label multiplicities, molecular graphs often require multiple iterations of Weisfeiler-Lehman or global features such as shortest paths. For social networks, the REDDIT-BINARY graphs are also effectively distinguished by Weisfeiler-Lehman refinement, while this is not possible for the two IMDB datasets. However, we observed that all the graphs in these two datasets that cannot be distinguished by WL_1



are in fact isomorphic. Therefore, a higher completeness ratio cannot be achieved by any permutation-invariant graph kernel.

We now consider the label completeness ratio depicted in Fig. 8. The label completion ratio generally shows the same trends, but higher values close to one are reached as expected. For the datasets IMDB-BINARY and IMDB-MULTI we have already observed that WL₁ distinguishes all non-isomorphic graphs. As we see in Fig. 8 these datasets contain a large number of isomorphic graphs that actually belong to different classes. Apparently, the information contained in the dataset is not sufficient to allow perfect classification. A general observations from the heatmaps is that WL (just as WL-OA) effectively distinguish most graphs after only few iterations of refinement. For some nonchallenging datasets even VL and EL are sufficient expressive. Therefore, these kernels are interesting baselines for accuracy experiments. In order to effectively learn with a graph kernel, it is not sufficient to just distinguish graphs, which may lead to strong overfitting, but to provide a smooth similarity measure that allows the classifier to generalize to unseen data.

Q2 Non-linear decision boundaries. We discuss the accuracy results of the classification experiments summarized in Tables 3 and 4. The classification accuracy of the simple kernels VL and EL can be drastically improved by combining them with the Gaussian RBF kernel for several datasets. A clear improvement is also achieved for GL3 on an average. For WL and WL-OA the Gaussian RBF kernel only leads to minor changes in classification accuracy for most datasets. However, a strong improvement is observed for WL and the dataset ENZYMES, even lifting the accuracy above the



	, , , , , , , , , , , , , , , , , , , ,	1 41-14 0 141	0		5	5	<u>जजरक सार्व्</u> स	5055)			
Dataset	>	٦	ш	I	S.	d_	1	۲	-ML-	OA	013	
[מינטרי	Klin	K _{RBF}	Klin	K _{RBF}	Klin	K _{RBF}	Klin	KRBF	Klin	K _{RBF}	Klin	K _{RBF}
NCI1	64.6土0.1	67.2±2.8	66.3±0.1	71.8土0.3	73.2±0.3	79.3土0.4	85.9±0.1	86.2±0.1	86.2±0.2	86.6 ±0.2	70.5±0.2	76.5土0.4
NCI109	63.6土0.2	68.9土1.4	64.9土0.1	71.4土0.5	72.7土0.3	77.6土0.3	85.9±0.3	86.0土0.3	86.2±0.2	86.4 ±0.2	69.3土0.2	76.0土0.4
PTC-FR	67.9±0.4	66.9±0.5	66.8土0.5	65.2±1.2	67.1±2.0	63.7±2.0	67.1土1.2	66.8土1.5	67.8土1.1	67.0土1.3	65.5土0.9	65.0土1.4
PTC-MR	57.8土0.9	59.4土1.4	56.7土1.6	60.5土1.8	58.8土2.2	62.0土1.8	60.4土1.5	62.7 ±2.0	62.6土1.5	62.7 ±1.0	57.4土1.6	60.4土1.6
PTC-FM	63.9土0.5	62.6土0.9	64.5土0.4	60.5土1.4	62.7土1.0	60.2土1.3	62.8土1.2	60.9±0.8	61.6土1.2	61.7土1.2	60.2±3.0	60.7±0.8
PTC-MM	66.6土0.8	64.7土0.4	64.1土1.0	62.7土1.6	63.3土1.2	63.2±0.8	67.8 ±2.1	67.7土1.3	66.4土1.1	66.3土1.7	61.4土1.7	61.3土1.4
MUTAG	85.4土0.7	82.9土1.0	83.6土1.0	88.4土2.2	83.1土1.3	85.2土1.4	86.6土0.6	87.9土1.0	87.5±2.1	87.3土1.7	87.2土1.1	87.8土1.1
Mutagenicity	67.0土0.2	73.9土0.3	72.4土0.1	80.3±0.3	77.4土0.2	80.1土0.2	83.6土0.2	84.5土0.3	84.2±0.2	84.7 ±0.4	79.8土0.2	82.7土0.3
AIDS	99.7 ±0.0	99.7 ±0.0	99.5土0.0	99.4土0.0	0.0千9.66	99.7 ±0.0	99.7 ±0.0	99.7 ±0.0	99.7 ±0.0	99.7 ±0.0	99.2土0.1	99.3土0.1
BZR	78.8土0.1	86.0土0.2	79.1土0.5	86.3土0.3	86.5土0.9	88.1土0.5	88.5 ±0.7	87.9土0.8	88.2土0.4	88.0土0.5	81.6土0.7	85.4土1.0
COX2	78.2土0.0	80.6土0.3	82.0土0.6	83.9 ±0.7	80.6土0.9	81.7±0.8	81.2土1.0	81.7±0.7	80.4土0.9	80.8土1.3	81.3土0.7	81.9土0.5
DHFR	60.9±0.2	74.8土1.2	67.9±0.6	73.2土0.9	77.5土0.6	80.7±0.7	82.7±0.4	83.5 ±0.6	83.0土1.0	83.3±0.6	74.7土0.6	81.2±1.0
DD	78.2土0.4	80.1土0.4	77.5土0.6	78.7±0.7	79.5土0.6	74.5土0.2	78.9±0.4	80.9土0.3	79.2土0.4	79.9土0.5	79.7土0.7	79.1土0.6
PROTEINS	71.9±0.4	74.7土0.4	73.4土0.3	75.2±0.5	75.9土0.4	74.0土0.3	75.5土0.3	73.9±0.7	76.2土0.4	75.9土0.6	72.7土0.6	73.0土0.6
ENZYMES	23.4土1.1	41.7土1.1	27.7±0.7	45.1土1.2	41.9土1.7	59.5土1.3	53.7土1.5	62.6土1.2	59.9土1.0	62.3土1.1	30.4土1.1	58.6土1.0
IMDB-BINARY	46.3土0.9	56.5土0.6	46.0土0.9	62.6土1.2	57.3土0.6	70.2土0.8	72.9土0.6	71.3土1.0	73.1土0.7	73.5 ±0.6	59.4土0.4	70.1土0.8
IMDB-MULTI	31.9±0.5	39.5土0.9	30.8土0.9	46.9土0.6	39.6土0.2	46.1土0.7	50.3土0.4	50.7土0.6	50.4土0.5	50.7±0.5	40.6土0.4	47.1土0.5
REDDIT-BINARY	75.3土0.1	77.6土0.2	75.1土0.1	79.4土0.1	81.7±0.2	67.8土0.2	80.9土0.4	83.9土0.5	89.3 ±0.2	88.9土0.1	60.1土0.2	73.6土0.1
MSRC-9	88.4土1.3	87.7±1.0	92.6 ±0.9	90.2±0.7	91.4土0.8	89.2土1.0	90.1土0.8	89.1土0.9	90.7±0.8	90.1土0.7	91.6土0.7	91.6土0.9
MSRC-21	89.4土0.3	90.0土0.5	89.5土0.3	87.3土0.4	89.4土0.6	37.4土1.2	89.3土0.6	89.8土0.4	9.0干0.06	90.5土0.4	90.5±0.7	85.1土0.6
MSRC-21C	81.2土1.2	80.8土1.7	84.5土0.8	81.8土1.3	83.8土1.2	78.3土1.3	81.9±0.9	82.1土1.1	84.9土0.8	84.5土1.0	84.0土1.7	82.6土1.0
Tox21-AR	95.9±0.0	96.4土0.0	95.9±0.0	97.5±0.0	97.1土0.0	97.5土0.0	97.9±0.0	98.0 ±0.0	98.0 ±0.0	98.0 ±0.0	96.4土0.0	97.6土0.0
Tox21-MMP	84.3土0.0	86.5土0.1	84.5土0.0	89.7±0.2	86.4土0.1	90.7±0.2	92.5土0.1	93.0 ±0.2	92.7±0.1	92.8土0.1	87.3土0.1	91.2±0.2
Tox21-AHR	88.4土0.0	89.1土0.2	88.6土0.1	91.4土0.2	88.4土0.0	91.9±0.1	93.4土0.1	93.7 ±0.1	93.5±0.1	93.6土0.1	89.7土0.1	92.8±0.2
Average	71.2	74.5	72.2	76.2	75.6	74.9	79.6	80.2	80.5	80.6	73.8	77.5
The highest accuracy	/ value is highlight	ted in bold for each	h dataset									

