# Marginalized Kernels Between Labeled Graphs 

Hisashi Kashima<br>HKASHIMA@JP.IBM.COM<br>IBM Tokyo Research Laboratory, 1623-14 Shimotsuruma, Yamato-shi, 242-8502 Kanagawa, Japan<br>Koji Tsuda<br>KOJI.TSUDA@TUEBINGEN.MPG.DE<br>Max Planck Institute for Biological Cybernetics, Spemannstr. 38, 72076 Tübingen, Germany; and<br>AIST Computational Biology Research Center, 2-43 Aomi, Koto-ku, 135-0064 Tokyo, Japan

Akihiro Inokuchi<br>INOKUCHI@JP.IBM.COM IBM Tokyo Research Laboratory, 1623-14 Shimotsuruma, Yamato-shi, 242-8502 Kanagawa, Japan


#### Abstract

A new kernel function between two labeled graphs is presented. Feature vectors are defined as the counts of label paths produced by random walks on graphs. The kernel computation finally boils down to obtaining the stationary state of a discrete-time linear system, thus is efficiently performed by solving simultaneous linear equations. Our kernel is based on an infinite dimensional feature space, so it is fundamentally different from other string or tree kernels based on dynamic programming. We will present promising empirical results in classification of chemical compounds. ${ }^{1}$


## 1. Introduction

A large amount of the research in machine learning is concerned with classification and regression for realvalued vectors (Vapnik, 1998). However, much of the real world data is represented not as vectors, but as graphs including sequences and trees, for example, biological sequences (Durbin et al., 1998), natural language texts (Manning \& Schütze, 1999), semistructured data such as HTML and XML (Abiteboul et al., 2000), and so on. Especially, in the pharmaceutical area, the chemical compounds are represented as labeled graphs, and their automatic classification is of crucial importance in the rationalization of drug discovery processes to predict the effectiveness or toxicity

[^0]of drugs from their chemical structures (Kramer \& De Raedt, 2001; Inokuchi et al., 2000).
Kernel methods such as support vector machines are becoming increasingly popular for their high performance (Schölkopf \& Smola, 2002). In kernel methods, all computations are done via a kernel function, which is the inner product of two vectors in a feature space. In order to apply kernel methods to graph classification, we first need to define a kernel function between the graphs. However, defining a kernel function is not an easy task, because it must be designed to be positive semidefinite ${ }^{2}$. Following the pioneering work by Haussler (1999), a number of kernels were proposed for structured data, for example, Watkins (2000), Jaakkola et al. (2000), Leslie et al. (2003), Lodhi et al. (2002), and Tsuda et al. (2002) for sequences, and Vishwanathan and Smola (2003), Collins and Duffy (2001), and Kashima and Koyanagi (2002) for trees. Most of them are based on the idea of an object decomposed into substructures (i.e. subsequences, subtrees or subgraphs) and a feature vector is composed of the counts of the substructures. As the dimensionality of feature vectors is typically very high, they deliberately avoid explicit computations of feature values, and adopt efficient procedures such as dynamic programming or suffix trees.
In this paper, we will construct a kernel function between two graphs, which is distinguished from the kernel between two vertices in a graph, e.g., diffusion kernels (Kondor \& Lafferty, 2002; Kandola et al., 2003; Lafferty \& Lebanon, 2003) or the kernel between two paths in a graph, e.g., path kernels (Takimoto \& Warmuth, 2002). There has been almost no significant works for designing kernels between two graphs except

[^1]for the work by Gärtner (2002). We will discuss about his kernels in Appendix.
One existing way to describe a labeled graph as a feature vector is to count label paths appearing in the graph. For example, the pattern discovery algorithm (Kramer \& De Raedt, 2001; Inokuchi et al., 2000) explicitly constructs a feature vector from the counts of frequently appearing label paths in a graph. For the labeled graph shown in Figure 1, a label path is produced by traversing the vertices, and looks like
$$
(A, e, A, d, D, a, B, c, D)
$$
where the vertex labels $A, B, C, D$ and the edge labels $a, b, c, d, e$ appear alternately. Numerous label paths are produced by traversing the nodes in every possible way. Especially when the graph has a loop, the dimensionality of the count vector is infinite, because traversing may never end. In order to explicitly construct the feature vector, we have to select features to keep the dimensionality finite. The label paths may be selected simply by limiting the path length or, more intelligently, the pattern discovery algorithm identifies the label paths which appear frequently in the training set of graphs. In any case, there is an additional unwanted parameter (i.e. a threshold on path length or path frequency) that has to be determined by a model selection criterion (e.g. the cross validation error).
We will propose a kernel function based on the inner product between infinite dimensional path count vectors. A label path is produced by random walks on graphs, and thus is regarded as a random variable. Our kernel is defined as the inner product of the count vectors averaged over all possible label paths, which is regarded as a special case of marginalized kernels (Tsuda et al., 2002). The kernel computation boils down to finding the stationary state of a discrete-time linear system (Rugh, 1995), which can be done efficiently by solving simultaneous linear equations with a sparse coefficient matrix. We especially notice that this computational trick is fundamentally different from the dynamic programming techniques adopted in other kernels (e.g. Watkins (2000), Lodhi et al. (2002), Collins and Duffy (2001), Kashima and Koyanagi (2002)). These kernels deal with very large but still finite dimensional feature spaces and have parameters to constrain the dimensionality (e.g. the maximum length of subsequences in Lodhi et al. (2002)). In order to investigate how our kernel performs well on the real data, we will show promising results on predicting the properties of chemical compounds.
This paper is organized as follows. In Section 2, we review marginalized kernels as the theoretical foundation. Then, a new kernel between graphs is presented in Section 3. In Section 4, we summarize the results of our experiments on the classification of chemical compounds. Finally, we conclude with Section 5.