Table 3 Classification accuracy and standard deviation for several kernels and their variant when plugged into the Gaussian RBF kernel

Kriege et al. Applied Network Science (2020) 5:6

						-)					
Datacot	MI	K-IL	MI	K-L	PI	Σ	Dee	pGK	DB	3R a	Prop	
Dataset	Klin	K _{RBF}	Klin	K _{RBF}	Klin	K _{RBF}	Klin	K _{RBF}	\mathcal{K}_{lin}	K _{RBF}	Klin	K _{RBF}
NCI1	76.8±0.3	78.1±0.3	72.8±0.3	75.5±0.3	73.3±0.3	80.0土0.3	74.9±0.2	78.4土0.3	67.4±0.3	76.0土0.2	84.6土0.2	85.6±0.2
NCI109	75.4土0.3	75.9土0.3	71.9±0.3	74.1土0.4	71.1土0.3	78.8土0.2	73.3±0.2	77.7±0.3	66.6土0.2	75.3土0.2	84.1土0.2	84.7±0.2
PTC-FR	69.1 ±0.6	68.6土0.7	68.4土0.7	68.1土0.9	65.9±0.8	65.0土1.1	66.4土1.2	63.8土1.5	65.3土0.5	64.0土0.7	66.0土1.5	65.6土1.0
PTC-MR	60.6土1.0	60.5土1.1	59.1土1.7	59.7土1.4	61.5土1.3	59.5土2.0	59.9土1.5	60.9土1.7	53.7土1.3	55.1土1.9	59.9土1.6	61.3土2.1
PTC-FM	58.6土1.8	64.7 ±0.6	60.4土0.9	61.4土1.6	59.7土1.4	62.2土0.8	62.6土0.9	60.9土1.1	56.2土1.5	59.8土1.5	60.9土1.6	61.7土1.6
PTC-MM	62.1土1.6	65.0土1.4	63.8土1.2	63.1±0.7	62.9土1.3	62.2土1.0	63.3土0.9	61.8土1.4	59.4土1.1	63.5土1.0	63.9土1.0	64.6土1.9
MUTAG	82.8土1.4	83.7土0.8	83.1土1.3	83.5土1.0	84.9土1.2	86.7±0.8	85.1土1.5	84.4土0.7	86.2土1.6	84.6土0.7	90.3 ±0.9	86.1土1.1
Mutagenicity	70.7±0.3	75.2±0.2	70.9±0.3	75.3土0.2	72.1土0.2	75.5±0.3	79.4土0.3	80.2±0.2	66.2±0.1	66.8土0.5	67.5土0.2	76.7土0.4
AIDS	0.0千9.66	99.6土0.1	0.0干9.66	0.0干9.66	99.7 ±0.0	99.7 ±0.0	0.0干9.66	0.0干9.66	99.3土0.1	99.7 ±0.0	99.7 ±0.0	99.7 ±0.0
BZR	88.1土0.8	88.2土0.8	88.1土0.6	87.8±0.7	84.5土1.0	85.5±0.7	86.5土0.6	87.8土0.6	82.8土0.9	84.1土0.8	87.1土0.5	87.7土1.0
COX2	81.2土1.0	81.1土0.5	80.5土0.8	80.5±0.7	80.7±0.5	80.3±0.7	80.4土1.1	81.4土0.6	78.1土0.1	77.3土0.7	81.7土0.8	81.5±0.9
DHFR	81.5土0.8	82.1土0.3	79.2±0.8	80.0土0.7	75.3土0.7	78.1土0.8	80.7土1.0	81.0土0.8	75.1土0.5	78.3土0.7	82.8土0.6	83.2±0.7
DD	78.3±0.3	78.2±0.3	77.3土0.4	77.3土0.4	78.7±0.3	79.2土0.9	79.4土0.4	71.0土0.2	78.8土0.6	78.2土0.6	78.9土0.3	81.6 ±0.5
PROTEINS	76.6土0.6	76.8 ±0.4	75.1土0.2	74.8土0.5	74.5土0.4	74.6土0.5	75.7±0.3	74.2土0.4			74.3土0.5	74.6土0.5
ENZYMES	64.1 ±1.3	63.5土1.1	61.6土1.2	62.0土1.2	40.2土1.0	49.3土1.1	42.3土1.0	58.9土1.1	37.6土0.7	39.5土1.3	49.0土1.6	62.6土0.9
IMDB-BINARY	69.4土0.6	69.9土0.5	70.6土0.5	70.1土0.4	70.7土0.6	71.1土0.9	60.5±0.3	70.2±0.7			73.5 ±0.3	71.2±0.7
IMDB-MULTI	46.1土0.7	47.0土0.5	47.1土0.6	47.6土0.4	47.8土0.6	47.8土0.6	40.8土1.1	46.1土0.7			49.8土0.6	51.0 ±0.7
REDDIT-BINARY					82.3土0.2	82.7土0.4	82.4土0.1	67.8土0.2			78.2土0.3	85.5±0.3
MSRC-9	90.9土1.0	90.4土0.5	90.4土1.2	90.4土0.8	90.4土1.4	90.1土1.0	91.8±0.8	88.2±1.2			89.4土1.3	89.7土0.9
MSRC-21	89.0十0.68	89.0十0.8	89.3土0.5	89.3土0.5	91.3 ±0.5	91.2土0.5	89.9±0.5	27.5土1.2			88.6土0.5	89.8土0.6
MSRC-21C	85.7 ±0.6	85.6土0.9	85.6土0.9	84.8土1.1	84.4土0.9	84.6土0.9	85.1土1.4	76.8土1.3			81.4土1.1	81.8土1.1
Tox21-AR	97.7±0.0	97.7±0.0	97.4±0.0	97.4土0.0	97.6土0.0	97.7±0.0	97.0±0.0	97.6土0.0			97.8土0.0	97.8土0.0
Tox21-MMP	86.9土0.1	87.2土0.1	86.8土0.1	87.2±0.1	86.6土0.2	89.7土0.1	86.3土0.1	90.4土0.1			84.7土0.1	89.6土0.2
Tox21-AHR	89.6土0.0	89.6土0.0	89.4土0.1	89.5土0.0	89.4土0.1	91.7土0.1	88.7土0.1	91.8土0.1			89.2±0.0	91.1土0.1
Average	77.4	78.2	76.9	77.3	76.1	77.6	76.3	74.1	69.5	71.6	77.6	78.1
^a Computation of the The highest accuracy	DBR kernel did nc	ot finish within 48h ed in bold for each	on some dataset: ה dataset.	s, as indicated by a	a line (—). The DBF	Rernel does not r	nake use of label i	information				

Table 4 Classification accuracy and standard deviation for several kernels and their variant when plugged into the Gaussian RBF kernel

(2020) 5:6 Kriege et al. Applied Network Science

value reached by WL-OA on the same dataset. However, for the dataset REDDIT-BINARY the accuracy of WL is improved, but still far below the accuracy obtained by WL-OA, which is based on the histogram intersection kernel applied to the WL feature vectors. A surprising result is that the trivial EL kernel combined with the Gaussian RBF kernel performs competitive to many sophisticated graph kernels. On an average it provides a higher accuracy than the (unmodified) SP, GL3 and PM kernel. The DBR kernel does not take labels into account and performs poorly on most datasets.

The application of the Gaussian RBF kernel introduces the hyper-parameter σ , which must be optimized, e.g., via grid search and cross-validation. This is computational demanding for large datasets, in particular, when the graph kernel also requires parameters that must be optimized. Therefore, we suggest to combine VL, EL and GL3 with a Gaussian RBF kernel as a base line. For WL and WL-OA the parameter h needs to be optimized and the accuracy gain is minor for most datasets, in particular for WL-OA. Therefore, their combination with an Gaussian RBF kernel cannot be generally recommended. Note that the combination with an Gaussian RBF kernel also complicates the application of fast linear classifiers, which are advisable for large datasets.

Q3 Accuracy. Tables 3 and 4 show that for almost every kernel there is at least one dataset, for which it provides the best accuracy. This is even true for the trivial kernels VL and EL on the datasets AIDS and MSRC-9; and also COX2 when combined with an Gaussian RBF kernel. Moreover, VL combined with the Gaussian RBF kernel almost reaches the accuracy of the best kernels for DD. The dataset AIDS is almost perfectly classified by VL, which suggests that this dataset is not an adequate benchmark dataset for graph kernel comparison. For the other two datasets (MSRC-9 and COX2), there are two possible reasons for the observed results. Either these datasets can be classified optimally without taking the graph structure into account, making them not adequate for graph kernel comparison. This would mean that the remaining error is dominated by irreducible error (label noise). Alternatively, current state-of-the-art kernels are not able to benefit from their structure; the remaining error is due to bias. If the second reason is true, these datasets are particularly challenging. In practice, for a finite dataset, it is hard to distinguish bias from noise conclusively, and it is likely that the full explanation is a combination of the two.

The kernels WL and WL-OA provide the best accuracy results for most datasets. WL-OA achieves the highest accuracy on an average even without combining it with the Gaussian RBF kernel. Since these kernels are also efficiently computed, they represent a suitable first approach when classifying new datasets. We suggest to use WL-OA for small and medium-sized datasets with kernel support vector machines and WL for large datasets with linear support vector machines.

The analysis of the label completeness ration depicted in Fig. 8 suggests that VL cannot perform well on ENZYMES, IMDB-BINARY, IMDB-MULTI and REDDIT-BINARY. EL shows weaknesses on IMDB-BINARY, IMDB-MULTI and REDDIT-BINARY and DBR on Mutagenicity. The WL and WL-OA kernels can effectively distinguish most non-isomorphic benchmark graphs. These observations are in accordance with the accuracy results observed. However, there is no clear relation between the label completeness ratio



and the prediction accuracy. This suggests that the ability of graph kernels to take features into account that allow to effectively distinguish graphs is only a minor issue for current benchmark datasets. Instead taking the features into account that allow the classifier to generalize to unseen data appears to be most relevant.

Q4 Agreement. The sheer number and variety of existing graph kernels suggest that there may be groups of kernels that are more similar to each other than to other kernels. In this section, we attempt to discover such groups by a qualitative comparison of the predictions (and errors) made by different kernels for a fixed set of graphs. Additionally, we examine the heterogeneity in errors made for the same set of graphs to assess the overall agreement between rivalling kernels.

We embed each kernel into a common geometric space based on their predictions on a set of benchmark graphs. Let each kernel $k_1, ..., k_m$ and each graph $G_1, ..., G_n$ in a dataset D index the rows and columns of a matrix $\mathbf{P}^D \in \mathbb{R}^{m \times n}$, respectively. Then, let P_{ij}^D represent the prediction made by kernel k_i on graph G_j after being trained on other graphs from D. We construct such matrices \mathbf{P}^l for multiple datasets $\{D_l\}_{l=1}^N$ and concatenate them to form $\mathbf{P} = [\mathbf{P}^1, ..., \mathbf{P}^N]$, a high-dimensional representation of the features captured by each kernel. Similarly, we construct matrices $\{E^l\}_{l=1}^N$ and $E = [\mathbf{P}^1, ..., \mathbf{P}^N]$, representing the prediction errors made by different kernels on different graphs. Specifically, we let $E_{ij}^l = \mathbbm{1}[P_{ij} \neq y_l(G_j)]$, where $y_l(G_j)$ is the class label of G_j . Here, we construct \mathbf{P} and \mathbf{E} from the predictions made by a large set of kernels and parameter settings (see Fig. 9 for a list) applied to the datasets MUTAG, ENZYMES and PTC-MR.

In Fig. 9, we illustrate the predictions of different kernels by projecting the rows of the prediction matrix P to \mathbb{R}^2 using t-SNE (Maaten Lvd and Hinton 2008). The position of each dot represents a projection of the predictions made by a single kernel. The color represents the kernel family and the size represents the average accuracy of the kernel in the considered datasets. For comparison, we include two additional variants of the RW kernel: one comparing only walks of a fixed length l (FL-RW), and one defined as the sum of such kernels up to a fixed length l (MFL-RW). We see that WL optimal assignment (WL-OA) and matching kernels (MK) predict similarly, compared to for example short-length RW kernels. However, despite small random walks and WL-OA with h = 0 representing very local features, they predict qualitatively different. We also see that RW kernels that sum up kernels of length l < L walks are very similar to kernels based on just length L walks and that EL, GL3 and short-length RW kernels predict similarly, as expected from their local scope.

Similarity between two rows $e_i = E_i$, $e_{i'} = E_{i'}$. of the error matrix *E* indicate that kernels k_i and $k_{i'}$ make similar predictive errors on the considered datasets. To assess the overall extent to which particular graphs are "easy" or "hard" for many kernels, we studied the variance of the columns of *E*. We find that the average zero-one loss across kernels on MUTAG (0.14), ENZYMES (0.57) and PTC-MR (0.42) correlates strongly with the mean absolute deviation around the median across kernels (0.07, 0.26, 0.23). The latter may be interpreted as the fraction of instances for which kernels disagree with the majority vote. We also evaluated the average *inter-agreement* between kernels as measured using *Fleiss' kappa* (Fleiss 1971). A high value of Fleiss' kappa indicates that different raters agree significantly more often than random raters with the same marginal label probability. On MUTAG, ENZYMES and PTC-MR, the kappa measure shows a trend similar (but inverse) to the standard deviation with values of (0.60, 0.28, 0.36).