## 2. Marginalized Kernels

A common way for constructing a kernel for structured data such as strings, trees, and graphs is to assume hidden variables and make use of the probability distribution of visible and hidden variables (Haussler, 1999; Watkins, 2000; Tsuda et al., 2002). For example, Watkins (2000) proposed conditional symmetric independence (CSI) kernels which are described as

$$
K\left(x, x^{\prime}\right)=\sum_{h} p(x \mid h) p\left(x^{\prime} \mid \boldsymbol{h}\right) p(h)
$$

where $\boldsymbol{h}$ is a hidden variable, and $\boldsymbol{x}$ and $\boldsymbol{x}^{\prime}$ are visible variables which correspond to structured data. This kernel requires $p(x \mid h)$, which means that the generation process of $\boldsymbol{x}$ from $\boldsymbol{h}$ needs to be known. However, it may be the case that $p(h \mid x)$ is known instead of $p(\boldsymbol{x} \mid \boldsymbol{h})$. Then we can use a marginalized kernel (Tsuda et al., 2002) which is described as

$$
\begin{equation*}
K\left(\boldsymbol{x}, \boldsymbol{x}^{\prime}\right)=\sum_{h} \sum_{h^{\prime}} K_{z}\left(z, z^{\prime}\right) p(h \mid x) p\left(h^{\prime} \mid \boldsymbol{x}^{\prime}\right) \tag{1}
\end{equation*}
$$

where $\boldsymbol{z}=[\boldsymbol{x}, h]$ and $K_{z}\left(\boldsymbol{z}, \boldsymbol{z}^{\prime}\right)$ is the joint kernel depending on both visible and hidden variables. The posterior probability $p(h \mid x)$ can be interpreted as a feature extractor that extracts informative features for classification from $x$. The marginalized kernel (1) is defined as the expectation of the joint kernel over all possible values of $\boldsymbol{h}$ and $\boldsymbol{h}^{\prime}$. In the following, we will construct a graph kernel in the context of the marginalized kernel. The hidden variable is a sequence of vertex indices, which is generated by random walks on the graph. Also the joint kernel $K_{z}$ is defined as a kernel between the sequences of vertex and edge labels traversed in the random walk.

## 3. Graph Kernel

In this section, we introduce a new kernel between graphs with vertex labels and edge labels. At the beginning, let us formally define a labeled graph. Denote by $G$ a labeled directed graph and by $|G|$ the number of vertices. Each vertex of the graph is uniquely indexed from 1 to $|G|$. Let $v_{i} \in \Sigma_{v}$ denote the label of vertex $i$ and $e_{i j} \in \Sigma_{E}$ denote the label of the edge from $i$ to $j$. Figure 1 shows an example of the graphs that we handle in this paper. We assume that there are no multiple edges between any vertices. Our task is to construct a kernel function $K\left(G, G^{\prime}\right)$ between two labeled graphs $G$ and $G^{\prime}$.

### 3.1. Random Walks on Graphs

The hidden variable $h=\left(h_{1}, \ldots, h_{\ell}\right)$ associated with graph $G$ is a sequence of natural numbers from 1 to $|G|$. Given a graph $G, h$ is generated by a random walk as follows: At the first step, $h_{1}$ is sampled from the


Figure 1. an example of graphs with labels
initial probability distribution $p_{s}(h)$. Subsequently, at the $i$-th step, the next vertex $h_{i}$ is sampled subject to the transition probability $p_{t}\left(h_{i} \mid h_{i-1}\right)$, but the random walk may also end with probability $p_{q}\left(h_{i-1}\right)$ :

$$
\begin{equation*}
\sum_{j=1}^{|G|} p_{t}(j \mid i)+p_{q}(i)=1 \tag{2}
\end{equation*}
$$

The posterior probability for the label path $\boldsymbol{h}$ is described as

$$
p(\boldsymbol{h} \mid G)=p_{s}\left(h_{1}\right) \prod_{i=2}^{\ell} p_{t}\left(h_{i} \mid h_{i-1}\right) p_{q}\left(h_{\ell}\right)
$$

where $\ell$ is the length of $h$.
When we do not have any prior knowledge, we cab set $p_{s}$ to be the uniform distribution, the transition probability $p_{t}$ to be an uniform distribution over the vertices adjacent to the current vertex, and the termination probability $p_{q}$ to be a constant. Also gaps (Lodhi et al., 2002) can be incorporated by setting the transition probabilities appropriately.

### 3.2. Joint Kernel

Next we define the joint kernel $K_{z}$ assuming that the hidden sequences $\boldsymbol{h}$ and $\boldsymbol{h}^{\prime}$ are given for two graphs $G$ and $G^{\prime}$, respectively. When the random walk is done as described in $\boldsymbol{h}$, the traversed labels are listed as

$$
v_{h_{1}} e_{h_{1} h_{2}} v_{h_{2}} e_{h_{2} h_{3}} v_{h_{3}} \cdots
$$

Assume that two kernel functions, $K\left(v, v^{\prime}\right)$ and $K\left(e, e^{\prime}\right)$, are readily defined between vertex labels and edge labels, respectively. We constrain both kernels $K\left(v, v^{\prime}\right), K\left(e, e^{\prime}\right) \geq 0$ to be nonnegative ${ }^{3}$. An example of the vertex label kernels is

$$
\begin{equation*}
K\left(v, v^{\prime}\right)=\delta\left(v=v^{\prime}\right) \tag{3}
\end{equation*}
$$

where $\delta$ is a function that returns 1 if its argument holds, 0 otherwise. If the labels are defined in $\mathbb{R}$, the Gaussian kernel

$$
\begin{equation*}
K\left(v, v^{\prime}\right)=\exp \left(-\left\|v-v^{\prime}\right\|^{2} / 2 \sigma^{2}\right) \tag{4}
\end{equation*}
$$

[^2]would be a natural choice (Schölkopf \& Smola, 2002). The joint kernel is defined as the product of the label kernels as follows:

$K_{z}\left(\boldsymbol{z}, \boldsymbol{z}^{\prime}\right)=\left\{\begin{array}{lr}0 & \left(\ell \neq \ell^{\prime}\right) \\ K\left(v_{h_{1}}, v_{h_{1}^{\prime}}^{\prime}\right) \prod_{i=2}^{\ell} K\left(e_{h_{i-1} h_{i}}, e_{h_{i-1}^{\prime} h_{i}^{\prime}}^{\prime}\right) \times \\ K\left(v_{h_{i}}, v_{h_{i}^{\prime}}^{\prime}\right) & \left(\ell=\ell^{\prime}\right)\end{array}\right.$,
where $\boldsymbol{z}=(G, \boldsymbol{h})$.

### 3.3. Efficient Computation

The marginalized kernel (1) is the expectation of $K_{z}$ over all possible $\boldsymbol{h}$ and $\boldsymbol{h}^{\prime}$, which is described as

$$
\begin{aligned}
& K\left(G, G^{\prime}\right) \\
& =\sum_{\ell=1}^{\infty} \sum_{h} \sum_{h^{\prime}} p_{s}\left(h_{1}\right) \prod_{i=2}^{\ell} p_{t}\left(h_{i} \mid h_{i-1}\right) p_{q}\left(h_{l}\right) \times \\
& p_{s}^{\prime}\left(h_{1}^{\prime}\right) \prod_{j=2}^{\ell} p_{t}^{\prime}\left(h_{j}^{\prime} \mid h_{j-1}^{\prime}\right) p_{q}^{\prime}\left(h_{\ell}^{\prime}\right) \times \\
& K\left(v_{h_{1}}, v_{h_{1}^{\prime}}^{\prime}\right) \prod_{k=2}^{\ell} K\left(e_{h_{k-1} h_{k}}, e_{h_{k-1}^{\prime}, h_{k}^{\prime}}^{\prime}\right) K\left(v_{h_{k}}, v_{h_{k}^{\prime}}^{\prime}\right),
\end{aligned}
$$

where $\sum_{h}:=\sum_{h_{1}=1}^{|G|} \cdots \sum_{h_{\ell}=1}^{|G|}$. The straightforward enumeration is obviously impossible, because $\ell$ spans from 1 to infinity. Nevertheless we have an efficient way to compute this kernel as shown below. By rearranging the terms, the kernel has the following nested structure.

$$
\begin{align*}
& K\left(G, G^{\prime}\right)=\lim _{L \rightarrow \infty} \sum_{\ell=1}^{L} \sum_{h_{1}, h_{1}^{\prime}} s\left(h_{1}, h_{1}^{\prime}\right) \times  \tag{5}\\
& \left(\sum _ { h _ { 2 } , h _ { 2 } ^ { \prime } } t ( h _ { 2 } , h _ { 2 } ^ { \prime } , h _ { 1 } , h _ { 1 } ^ { \prime } ) \left(\sum_{h_{3}, h_{3}^{\prime}} t\left(h_{3}, h_{3}^{\prime}, h_{2}, h_{2}^{\prime}\right) \times\right.\right. \\
& \left.\left(\cdots\left(\sum_{h_{\ell}, h_{\ell}^{\prime}} t\left(h_{\ell}, h_{\ell}^{\prime}, h_{\ell-1}, h_{\ell-1}^{\prime}\right) q\left(h_{\ell}, h_{\ell}^{\prime}\right)\right)\right) \cdots\right)
\end{align*}
$$

where

$$
\begin{gathered}
s\left(h_{1}, h_{1}^{\prime}\right):=\quad p_{s}\left(h_{1}\right) p_{s}^{\prime}\left(h_{1}^{\prime}\right) K\left(v_{h_{1}}, v_{h_{1}^{\prime}}^{\prime}\right) \\
t\left(h_{i}, h_{i}^{\prime}, h_{i-1}, h_{i-1}^{\prime}\right) \quad:=p_{t}\left(h_{i} \mid h_{i-1}\right) p_{t}^{\prime}\left(h_{i}^{\prime} \mid h_{i-1}^{\prime}\right) \times \\
K\left(v_{h_{i}}, v_{h_{i}^{\prime}}^{\prime}\right) K\left(e_{h_{i-1} h_{i}}, e_{h_{i-1}^{\prime} h_{i}^{\prime}}\right) \\
q\left(h_{\ell}, h_{\ell}^{\prime}\right):=p_{q}\left(h_{\ell}\right) p_{q}^{\prime}\left(h_{\ell}^{\prime}\right)
\end{gathered}
$$

Let us further simplify (5) as

$$
K\left(G, G^{\prime}\right)=\lim _{L \rightarrow \infty} \sum_{\ell=1}^{L} \sum_{h_{1}, h_{1}^{\prime}} s\left(h_{1}, h_{1}^{\prime}\right) r_{\ell}\left(h_{1}, h_{1}^{\prime}\right)
$$

where for $\ell \geq 2$,

$$
\begin{aligned}
& r_{\ell}\left(h_{1}, h_{1}^{\prime}\right) \\
& :=\left(\sum _ { h _ { 2 } , h _ { 2 } ^ { \prime } } t ( h _ { 2 } , h _ { 2 } ^ { \prime } , h _ { 1 } , h _ { 1 } ^ { \prime } ) \left(\sum_{h_{3}, h_{3}^{\prime}} t\left(h_{3}, h_{3}^{\prime}, h_{2}, h_{2}^{\prime}\right) \times\right.\right. \\
& \left.\left(\cdots\left(\sum_{h_{\ell}, h_{\ell}^{\prime}} t\left(h_{\ell}, h_{\ell}^{\prime}, h_{\ell-1}, h_{\ell-1}^{\prime}\right) q\left(h_{\ell}, h_{\ell}^{\prime}\right)\right)\right) \cdots\right),
\end{aligned}
$$

and $r_{1}\left(h_{1}, h_{1}^{\prime}\right):=q\left(h_{1}, h_{1}^{\prime}\right)$. Replacing the order of summation, we have the following:

$$
\begin{align*}
K\left(G, G^{\prime}\right) & =\sum_{h_{1}, h_{1}^{\prime}} s\left(h_{1}, h_{1}^{\prime}\right) \lim _{L \rightarrow \infty} \sum_{\ell=1}^{L} r_{\ell}\left(h_{1}, h_{1}^{\prime}\right) \\
& =\sum_{h_{1}, h_{1}^{\prime}} s\left(h_{1}, h_{1}^{\prime}\right) \lim _{L \rightarrow \infty} R_{L}\left(h_{1}, h_{1}^{\prime}\right) \tag{6}
\end{align*}
$$

where

$$
R_{L}\left(h_{1}, h_{1}^{\prime}\right):=\sum_{\ell=1}^{L} r_{\ell}\left(h_{1}, h_{1}^{\prime}\right)
$$

Thus we need to compute $R_{\infty}\left(h_{1}, h_{1}^{\prime}\right)$ to obtain $K\left(G, G^{\prime}\right)$.