We conclude that, on these examples, the more difficult the classification task, the more varied the predictive errors. Indeed, if the average error across kernels was 0.0, all models would agree everywhere. However, if different kernels had similar biases, the reverse would not necessarily be true. Instead, these results confirm our intuition that different kernels encode different biases and may be appropriate for different datasets as a result.

Q5 Continuous attributes. As can be seen in Table 5, on all datasets, excluding the FRANKENSTEIN dataset, one variant of the hash graph kernel framework achieves state-of-the-art results. This is in line with the theoretical results outlined in (Morris et al. 2016), i.e., they show how to approximate well-known graph kernels for

Table 5 Classification accuracies in percent and standard deviations (Number of iterations for HGK-WL and HGK-SP: 20 (100 for SYNTHIE), OOM— Out of Memory

Kernel	Dataset					
Kenner	ENZYMES	FRANKENSTEIN	PROTEINS	SyntheticNew	Synthie	rweitage
SP+RBF	71.0±0.8	72.8±0.2	76.6±0.5	96.2±0.4	52.8±1.8	73.9
HGK-SP	71.3 ±0.9	70.1±0.3	77.5 ±0.4	96.5±0.6	94.3±0.5	81.9
HGK-WL	67.6±1.0	73.6±0.4	76.7±0.4	98.8 ±0.3	96.8 ±0.5	82.7
GH	68.8±1.0	68.5±0.3	72.3±0.3	85.1±1.0	73.2 ± 0.8	73.6
GI	71.7 ±0.8	76.3 ±0.3	76.9±0.5	83.1±1.1	95.8±0.5	80.8
P2K	69.2±0.4	OOM	73.5±0.5	91.7±0.9	50.2 ± 1.9	71.2

The highest accuracy value is highlighted in bold for each dataset

graphs with vertex attributes up to some arbitrarily small error (depending on the number of iterations). However, the results are already achieved with a small number of iterations. This is likely a property of the employed datasets, i.e., a coarse-grained comparison of the attributes is sufficient. Moreover, together with the propagation kernel, the instances of the hash graph kernel framework achieve a much lower running time compared to the other implicit approaches. The lower performance of the hash graph kernel instances on the FRANKENSTEIN dataset is likely due to the high-dimensional vertex attributes, which are hard to compare using hash functions.

A practitioner's guide

Because of the limited theoretical knowledge we have about the expressivity of different kernels and the challenge of assessing this a priori, it is difficult to predict which kernel will perform well for a given problem. Nevertheless, it is often the case that some of the kernels in the literature are less or more well suited to the problem at hand. For example, kernels with high time complexity w.r.t. vertex count are expensive to compute for very large graphs; kernels that do not support vertex attributes are ill-suited in learning problems where these are highly significant.

Below, we give and motivate general guidelines for prioritizing and deprioritizing kernels based on four properties of the problem at hand: the importance and nature of vertex attributes, the size and density of graphs, the importance of global structure, and the number of graphs in the available dataset. Examples of appropriate and unappropriate kernels are given for extreme cases of each property, and the resulting guidelines are illustrated in Fig. 10. The chosen set of properties is certainly a subset of those that may be predictive of a kernel's performance in a given task. For example, the density and number of vertices of a graph are very crude measures of the graph's structure. On the other hand, these features are generally applicable and easy to compute for any sets of graphs. In some fixed domain, more specific structural properties such as girth or diameter may be important



Fig. 10 Guidelines for prioritizing kernels for consideration based on known properties of the graph learning problem. In the "A practitioner's guide" section, we justify these recommendations based on the graph kernel literature

and could guide the choice of kernel further. In this work, however, we limit ourselves to the more general case.

Vertex attributes Almost all established benchmarks for graph classification contain vertex labels and almost all graph kernels support the use of them in some way. In fact, any kernel can be made sensitive to vertex and edge attribute through multiplication by a label kernel, although this approach will not take into account the dependencies between labels and structure. Hence, one of the great contributions of the Weisfeiler-Lehman (Shervashidze et al. 2011) and related kernels (e.g. Propagation kernels (Neumann et al. 2016)) is that they capture such dependencies in transformed graphs that are beneficial to all kernels that support vertex labels. It has therefore become standard practice to perform a WL-like transform on labeled graphs before application of other kernels. For this reason, we consider WL-kernels a first choice for applications where vertex labels are important. Propagation kernels also naturally couple structure and attributes, but are generally more expensive to compute. The assignment step of OA kernels matches vertices based on both structure and attribute, depending on implementation. In contrast, the original Lovász, SVM-theta and graphlet kernels have no standard way of incorporating vertex labels. The graphlet kernel may be modified to do so by considering subgraph patterns as different if they have different labels. An important special-case of attributed graphs is graphs with non-discrete vertex attributes; these require special consideration. The GraphHopper, GraphInvariant and Hash Graph kernels as well as neural network-based approaches excel at making use of such attributes. In contrast, subtree kernels and shortest-path kernels become prohibitively expensive to compute when combined with continuous attributes.

Large graphs Early graph kernels such as the RW and SP kernels were plagued by worst-case running time complexities that were prohibitively high for large graphs: $\mathcal{O}(n^6)$ and $\mathcal{O}(n^4)$ for pairs of graphs with *n* the largest number of vertices. Also expensive to compute, the subgraph matching kernel has complexity $\mathcal{O}(kn^{2(k+1)})$ where k is the size of the considered subgraphs. In practice, even a complexity quadratic in the number of vertices is too high for large-scale learning-the goal is often to achieve complexity linear in the largest number of edges, m. This goal puts fundamental limitations on expressivity, as linear complexity is unachievable if the attributes of each edge of one graph has to be independently compared to those of each edge in another. However, when speed is of utmost importance, we recommend using efficient alternatives such as fast subtree kernels with complexity $\mathcal{O}(hm)$ where h the depth of the deepest subtree. Additionally, a single WL iteration may be computed in $\mathcal{O}(m)$ time and the WL label propagation may be used as-is with an already fast kernel at a constant multiplicative cost h, equal to the number of WL iterations. As a result, to improve a kernel's sensitivity to vertex label structure is often relatively cheap. Finally, for settings when a particular kernel is preferred for its expressivity but not for its running time, authors have proposed approximation schemes that reduce running time based on sampling or approximate optimization. For example, the time to compute the k-graphlet spectrum for a graph, with worst-case complexity $O(nd^{k-1})$ and d the maximum degree, may be significantly reduced for dense graphs by sampling subgraphs to produce an unbiased estimate of the kernel; The Lovász kernel, with complexity $\mathcal{O}(n^6)$, was approximated with the SVM-theta kernel with $\mathcal{O}(n^2)$; The random walk kernel may be approximated by the *p*-random walk kernel where walks are limited to length *p*. Similar approximations may be derived also for other kernels. For very large graphs, simple alternatives like the edge label and vertex label kernels may be useful baselines but neglect the graph structure completely.