Now let us restate this problem in terms of linear system theory (Rugh, 1995). The following recursive relationship holds between $r_{k}$ and $r_{k-1}(k \geq 2)$ :

$$
\begin{equation*}
r_{k}\left(h_{1}, h_{1}^{\prime}\right)=\sum_{i, j} t\left(i, j, h_{1}, h_{1}^{\prime}\right) r_{k-1}(i, j) \tag{7}
\end{equation*}
$$

Using (7), the recursive relationship for $R_{L}$ also holds as follows:

$$
\begin{align*}
R_{L}\left(h_{1}, h_{1}^{\prime}\right) & =r_{1}\left(h_{1}, h_{1}^{\prime}\right)+\sum_{k=2}^{T} r_{k}\left(h_{1}, h_{1}^{\prime}\right) \\
& =r_{1}\left(h_{1}, h_{1}^{\prime}\right)+\sum_{k=2}^{T} \sum_{i, j} t\left(i, j, h_{1}, h_{1}^{\prime}\right) r_{k-1}(i, j) \\
& =r_{1}\left(h_{1}, h_{1}^{\prime}\right)+\sum_{i, j} t\left(i, j, h_{1}, h_{1}^{\prime}\right) R_{L-1}(i, j) . \tag{8}
\end{align*}
$$

Thus $R_{L}$ is perceived as a discrete-time linear system (Rugh, 1995) evolving as the time $L$ increases. Assuming that $R_{L}$ converges, (see Sec. 3.4 for the convergence condition), we have the following equilibrium equation:

$$
\begin{equation*}
R_{\infty}\left(h_{1}, h_{1}^{\prime}\right)=r_{1}\left(h_{1}, h_{1}^{\prime}\right)+\sum_{i, j} t\left(i, j, h_{1}, h_{1}^{\prime}\right) R_{\infty}(i, j) \tag{9}
\end{equation*}
$$

Therefore, the computation of the marginalized kernel finally comes down to solving linear simultaneous equations (9) and substituting the solutions into (6).

Computing the kernel requires solving a linear equation with a $|G|\left|G^{\prime}\right| \times|G|\left|G^{\prime}\right|$ coefficient matrix. However, the matrix is actually sparse because the number of non-zero elements is less than $c^{2}|G|\left|G^{\prime}\right|$ where $c$ is the maximum degree. Therefore, we can employ various kinds of efficient numerical algorithms that exploit sparsity (Barrett et al., 1994). In our implementation, we employed a simple iterative method that updates current solutions by using (8) until convergence starting from $R_{1}\left(h_{1}, h_{1}^{\prime}\right)=r_{1}\left(h_{1}, h_{1}^{\prime}\right)$.

### 3.4. Convergence Condition

The convergence condition needed to justify (9) is described as follows:
Theorem $1 . \quad$ The infinite
$\lim _{L \rightarrow \infty} R_{L}\left(h_{1}, h_{1}^{\prime}\right) \quad$ sequence
$h_{1} \quad \in$ $\{1, \cdots,|G|\}$ and $h_{1}^{\prime} \in\left\{1, \cdots,\left|G^{\prime}\right|\right\}$, if the following inequality holds for $i_{0} \in\{1, \cdots,|G|\}$ and $j_{0} \in\left\{1, \cdots,\left|G^{\prime}\right|\right\}$,

$$
\begin{equation*}
\sum_{i=1}^{|G|} \sum_{j=1}^{\mid G \prime} \mid t\left(i, j, i_{0}, j_{0}\right) q(i, j)<q\left(i_{0}, j_{0}\right) \tag{10}
\end{equation*}
$$

(proof) $R_{\ell}\left(h_{1}, h_{1}^{\prime}\right)$ can also be described as

$$
\begin{aligned}
& R_{\ell}\left(h_{1}, h_{1}^{\prime}\right)=q\left(h_{1}, h_{1}^{\prime}\right)+ \\
& \sum_{i=2}^{\ell} \sum_{h_{2}, h_{2}^{\prime}} t\left(h_{2}, h_{2}^{\prime}, h_{1}, h_{1}^{\prime}\right)\left(\sum_{h_{3}, h_{3}^{\prime}} t\left(h_{3}, h_{3}^{\prime}, h_{2}, h_{2}^{\prime}\right) \times\right. \\
& \left(\cdots\left(\sum_{h_{i}, h_{i}^{\prime}} t\left(h_{i}, h_{i}^{\prime}, h_{i-1}, h_{i-1}^{\prime}\right) q\left(h_{i}, h_{i}^{\prime}\right)\right) \cdots\right) \\
& :=q\left(h_{1}, h_{1}^{\prime}\right)+\sum_{i=2}^{\ell} u_{i}\left(h_{1}, h_{1}^{\prime}\right)
\end{aligned}
$$

where $u_{i}(i \geq 2)$ denotes the $i$-th term of the series. According to the ratio test, this series converges if $\limsup \mathrm{m}_{i \rightarrow \infty}\left|u_{i}\left(h_{1}, h_{1}^{\prime}\right)\right| /\left|u_{i-1}\left(h_{1}, h_{1}^{\prime}\right)\right|<1$. Since $K\left(v, v^{\prime}\right)$ and $K\left(e, e^{\prime}\right)$ are nonnegative as previously defined, $u_{i}\left(h_{1}, h_{1}^{\prime}\right) \geq 0$, and thus what. we need to prove is limsup $\operatorname{sum}_{i \rightarrow \infty} u_{i}\left(\bar{h}_{1}, h_{1}^{\prime}\right) / u_{i-1}\left(h_{1}, h_{1}^{\prime}\right)<1$. For this purpose, it is sufficient to show

$$
\begin{equation*}
\frac{u_{i}\left(h_{1}, h_{1}^{\prime}\right)}{u_{i-1}\left(h_{1}, h_{1}^{\prime}\right)}<1 \tag{11}
\end{equation*}
$$