Global structure Global properties of graphs are properties that are not well described by statistics of (small) subgraphs (Johansson et al. 2014). It has been shown, for example, that there exist graphs for which all small subgraphs are trees, but the overall graph has high girth and high chromatic number (Alon and Spencer 2004). Although the graph kernel literature has often left the precise interpretation of "global" to the reader, kernels such as the Lovász kernels and the Glocalized WL kernel, have been proposed with guarantees of capturing specific properties that are considered global by the authors (see in the "Other approaches" section). Beside these kernels, if domain knowledge suggests that global structure is important to the task at hand, we recommend prioritizing kernels that compute features from larger subgraph patterns, walks or paths. This rules out the use of Graphlet kernels, since counting large graphlets is often prohibitively expensive, and (small) neighborhood aggregation methods such as the Weisfeiler-Lehman kernel for small numbers of iterations. On the other hand, the shortest-path kernel, long-walk RW and high-iteration WL kernels compute features based on patterns spanning large portions of graphs.

Large datasets A drawback of kernel methods in general is that they require computation and storage of the full $N \times N$ kernel matrix for each pair of instances in a dataset of N graphs. This can be alleviated significantly if the chosen kernel admits an explicit *d*-dimensional representation with $d \ll N$, such as the vertex label, Weisfeiler-Lehman and graphlet kernels. In this case, only the $N \times d$ feature matrix is necessary for learning. Thus, if many graphs are available to learn from, we recommend starting with kernels that admit an explicit feature representation, such as the WL, GL and subtree kernels. However, this is not always possible, such as when continuous vertex attributes are important, and vertices are compared with a distance metric. Instead, computations using implicit kernels may be approximated using the prototypes method described in the "Assignment- and matching-based approaches" section in which a subset of d graphs are selected and compared to each instance in the dataset. Under certain conditions on the prototype selection, this gives an unbiased estimator of the kernel matrix which can be used in place of its exact version. Finally, in most cases, more efficient learning methods are applicable when explicit feature representations are available. For classification with support vector machines, for example, the software package LIBSVM (Chang and Lin 2011) is commonly used for learning with (implicit) kernels. When explicit feature representations are available, the software LIBLINEAR (Fan et al. 2008), which scales to very large datasets, can be used as an alternative.

Conclusion

We gave an overview over the graph kernel literature. We hope that this survey will spark further progress in the area of graph kernel design and graph classification in general. Moreover, we hope that this article is valuable for the practitioner applying graph classification methods to solve real-world problems.

Acknowledgements

We thank Pinar Yanardag, Lu Bai, Giannis Nikolentzos, Marion Neumann, and Franceso Orsini for providing their graph kernel source code.

Authors' contributions

NMK implemented several of the graph kernels used in the "Experimental study" section and was the main responsible for the experimental evaluation regarding Q1, Q2 and Q3. FDJ was the main responsible for the implementation and/or adaptation of the MK, PM, DeepGK, Prop and DBR kernels for use in the "Experimental study" section and for experimental evaluation with regards to Q4. CM was the main responsible for the experimental evaluation regarding Q5. All authors contributed to the writing of the manuscript, read and approved the final version.

Funding

NMK and CM have been supported by the German Research Foundation (DFG) within the Collaborative Research Center SFB 876 "Providing Information by Resource-Constrained Data Analysis", project A6 "Resource-efficient Graph Mining".

Availability of data and materials

The datasets analyzed during the current study are available in the "Benchmark Data Sets for Graph Kernels" repository, http://graphkernels.cs.tu-dortmund.de (Kersting et al. 2016).

Competing interests

The authors declare that they have no competing interests.

Author details

¹Department of Computer Science, TU Dortmund University, Otto-Hahn-Str. 14, 44227 Dortmund, Germany. ²Department of Computer Science and Engineering, Chalmers University of Technology, Rännvägen 6, 41258 Göteborg, Sweden.

Received: 21 March 2019 Accepted: 23 August 2019 Published online: 14 January 2020

References

- Adamson GW, Bush JA (1973) A method for the automatic classification of chemical structures. Inf Storage Retrieval 9(10):561–568. http://dx.doi.org/10.1016/0020-0271(73)90059-4
- Ahmed NK, Willke T, Rossi RA (2016) Estimation of local subgraph counts. In: IEEE International Conference on Big Data. pp 1–10. https://doi.org/10.1109/bigdata.2016.7840651
- Aiolli F, Donini M, Navarin N, Sperduti A (2015) Multiple graph-kernel learning. In: IEEE Symposium Series on Computational Intelligence. pp 1607–1614. https://doi.org/10.1109/ssci.2015.226
- Alon N, Spencer JH (2004) The probabilistic method. Wiley. https://doi.org/10.1002/0471722154.ch1
- Babai L, Kucera L (1979) Canonical labelling of graphs in linear average time. In: Annual Symposium on Foundations of Computer Science. pp 39–46. https://doi.org/10.1109/sfcs.1979.8
- Bai L, Ren P, Bai X, Hancock ER (2014) A graph kernel from the depth-based representation. In: Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition and Structural and Syntactic Pattern Recognition. pp 1–11. https://doi.org/10.1007/978-3-662-44415-3_1
- Bai L, Rossi L, Zhang Z, Hancock ER (2015) An aligned subtree kernel for weighted graphs. In: International Conference on Machine Learning. pp 30–39. https://doi.org/10.1109/icpr.2016.7899666
- Balcan MF, Blum A, Srebro N (2008) A theory of learning with similarity functions. Mach Learn 72(1-2):89–112 Borgwardt KM (2007) Graph kernels. Phd thesis, Ludwig Maximilians University Munich
- Borgwardt KM, Kriegel HP (2005) Shortest-path kernels on graphs. In: IEEE International Conference on Data Mining. pp 74–81. https://doi.org/10.1109/icdm.2005.132
- Borgwardt KM, Ong CS, Schönauer S, Vishwanathan SVN, Smola AJ, Kriegel HP (2005) Protein function prediction via graph kernels. Bioinformatics 21(Supplement 1):i47–i56
- Borgwardt KM, Kriegel HP, Vishwanathan S, Schraudolphs NN (2007) Graph kernels for disease outcome prediction from protein-protein interaction networks. In: Biocomputing 2007, World Scientific. pp 4–15. https://doi.org/10.1142/9789812772435_0002
- Bressan M, Chierichetti F, Kumar R, Leucci S, Panconesi A (2017) Counting graphlets: Space vs time. In: ACM International Conference on Web Search and Data Mining. pp 557–566. https://doi.org/10.1145/3018661.3018732
- Brown N (2009) Chemoinformatics an introduction for computer scientists. ACM Comput Surv 41(2). https://doi.org/10. 1145/1459352.1459353
- Ceroni A, Costa F, Frasconi P (2007) Classification of small molecules by two- and three-dimensional decomposition kernels. Bioinformatics 23(16):2038–2045. http://dx.doi.org/10.1093/bioinformatics/btm298

Chang CC, Lin CJ (2011) LIBSVM: A library for support vector machines. ACM Trans Intell Syst Technol 2(3):27:1–27:27 Chen X, Lui JCS (2016) Mining graphlet counts in online social networks. In: IEEE International Conference on Data Mining. pp 71–80. https://doi.org/10.1109/icdm.2016.0018

Cortes C, Vapnik V (1995) Support-vector networks. Mach Learn 20(3):273-297

Costa F, De Grave K (2010) Fast Neighborhood Subgraph Pairwise Distance Kernel. In: Fürnkranz J, Joachims T (eds). Proceedings of the 27th International Conference on Machine Learning (ICML-10). Omnipress, Haifa. pp 255–262. http://www.icml2010.org/papers/347.pdf

- Da San Martino G, Navarin N, Sperduti A (2012a) A memory efficient graph kernel. In: International Joint Conference on Neural Networks. pp 1–7. https://doi.org/10.1109/ijcnn.2012.6252831
- Da San Martino, G, Navarin N, Sperduti A (2012b) A tree-based kernel for graphs. In: SIAM Conference of Data Mining. pp 975–986. https://doi.org/10.1137/1.9781611972825.84

Daylight CIS (2008) Daylight theory manual v4.9. http://www.daylight.com/dayhtml/doc/theory

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

de Vries GKD (2013) A fast approximation of the Weisfeiler-Lehman graph kernel for rdf data. In: European Conference on Machine Learning & Principles and Practice of Knowledge Discovery in Databases. pp 606–621. https://doi.org/10. 1007/978-3-642-40988-2_39

Dobson PD, Doig AJ (2003) Distinguishing enzyme structures from non-enzymes without alignments. J Mol Biol 330(4):771–783

Durant JL, Leland BA, Henry DR, Nourse JG (2002) Reoptimization of mdl keys for use in drug discovery. J Chem Inf Comput Sci 42(5):1273–1280

Duvenaud DK, Maclaurin D, Iparraguirre J, Bombarell R, Hirzel T, Aspuru-Guzik A, Adams RP (2015) Convolutional networks on graphs for learning molecular fingerprints. In: Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada. pp 2224–2232

Dwork C, Roth A, et al. (2014) The algorithmic foundations of differential privacy. Found Trends® Theor Comput Sci 9(3–4):211–407

Fan RE, Chang KW, Hsieh CJ, Wang XR, Lin CJ (2008) Liblinear: A library for large linear classification. J Mach Learn Res 9:1871–1874

Feragen A, Kasenburg N, Petersen J, Bruijne MD, M BK (2013) Scalable kernels for graphs with continuous attributes. In: Advances in Neural Information Processing Systems. pp 216–224. erratum available at http://image.diku.dk/aasa/ papers/graphkernels_nips_erratum.pdf

Fey M, Lenssen JE, Weichert F, Müller H (2018) SplineCNN: Fast geometric deep learning with continuous b-spline kernels. In: IEEE Conference on Computer Vision and Pattern Recognition. pp 869–877. https://doi.org/10.1109/cvpr. 2018.00097

Fleiss JL (1971) Measuring nominal scale agreement among many raters. Psychol Bull 76(5):378

Fröhlich H, Wegner JK, Sieker F, Zell A (2005) Optimal assignment kernels for attributed molecular graphs. In: International Conference on Machine learning. pp 225–232. https://doi.org/10.1145/1102351.1102380

Gärtner T, Flach P, Wrobel S (2003) On graph kernels: Hardness results and efficient alternatives. In: Learning Theory and Kernel Machines. Springer. pp 129–143. https://doi.org/10.1007/978-3-540-45167-9_11

Ghosh S, Das N, Gonçalves T, Quaresma P, Kundu M (2018) The journey of graph kernels through two decades. Comput Sci Rev 27:88–111

Gilmer J, Schoenholz SS, Riley PF, Vinyals O, Dahl GE (2017) Neural Message Passing for Quantum Chemistry. In: Precup D, Whye Teh Y (eds). Proceedings of the 34th International Conference on Machine Learning. PMLR, Sydney. http:// proceedings.mlr.press/v70/gilmer17a.html

Grauman K, Darrell T (2007a) Approximate correspondences in high dimensions. In: Advances in Neural Information Processing Systems. pp 505–512. https://doi.org/10.7551/mitpress/7503.003.0068

Grauman K, Darrell T (2007b) The pyramid match kernel: Efficient learning with sets of features. J Mach Learn Res 8(Apr):725–760

Hamilton WL, Ying R, Leskovec J (2017) Inductive representation learning on large graphs. CoRR abs/1706.02216:1025–1035. http://arxiv.org/abs/1706.02216

Harchaoui Z, Bach F (2007) Image classification with segmentation graph kernels. In: IEEE Conference on Computer Vision and Pattern Recognition. pp 1–8. https://doi.org/10.1109/cvpr.2007.383049

Haussler D (1999) Convolution kernels on discrete structures. Tech. Rep. UCS-CRL-99-10, University of California at Santa Cruz

Helma C, King RD, Kramer S, Srinivasan A (2001) The predictive toxicology challenge 2000–2001. Bioinformatics 17(1):107–108

Hermansson L, Kerola T, Johansson F, Jethava V, Dubhashi D (2013) Entity disambiguation in anonymized graphs using graph kernels. In: ACM International Conference on Information & Knowledge Management. pp 1037–1046. https://doi.org/10.1145/2505515.2505565

Hermansson L, Johansson FD, Watanabe O (2015) Generalized shortest path kernel on graphs. In: Discovery Science: International Conference. pp 78–85. https://doi.org/10.1007/978-3-319-24282-8_8

Hido S, Kashima H (2009) A linear-time graph kernel. In: IEEE International Conference on Data Mining. pp 179–188. https://doi.org/10.1109/icdm.2009.30

Horváth T, Gärtner T, Wrobel S (2004) Cyclic pattern kernels for predictive graph mining. In: ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp 158–167. https://doi.org/10.1145/1014052.1014072

Horváth T, Ramon J, Wrobel S (2010) Frequent subgraph mining in outerplanar graphs. Data Min Knowl Discov 21:472–508. https://doi.org/10.1007/s10618-009-0162-1

Jie B, Liu M, Jiang X, Zhang D (2016) Sub-network based kernels for brain network classification. In: ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics. pp 622–629. https://doi.org/10.1145/2975167.2985687

Johansson FD, Dubhashi D (2015) Learning with similarity functions on graphs using matchings of geometric embeddings. In: ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. pp 467–476. https://doi.org/10.1145/2783258.2783341

Johansson FD, Jethava V, Dubhashi DP, Bhattacharyya C (2014) Global graph kernels using geometric embeddings. In: International Conference on Machine Learning. pp 694–702

Johansson FD, Frost O, Retzner C, Dubhashi D (2015) Classifying large graphs with differential privacy. In: Modeling Decisions for Artificial Intelligence. Springer. pp 3–17. https://doi.org/10.1007/98-3-319-23240-9_1

Johnson DS (2005) The NP-completeness column. ACM Trans Algorithms 1(1):160–176. https://doi.org/10.1145/1077464. 1077476

Kang U, Tong H, Sun J (2012) Fast random walk graph kernel. In: SIAM International Conference on Data Mining. pp 828–838. https://doi.org/10.1137/1.9781611972825.71

Kashima H, Tsuda K, Inokuchi A (2003) Marginalized kernels between labeled graphs. In: International Conference on Machine Learning. pp 321–328

Kazius J, McGuire R, Bursi R (2005) Derivation and validation of toxicophores for mutagenicity prediction. J Med Chem 48(13):312–320

Kersting K, Kriege NM, Morris C, Mutzel P, Neumann M (2016) Benchmark data sets for graph kernels. http://graphkernels. cs.tu-dortmund.de