for all $i$. The two terms in (11) are written as

$$
\begin{align*}
u_{i}\left(h_{1}, h_{1}^{\prime}\right)= & \sum_{h_{i-1}, h_{i-1}^{\prime}} a_{i-1}\left(h_{i-1}, h_{i-1}^{\prime}, h_{1}, h_{1}^{\prime}\right) \times  \tag{12}\\
& \sum_{h_{1}, h_{i}^{\prime}} t\left(h_{i}, h_{i}^{\prime}, h_{i-1}, h_{i-1}^{\prime}\right) q\left(h_{i}, h_{i}^{\prime}\right) \\
u_{i-1}\left(h_{1}, h_{1}^{\prime}\right)= & \sum_{h_{i-1}, h_{i-1}^{\prime}} a_{i-1}\left(h_{i-1}, h_{i-1}^{\prime}, h_{1}, h_{1}^{\prime}\right) q\left(h_{i-1}, h_{i-1}^{\prime}\right) \tag{13}
\end{align*}
$$

where

$$
\begin{aligned}
& a_{i-1}\left(h_{i-1}, h_{i-1}^{\prime}, h_{1}, h_{1}^{\prime}\right):= \\
& \sum_{h_{2}, h_{2}^{\prime}} t\left(h_{2}, h_{2}^{\prime}, h_{1}, h_{1}^{\prime}\right)\left(\sum_{h_{3}, h_{3}^{\prime}} t\left(h_{3}, h_{3}^{\prime}, h_{2}, h_{2}^{\prime}\right) \times\right. \\
& \left(\cdots \left(\sum_{h_{i-2}, h_{i-2}^{\prime}} t\left(h_{i-2}, h_{i-2}^{\prime}, h_{i-3}, h_{i-3}^{\prime}\right) \times\right.\right. \\
& \left.\left.t\left(h_{i-1}, h_{i-1}^{\prime}, h_{i-2}, h_{i-2}^{\prime}\right)\right) \cdots\right)
\end{aligned}
$$

Notice that $a_{i-1}\left(h_{i-1}, h_{i-1}^{\prime}, h_{1}, h_{1}^{\prime}\right) \geq 0$ as well. Substituting (12) and (13) into (11), we have

$$
\begin{aligned}
& \sum_{h_{i-1}, h_{i-1}^{\prime}} a_{i-1}\left(h_{i-1}, h_{i-1}^{\prime}, h_{1}, h_{1}^{\prime}\right) \times \\
& \quad \sum_{h_{i}, h_{i}^{\prime}} t\left(h_{i}, h_{i}^{\prime}, h_{i-1}, h_{i-1}^{\prime}\right) q\left(h_{i}, h_{i}^{\prime}\right) \\
& \quad<\sum_{h_{i-1}, h_{i-1}^{\prime}} a_{i-1}\left(h_{i-1}, h_{i-1}^{\prime}, h_{1}, h_{1}^{\prime}\right) q\left(h_{i-1}, h_{i-1}^{\prime}\right) .
\end{aligned}
$$

Both sides of this inequality are the linear combinations of $a_{i-1}\left(h_{i-1}, h_{i-1}^{\prime}, h_{1}, h_{1}^{\prime}\right)$ with nonnegative coefficients. In order to prove the inequality, it is sufficient to prove the following inequalities for coefficients:

$$
\sum_{i, j} t\left(i, j, i_{0}, j_{0}\right) q(i, j)<q\left(i_{0}, j_{0}\right), \quad \forall i_{0}, j_{0}
$$

which we have assumed to hold in (10).
The condition (10) seems rather complicated, but we can have a simpler condition, if the termination probabilities are constant over all vertices.
Lemma 1. If $p_{q}(i)=p_{q}^{\prime}(j)=\gamma$ for any $i$ and $j$, the infinite sequence $\lim _{L \rightarrow \infty} R_{L}\left(h_{1}, h_{1}^{\prime}\right)$ converges if

$$
\begin{equation*}
K\left(v, v^{\prime}\right) K\left(e, e^{\prime}\right)<\frac{1}{(1-\gamma)^{2}} \tag{14}
\end{equation*}
$$

(proof) Due to the assumption, $q(i, j)=q\left(i_{0}, j_{0}\right)=\gamma^{2}$. Thus it is sufficient to prove $\sum_{i, j} t\left(i, j, i_{0}, j_{0}\right)<1$, that is,

$$
\begin{equation*}
\sum_{i, j} p_{t}\left(i \mid i_{0}\right) p_{t}^{\prime}\left(j \mid j_{0}\right) K\left(v_{i}, v_{j}\right) K\left(e_{i_{0} i}, e_{j_{0} j}\right)<1 \tag{15}
\end{equation*}
$$

Due to Relation (2), we observe that
$\sum_{i, j} p_{t}\left(i \mid i_{0}\right) p_{t}^{\prime}\left(j \mid j_{0}\right)=\sum_{i} p_{t}\left(i \mid i_{0}\right) \sum_{j} p_{t}^{\prime}\left(j \mid j_{0}\right)=(1-\gamma)^{2}$.
Thus (15) holds if all the coefficients satisfy $K\left(v_{i}, v_{j}\right) K\left(e_{i_{0} i}, e_{j_{0} j}\right)<1 /(1-\gamma)^{2}$.
Apparently, the above lemma holds for $\lambda>0$ if $K(\cdot, \cdot) \leq 1$. Standard label kernels such as (3) and (4) satistfy this condition.


Figure 2. A chemical compound is conventionally represented as an undirected graph (left). Atom types and bond types correspond to vertex labels and edge labels, respectively. The edge labels 's' and 'd' denote single and double bonds, respectively. As our kernel assumes a directed graph, undirected edges are replaced by directed edges (right).

## 4. Experiments

We applied our kernel to prediction of the properties of chemical compounds. A chemical compound can naturally be represented as an undirected graph by considering the atom types as the vertex labels, e.g. $\mathrm{C}, \mathrm{Cl}$ and H , and the bond types as the edge labels, e.g. $s$ (single bond) and $d$ (double bond). For our graph kernel, we replaced an undirected edge by two directed edges (Figure 2) since the kernel assumes directed graphs.

### 4.1. Pattern Discovery Algorithm

We compare our graph kernel with the patterndiscovery (PD) method by Kramer and De Raedt (2001) that is one of the best state-of-the-art methods in predictive toxicology. As in our graph kernel, each feature is constructed as the count that a particular label path appears in the graph ${ }^{4}$. There are other methods which count more complicated substructures such as subgraphs (Inokuchi et al., 2000), but we focus on Kramer and De Raedt (2001) whose features are similar to ours.