- Kipf TN, Welling M (2017) Semi-supervised classification with graph convolutional networks. In: International Conference on Learning Representations
- Kondor R, Pan H (2016) The multiscale laplacian graph kernel. In: Advances in Neural Information Processing Systems. pp 2982–2990
- Kondor R, Shervashidze N, Borgwardt KM (2009) The graphlet spectrum. In: International Conference on Machine Learning. pp 529–536. https://doi.org/10.1145/1553374.1553443
- Kriege N, Mutzel P (2012) Subgraph matching kernels for attributed graphs. In: International Conference on Machine Learning
- Kriege N, Neumann M, Kersting K, Mutzel M (2014) Explicit versus implicit graph feature maps: A computational phase transition for walk kernels. In: IEEE International Conference on Data Mining. pp 881–886. https://doi.org/10.1109/ icdm.2014.129

Kriege NM (2015) Comparing graphs: Algorithms & applications. Phd thesis, TU Dortmund University

Kriege NM (2019) Deep Weisfeiler-Lehman assignment kernels via multiple kernel learning. In: 27th European Symposium on Artificial Neural Networks, ESANN 2019

- Kriege NM, Giscard PL, Wilson RC (2016) On valid optimal assignment kernels and applications to graph classification. In: Advances in Neural Information Processing Systems. pp 1615–1623
- Kriege NM, Neumann M, Morris C, Kersting K, Mutzel P (2019) A unifying view of explicit and implicit feature maps of graph kernels. Data Mining and Knowledge Discovery 33(6):1505–1547. https://doi.org/10.1007/s10618-019-00652-0 Kriege NM, Morris C, Rey A, Sohler C (2018) A property testing framework for the theoretical expressivity of graph kernels.
- In: International Joint Conference on Artificial Intelligence. pp 2348–2354. https://doi.org/10.24963/ijcai.2018/325 Li B, Zhu X, Chi L, Zhang C (2012) Nested subtree hash kernels for large-scale graph classification over streams. In: IEEE
- International Conference on Data Mining. pp 399–408. https://doi.org/10.1109/icdm.2012.101 Li L, Tong H, Xiao Y, Fan W (2015) *Cheetah*: Fast graph kernel tracking on dynamic graphs. In: SIAM International
- Conference on Data Mining. pp 280–288. https://doi.org/10.1137/1.9781611974010.32 Li W, Saidi H, Sanchez H, Schäf M, Schweitzer P (2016) Detecting similar programs via the Weisfeiler-Leman graph kernel.
- In: International Conference on Software Reuse. pp 315–330. https://doi.org/10.1007/978-3-319-35122-3_21
- Loosli G, Canu S, Ong CS (2015) Learning svm in krein spaces. IEEE Trans Pattern Anal Mach Intell PP(99):1–1. https://doi. org/10.1109/TPAMI.2015.2477830
- Lovász L (2006) On the shannon capacity of a graph. IEEE Trans Inf Theory 25(1):1–7
- Maaten Lvd, Hinton G (2008) Visualizing data using t-sne. J Mach Learn Res 9(Nov):2579–2605
- Mahé P, Vert JP (2009) Graph kernels based on tree patterns for molecules. Mach Learn 75(1):3–35
- Mahé P, Ueda N, Akutsu T, Perret JL, Vert JP (2004) Extensions of marginalized graph kernels. In: International Conference on Machine Learning. pp 552–559. https://doi.org/10.1145/1015330.1015446
- Mahé P, Ueda N, Akutsu T, Perret JL, Vert JP (2005) Graph kernels for molecular structure-activity relationship analysis with support vector machines. J Chem Inf Model 45(4):939–951
- Mahé P, Ralaivola L, Stoven V, Vert JP (2006) The pharmacophore kernel for virtual screening with support vector machines. J Chem Inf Model 46(5):2003–2014

Massimo CM, Navarin N, Sperduti A (2016) Hyper-parameter tuning for graph kernels via multiple kernel learning. In: Advances in Neural Information Processing. pp 214–223. https://doi.org/10.1007/978-3-319-46672-9_25

- McKay BD, Piperno A (2014) Practical graph isomorphism, II. J Symb Comput 60(0):94–112. http://dx.doi.org/10.1016/j.jsc. 2013.09.003
- Merkwirth C, Lengauer T (2005) Automatic generation of complementary descriptors with molecular graph networks. J Chem Inf Model 45(5):1159–1168
- Mikolov T, Chen K, Corrado G, Dean J (2013) Efficient estimation of word representations in vector space. In: 1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings. https://dblp.org/rec/bib/journals/corr/abs-1301-3781

Mohri M, Rostamizadeh A, Talwalkar A (2012) Foundations of Machine Learning. MIT Press

- Morris C, Kriege NM, Kersting K, Mutzel P (2016) Faster kernel for graphs with continuous attributes via hashing. In: IEEE International Conference on Data Mining. pp 1095–1100. https://doi.org/10.1109/icdm.2016.0142
- Morris C, Kersting K, Mutzel P (2017) Glocalized Weisfeiler-Lehman kernel: Global-local feature maps of graphs. In: IEEE International Conference on Data Mining
- Morris C, Ritzert M, Fey M, Hamilton WL, Lenssen JE, Rattan G, Grohe M (2019) Weisfeiler and Leman go neural: Higher-order graph neural networks. In: AAAI Conference on Artificial Intelligence. p TBD. https://doi.org/10.1609/ aaai.v33i01.33014602

Neumann M (2015) Learning with graphs using kernels from propagated information. Phd thesis, University of Bonn Neumann M (2016) Propagation kernel (code). https://github.com/marionmari/propagation_kernels.git

- Neumann M, Moreno P, Antanas L, Garnett R, Kersting K (2013) Graph kernels for object category prediction in task–dependent robot grasping. In: Adamic L, Getoor L, Huang B, Leskovec J, McAuley J (eds). Working Notes of the International Workshop on Mining and Learning with Graphs at KDD 2013, Chicago
- Neumann M, Garnett R, Bauckhage C, Kersting K (2016) Propagation kernels: Efficient graph kernels from propagated information. Mach Learn 102(2):209–245

Nikolentzos G (2016) Pyramid match kernel. http://www.db-net.aueb.gr/nikolentzos/code/matchingnodes.zip

- Nikolentzos G, Vazirgiannis M (2018) Enhancing graph kernels via successive embeddings. In: ACM International Conference on Information and Knowledge Management. pp 1583–1586. https://doi.org/10.1145/3269206.3269289
- Nikolentzos G, Meladianos P, Rousseau F, Stavrakas Y, Vazirgiannis M (2017a) Shortest-path graph kernels for document similarity. In: Empirical Methods in Natural Language Processing. pp 1890–1900. https://doi.org/10.18653/v1/d17-1202

Nikolentzos G, Meladianos P, Vazirgiannis M (2017b) Matching node embeddings for graph similarity. In: AAAI Conference on Artificial Intelligence. pp 2429–2435

Nikolentzos G, Meladianos P, Limnios S, Vazirgiannis M (2018) A degeneracy framework for graph similarity. In:

International Joint Conference on Artificial Intelligenc, ijcai.org. pp 2595–2601. https://doi.org/10.24963/ijcai.2018/360 Oneto L, Navarin N, Donini M, Sperduti A, Aiolli F, Anguita D (2017) Measuring the expressivity of graph kernels through statistical learning theory. Neurocomputing 268(Supplement C):4–16

Orsini F, Frasconi P, De Raedt L (2015) Graph invariant kernels. In: International Joint Conference on Artificial Intelligence. pp 3756–3762

Pachauri D, Kondor R, Singh V (2013) Solving the multi-way matching problem by permutation synchronization. In: Advances in Neural Information Processing Systems. pp 1860–1868