Assume that we have $n$ graphs $G_{1}, \ldots, G_{n}$. Also let us define $\#(\boldsymbol{h}, G)$ as the number of appearances of a label path $\boldsymbol{h}$ in $G$. The PD method identifies a set of all label paths $\mathcal{H}$ which appear in more than $m$ graphs:

$$
\mathcal{H}=\left\{\boldsymbol{h} \mid \sum_{i=1}^{n} \delta\left(\#\left(\boldsymbol{h}, G_{i}\right)>0\right) \geq m\right\}
$$

where the parameter $m$ is called the minimum support parameter. Furthermore, it is possible to add extra conditions, e.g., selecting only the paths frequent in a certain class, and scarce in the other classes. Then the feature vector of $G$ based on the identified label paths

[^3]is built as
\[

$$
\begin{equation*}
G \rightarrow\left(\#\left(h_{1}, G\right), \ldots, \#\left(h_{|\mathcal{H}|}, G\right)\right) \tag{16}
\end{equation*}
$$

\]

whose dimensionality is the cardinality of $\mathcal{H}$. The PD method is useful for extracting comprehensive features. However, as the minimum support parameter gets smaller, the dimensionality of feature vectors becomes so huge that a prohibitive amount of computation is required. Therefore, the user has to control the minimum support parameter $m$, such that the feature space does not lose necessary information and, at the same time, computation stays feasible.
The PD method contrasts markedly with our method. Our kernel method puts emphasis on dealing with infinite, but less interpretable features, while the PD method trys to extract a relatively small number of meaningful features. Looking at the algorithms, our method is so simple that it is described by just one equation (9), while the PD method's algorithm is rather complicated (De Raedt \& Kramer, 2001).

### 4.2. Datasets

We used two datasets, the PTC dataset (Helma et al., 2001) and the Mutag dataset (Srinivasan et al., 1996).

The PTC dataset is the results of the following pharmaceutical experiments. Each of 417 compounds is given to four types of test animals: Male Mouse (MM), Female Mouse (FM), Male Rat (MR) and Female Rat (FR). According to their carcinogenicity, each compound is assigned one of the following labels: $\{E E, I S, E, C E, S E, P, N E, N\}$ where $C E, S E$ and $P$ indicate "relatively active", and NE and N indicate "relatively inactive", and EE, IS and E indicate "can not be decided". In order to simplify the problem, we relabeled CE, SE and P as "positive", and NE and N as "negative". The task is to predict whether a given compound is positive or negative for each type of test animals. Thus we eventually had four two-class problems.

In the Mutag dataset, the task is defined to be a twoclass classification problem to predict whether each of the 188 compounds has mutagenicity or not.
Each statistics of the datasets are summarized in Table 1.

### 4.3. Experimental Settings and Results

Assuming no prior knowledge, we defined the probability distributions for random walks as follows. The initial probabilities were simply uniform, i.e. $p_{s}(h)=$ $1 /|G|, \forall h$. The termination probabilities were determined as a constant $\gamma$ over all vertices. The transition probabilities $p_{t}\left(h \mid h_{0}\right)$ were set as uniform over adjacent vertices. We used Equation (3) as the label kernels. In solving the simultanous equations, we employed a simple iterative method (8). In our observation, 20-30

Table 1. Several statistics of the datasets such as numbers of positive examples (\#positive) and negative examples (\#negative), maximum degree (max. degree), maximum size of graphs (max. $|G|$ ), average size of graphs (avg. $|G|$ ), and numbers of vertex labels ( $|\Sigma|_{v}$ ) and edge labels $\left(|\Sigma|_{e}\right)$.

|  | MM | FM | MR | FR | MUTAG |
| :--- | :---: | :---: | :---: | :---: | :---: |
| \#POSITIVE | 129 | 143 | 152 | 121 | 125 |
| \#NEGATIVE | 207 | 206 | 192 | 230 | 63 |
| MAX. $\|G\|$ | 109 | 109 | 109 | 109 | 40 |
| AVG. $\|G\|$ | 25.0 | 25.2 | 25.6 | 26.1 | 31.4 |
| MAX. DEGREE | 4 | 4 | 4 | 4 | 4 |
| $\left\|\sum\right\| v$ | 21 | 19 | 19 | 20 | 8 |
| $\left.\sum\right\|_{e}$ | 4 | 4 | 4 | 4 | 4 |

Table 2. Classification accuracies (\%) of the pattern discovery method. 'MinSup' shows the ratio of the minimum support parameter to the number of compounds $\mathrm{m} / \mathrm{n}$.

| MinSup | MM | FM | MR | FR | Mutag |
| :--- | :---: | :---: | :---: | :---: | :---: |
| $0.5 \%$ | 60.1 | 57.6 | 61.3 | 66.7 | 88.3 |
| $\mathbf{1} \%$ | 61.0 | 61.0 | 62.8 | 63.2 | 87.8 |
| $3 \%$ | 58.3 | 55.9 | 60.2 | 63.2 | 89.9 |
| $5 \%$ | 60.7 | 55.6 | 57.3 | 63.0 | 86.2 |
| $10 \%$ | 58.9 | 58.7 | 57.8 | 60.1 | 84.6 |
| $20 \%$ | 61.0 | 55.3 | 56.1 | 61.3 | 83.5 |

iterations were enough for convergence in all cases. For the classification algorithm, we used the voted kernel perceptron (Freund \& Shapire, 1999), whose the performance is known to be comparable to SVMs. In the pattern discovery method, the minimum support parameter was determined as $0.5,1,3,5,10,20 \%$ of the number of compounds, and the simple dot product in the feature space (16) was used as a kernel. In our graph kernel, the termination probability $\gamma$ was changed from 0.1 to 0.9 .
Table 2 and Table 3 show the classification accuracies in the five two-class problems measured by leave-oneout cross validation. No general tendencies were found to conclude which method is better (the PD was better in MR, FR and Mutag, but our method was better in MM and FM). Thus it would be fair to say that the performances were comparable in this small set of experiments. Even though we could not show that our method is constantly better, this result is still appealing, because the advantage of our method lies in its simplicity both in concepts and in computational procedures.

Table 3. Classification accuracies (\%) of our graph kernel. The parameter $\gamma$ is the termination probability of random walks, which controls the effect of the length of label paths.

| $\boldsymbol{\gamma}$ | MM | FM | MR | FR | MUTAG |
| :--- | :--- | :--- | :--- | :---: | :---: |
| 0.1 | 62.2 | 59.3 | 57.0 | 62.1 | 84.3 |
| 0.2 | 62.2 | 61.0 | 57.0 | 62.4 | $\mathbf{8 3 . 5}$ |
| 0.3 | 64.0 | 61.3 | 56.7 | 62.1 | $\mathbf{8 5 . 1}$ |
| 0.4 | $\mathbf{6 4 . 3}$ | 61.9 | 56.1 | 63.0 | $\mathbf{8 5 . 1}$ |
| 0.5 | 64.0 | 61.3 | 56.1 | 64.4 | 83.5 |
| 0.6 | 62.8 | 61.9 | 54.4 | 65.8 | 83.0 |
| 0.7 | 63.1 | 62.5 | 54.1 | 63.2 | 81.9 |
| 0.8 | 63.4 | $\mathbf{6 3 . 4}$ | 54.9 | 64.1 | 79.8 |
| 0.9 | 62.8 | 61.6 | $\mathbf{5 8 . 4}$ | $\mathbf{6 6 . 1}$ | $\mathbf{7 8 . 7}$ |

## 5. Conclusion

In this paper, we introduced a new kernel between graphs with vertex labels and edge labels in the framework of the marginalized kernel. We defined the kernel by using random walks on graphs, and reduced the computation of the kernel to solving a system of linear simulataneous equations. As contrasted with the pattern-discovery method, our kernel takes into account all possible label paths without computing feature values explicitly. The structure we dealt with in this paper is fairly general, as it includes sequences, trees, DAGs and so on. Thus, applications of this kernel are by no means limited to chemical compounds. Promising targets would be DNA and RNA sequences with remote correlations, HTML and XML documents, distance graphs of 3-D protein structures, and so on. Recently, we found that our kernel can also be interpreted as an instance of the rational kernels (Cortes et al., 2003) over probability semiring. In this interpretation, the probability distributions over a graph are considered as a weighted automaton. In furure works, we would like to explore how to deal with infinite dimensions when another semiring is adopted.

## Appendix: Relation to Exponential and Geometric Kernels for Labeled Graphs

Here we review the exponential and geometric kernels between graphs proposed by Gärtner (2002) and describe the relation to our kernel. Although his kernels are defined in a quite different formulation from ours, we relate them by using the same notations.

His kernels are defined as the inner product of two $\left|\Sigma_{v}\right| \times\left|\Sigma_{v}\right|$ feature matrices $M(G)$ and $M\left(G^{\prime}\right)$,

$$
K\left(G, G^{\prime}\right)=\sum_{a_{1}=1}^{\left|\Sigma_{v}\right|} \sum_{a_{2}=1}^{\left|\Sigma_{v}\right|} M_{a_{1}, a_{2}}(G) m_{a_{1}, a_{2}}^{\prime}\left(G^{\prime}\right)
$$

where the feature matrix is defined as

$$
\begin{equation*}
M_{a_{1}, a_{2}}(G)=\sum_{\ell=1}^{\infty} \sum_{h_{1}, h_{\ell}} p\left(h_{1}, h_{\ell}\right) \delta\left(v_{h_{1}}=\sigma_{a_{1}}\right) \delta\left(v_{h_{\ell}}=\sigma_{a_{2}}\right) \tag{17}
\end{equation*}
$$

and $\Sigma_{v}=\left\{\sigma_{1}, \sigma_{2}, \ldots, \sigma_{\left|\Sigma_{v}\right|}\right\}$.
Notice that the probability $p\left(h_{1}, h_{\ell}\right)$ is denoted as

$$
\begin{aligned}
p\left(h_{1}, h_{\ell}\right)= & p_{s}\left(h_{1}\right) p_{q}\left(h_{\ell}\right) \times \\
& \sum_{h_{2}} \cdots \sum_{h_{\ell}} p_{t}\left(h_{2} \mid h_{1}\right) \cdots p_{t}\left(h_{\ell} \mid h_{\ell-1}\right) .
\end{aligned}
$$

Let $T$ be a $|G| \times|G|$ matrix where $T_{i j}=p_{t}(j \mid i)$. Also, let $S$ and $Q$ be $|G| \times\left|\Sigma_{v}\right|$ matrices where $S_{i j}=$ $p_{s}(i) \delta\left(v_{i}=\sigma_{j}\right)$ and $Q_{i j}=p_{q}(i) \delta\left(v_{i}=\sigma_{j}\right)$, respectively. By means of these matrices, the feature matrix is rewritten as

$$
M(G)=\sum_{\ell=1}^{\infty} S^{\top} T^{\ell-1} Q
$$

In his settings, the parameters $p_{s}, p_{t}$ and $p_{q}$ do not have stochastic constraints, i.e. $p_{s}(i)=1$ and $p_{q}(i)=$ 1 for any vertex $i$. Therefore, $S=Q:=L$, where $L_{i j}=\delta\left(v_{i}=j\right)$.
The feature vector of geometric kernel, $M(G)=$ $\sum_{\ell=1}^{\infty} L^{\top}(\beta E)^{\ell} L$, is obtained by setting $T=\beta E$, where $E$ is the adjacency matrix of $G$ and $\beta$ is a constant to assure convergence.
In the case of exponential kernel, $M$ is defined as $M(G)=\sum_{\ell=1}^{\infty}\left(L^{\top}(\beta E)^{\ell-1} L\right) /(\ell-1)!$, which still results in a similar form.

As seen in (17), the primary difference from our approach is that they only take the vertex labels at both ends ( $h_{1}, h_{\ell}$ ) into account. Although this simplification allows to keep the size of the matrices small, it might be harmful when contiguous labels are essential as in chemical compounds. Another difference is that they do not use an infinite dimensional feature space. Although they also compute the infinite sum in Equation (17), the infinite sum is for obtaining $\left|\Sigma_{v}\right|^{2}$ dimensional feature vectors.