Ralaivola L, Swamidass SJ, Saigo H, Baldi P (2005) Graph kernels for chemical informatics. Neural Netw 18(8):1093–1110. https://doi.org/10.1016/j.neunet.2005.07.009. Neural Networks and Kernel Methods for Structured Domains

Ramon J, Bruynooghe M (2001) A polynomial time computable metric between point sets. Acta Inform 37(10):765–780. https://doi.org/10.1007/PL00013304

Ramon J, Gärtner T (2003) Expressivity versus efficiency of graph kernels. In: International Workshop on Mining Graphs, Trees and Sequences. pp 65–74

Rasmussen CE (2004) Gaussian processes in machine learning. In: Advanced lectures on machine learning. Springer. pp 63–71

Riesen K, Bunke H (2008) lam graph database repository for graph based pattern recognition and machine learning. In: Structural, Syntactic, and Statistical Pattern Recognition: Joint IAPR International Workshop. pp 287–297. https://doi. org/10.1007/978-3-540-89689-0_33

Rogers D, Hahn M (2010) Extended-connectivity fingerprints. J Chem Inf Model 50(5):742–754. http://dx.doi.org/10.1021/ ci100050t

Schiavinato M, Gasparetto A, Torsello A (2015) Transitive assignment kernels for structural classification. In: Similarity-Based Pattern Recognition: Third International Workshop. pp 146–159. https://doi.org/10.1007/978-3-319-24261-3_12

Schölkopf B, Smola AJ (2001) Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond. MIT Press, Cambridge

Schölkopf B, Smola A, Müller KR (1997) Kernel principal component analysis. In: International Conference on Artificial Neural Networks. Springer. pp 583–588

Schomburg I, Chang A, Ebeling C, Gremse M, Heldt C, Huhn G, Schomburg D (2004) Brenda, the enzyme database: updates and major new developments. Nucleic Acids Res 32:431–433. https://doi.org/10.1093/nar/gkh081

Shawe-Taylor J, Cristianini N (2004) Kernel Methods for Pattern Analysis. Cambridge University Press, New York Shervashidze N (2012) Scalable graph kernels. Phd thesis

Shervashidze N, Vishwanathan SVN, Petri TH, Mehlhorn K, Borgwardt KM (2009) Efficient graphlet kernels for large graph comparison. In: International Conference on Artificial Intelligence and Statistics. pp 488–495

Shervashidze N, Schweitzer P, van Leeuwen EJ, Mehlhorn K, Borgwardt KM (2011) Weisfeiler-Lehman graph kernels. J Mach Learn Res 12:2539–2561

Shin K, Kuboyama T (2008) A generalization of haussler's convolution kernel: mapping kernel. In: International conference on Machine learning. ACM. pp 944–951. https://doi.org/10.1145/1390156.1390275

Silverman BW (1986) Density Estimation for Statistics and Data Analysis. Chapman & Hall/CRC, London

Su Y, Han F, Harang RE, Yan X (2016) A fast kernel for attributed graphs. In: SIAM International Conference on Data Mining. pp 486–494. https://doi.org/10.1137/1.9781611974348.55

Sugiyama M, Borgwardt KM (2015) Halting in random walk kernels. In: Advances in Neural Information Processing Systems. 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. https://doi.org/10.1021/ci034143r

Swamidass SJ, Chen J, Bruand J, Phung P, Ralaivola L, Baldi P (2005) Kernels for small molecules and the prediction of mutagenicity, toxicity and anti-cancer activity. Bioinformatics 21(Suppl 1):i359–i368

Takerkart S, Auzias G, Thirion B, Ralaivola L (2014) Graph-based inter-subject pattern analysis of fmri data. PLoS ONE 9(8):1–14. https://doi.org/10.1371/journal.pone.0104586

Tox21 Data Challenge (2014). https://tripod.nih.gov/tox21/challenge/data.jsp

Vega-Pons S, Avesani P (2013) Brain decoding via graph kernels. In: Proceedings of the 2013 International Workshop on Pattern Recognition in Neuroimaging, IEEE Computer Society, Washington, DC, USA, PRNI '13. pp 136–139. https:// doi.org/10.1109/PRNI.2013.43

 Vega-Pons S, Avesani P, Andric M, Hasson U (2014) Classification of inter-subject fmri data based on graph kernels. In: International Workshop on Pattern Recognition in Neuroimaging. pp 1–4. https://doi.org/10.1109/PRNI.2014.6858549
 Vert J (2008) The optimal assignment kernel is not positive definite. CoRR:abs/0801.4061. http://arxiv.org/abs/0801.4061

Vishwanathan SVN, Schraudolph NN, Kondor R, Borgwardt KM (2010) Graph kernels. J Mach Learn Res 11:1201–1242 Wale N, Watson IA, Karypis G (2008) Comparison of descriptor spaces for chemical compound retrieval and classification. Knowl Inf Syst 14(3):347–375

Wang J, Wilson RC, Hancock ER (2016) fmri activation network analysis using bose-einstein entropy. In: Robles-Kelly A, Loog M, Biggio B, Escolano F, Wilson R (eds) Structural, Syntactic, and Statistical Pattern Recognition. Springer International Publishing, Cham. pp 218–228. https://doi.org/10.1007/978-3-319-49055-7_20

Willett P, Winterman V (1986) A comparison of some measures for the determination of inter-molecular structural similarity measures of inter-molecular structural similarity. Quant Struct-Act Relationsh 5(1):18–25. https://doi.org/10. 1002/qsar.19860050105

Woźnica A, Kalousis A, Hilario M (2010) Adaptive matching based kernels for labelled graphs. In: Advances in Knowledge Discovery and Data Mining, Lecture Notes in Computer Science, vol 6119. pp 374–385. https://doi.org/10.1007/978-3-642-13672-6_37

Wu B, Yuan C, Hu W (2014) Human action recognition based on context-dependent graph kernels. In: IEEE Conference on Computer Vision and Pattern Recognition. pp 2609–2616. https://doi.org/10.1109/CVPR.2014.334 Yamaguchi A, Aoki KF, Mamitsuka H (2003) Graph complexity of chemical compounds in biological pathways. Genome Inf 14:376–377

Yanardag P (2015) Deep graph kernels (code). http://www.mit.edu/pinary/kdd/DEEP_GRAPH_KERNELS_CODE.tar.gz Yanardag P, Vishwanathan SVN (2015a) Deep graph kernels. In: ACM SIGKDD International Conference on Knowledge

Discovery and Data Mining. pp 1365–1374. https://doi.org/10.1145/2783258.2783417 Yanardag P, Vishwanathan SVN (2015b) A structural smoothing framework for robust graph comparison. In: Advances in

Neural Information Processing Systems. pp 2134–2142 Zhang Y, Wang L, Wang L (2018a) A comprehensive evaluation of graph kernels for unattributed graphs. Entropy 20(12):984

Zhang Z, Wang M, Xiang Y, Huang Y, Nehorai A (2018b) Retgk: Graph kernels based on return probabilities of random walks. In: Advances in Neural Information Processing Systems. pp 3964–3974

Publisher's Note

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

Submit your manuscript to a SpringerOpen[●] journal and benefit from:

- Convenient online submission
- ► Rigorous peer review
- ► Open access: articles freely available online
- ► High visibility within the field
- ► Retaining the copyright to your article

Submit your next manuscript at > springeropen.com