## References

Abiteboul, S., Buneman, P., \& Suciu, D. (2000). Data on the Web: from relations to semistructured data and XML. Morgan Kaufmann.

Bahlmann, C., Haasdonk, B., \& Burkhardt, H. (2002). On-line handwriting recognition with support vector machines - a kernel approach. Proc. of the 8th Int. Workshop on Frontiers in Handwriting Recognition (IWFHR) (pp. 49-54).

Barrett, R., Berry, M., Chan, T. F., Demmel, J., Donato, J., Dongarra, J., Eijkhout, V., Pozo, R.,

Romine, C., \& der Vorst, H. V. (1994). Templates for the solution of linear systems: Building blocks for iterative methods, 2nd edition. Philadelphia, PA: SIAM.

Collins, M., \& Duffy, N. (2001). Convolution kernel for natural language. Proc. of the 14 th NIPS.

Cortes, C., Haffner, P., \& Mohri, M. (2003). Rational kernels. Advances in Neural Information Processing Systems 15. MIT Press. In press.

De Raedt, L., \& Kramer, S. (2001). The levelwise version space algorithm and its application to molecular fragment finding. Proceedings of the 17 th International Joint Conference on Artificial Intelligence (pp. 853-862). Morgan Kaufmann.

Durbin, R., Eddy, S., Krogh, A., \& Mitchison, G. (1998). Biological sequence analysis: Probabilistic models of proteins and nucleic acids. Cambridge University Press.

Freund, Y., \& Shapire, R. (1999). Large margin classification using the perceptron algorithm. Machine Learning, 37, 277-296.

Gärtner, T. (2002). Exponential and geometric kernels for graphs. NIPS*02 Workshop on Unreal Data: Principles of Modeling Nonvectorial Data. Available from http://mlg.anu.edu.au/unrealdata/.

Haussler, D. (1999). Convolution kernels on discrete structures (Technical Report UCSC-CRL-9910). University of Calfornia in Santa Cruz.

Helma, C., King, R., Kramer, S., \& Srinivasan, A. (2001). The Predictive Toxicology Challenge 20002001. Bioinformatics, 17, 107-108.

Inokuchi, A., Washio, T., \& Motoda, H. (2000). An Apriori-based algorithm for mining frequent substructures from graph data. Proc. of the 4 th PKDD (pp. 13-23).
Jaakkola, T., Diekhans, M., \& Haussler, D. (2000). A discriminative framework for detecting remote protein homologies. Journal of Computational Biology, 7, 95-114.

Kandola, J., Shawe-Taylor, J., \& Cristianini, N. (2003). Learning semantic similarity. Advances in Neural Information Processing Systems 15. MIT Press. In press.

Kashima, H., \& Inokuchi, A. (2002). Kernels for graph classification. IEEE ICDM Workshop on Active Mining. Maebashi, Japan.

Kashima, H., \& Koyanagi, T. (2002). Kernels for semistructured data. Proc. of the 19th ICML (pp. 291298).

Kondor, R. I., \& Lafferty, J. (2002). Diffusion kernels on graphs and other discrete input spaces. Proc. of the 19th ICML (pp. 315-322).

Kramer, S., \& De Raedt, L. (2001). Feature construction with version spaces for biochemical application. Proc. of the 18 th ICML (pp. 258-265).
Lafferty, J., \& Lebanon, G. (2003). Information diffusion kernels. Advances in Neural Information Processing Systems 15. MIT Press. In press.

Leslie, C., Eskin, E., Weston, J., \& Noble, W. (2003). Mismatch string kernels for svm protein classification. Advances in Neural Information Processing Systems 15. MIT Press. In press.
Lodhi, H., Saunders, C., Shawe-Taylor, J., Cristianini, N., \& Watkins, C. (2002). Text classification using string kernels. Journal of Machine Learning Research, 2, 419-444.

Manning, C. D., \& Schütze, H. (1999). Foundations of statistical natural language processing. Cambridge, Massachusetts: The MIT Press.

Rugh, W. (1995). Linear system theory. Prentice Hall.
Schölkopf, B., \& Smola, A. J. (2002). Learning with kernels. Cambridge, MA: MIT Press.

Shimodaira, H., Noma, K.-I., Nakai, M., \& Sagayama, S. (2002). Dynamic time-alignment kernel in support vector machine. Advances in Neural Information Processing Systems 14 (pp. 921-928). Cambridge, MA: MIT Press.

Srinivasan, A., Muggleton, S., King, R. D., \& Sternberg, M. (1996). Theories for mutagenicity: a study of first-order and feature based induction. Artificial Intelligence, 85, 277-299.
Takimoto, E., \& Warmuth, M. K. (2002). Path kernels and multiplicative updates. Proc. of the 15th Annual Conference on Computational Learning Theory (pp. 74-89).
Tsuda, K., Kin, T., \& Asai, K. (2002). Marginalized kernels for biological sequences. Bioinformatics, 18, S268-S275.

Vapnik, V. (1998). Statistical learning theory. New York: Wiley.

Vishwanathan, S., \& Smola, A. (2003). Fast kernels for string and tree matching. Advances in Neural Information Processing Systems 15. MIT Press. In press.

Watkins, C. (2000). Dynamic alignment kernels. In A. Smola, P. Bartlett, B. Schölkopf and D. Schuurmans (Eds.), Advances in large margin classifiers, 39-50. MIT Press.


[^0]:    ${ }^{1}$ An extended abstract of this research has been presented in IEEE ICDM International Workshop on Active Mining at Maebashi, Japan (Kashima \& Inokuchi, 2002). But this full paper contains much richer theoretical analysises such as interpretation as marginalized kernels, convergence conditions, relationship to linear systems, and so on.

[^1]:    ${ }^{2}$ Ad hoc similarity functions are not always positive semidefinite, e.g. Shimodaira et al. (2002) and Bahlmann et al. (2002).

[^2]:    ${ }^{3}$ This constraint will play an important role in proving the convergence of the marginalized kernel in Section 3.4.

[^3]:    ${ }^{4}$ Notice that the definition of label paths is different from ours in several points, e.g. a vertex will not be visited twice in a path. See Kramer and De Raedt (2001) for details.

