

Machine learning for *in silico* virtual screening and chemical genomics: new strategies

Jean-Philippe Vert* and Laurent Jacob

Centre for Computational Biology, Ecole des Mines de Paris

35, rue Saint Honoré, 77305 Fontainebleau, France

{jean-philippe.vert, laurent.jacob}@ensmp.fr

Tel : +33-1-6469-4782

Fax : +33-1-6469-4705

January 28, 2008

Abstract

Support vector machines and kernel methods are a class of machine learning algorithms that have recently become prominent in both computational biology and chemistry, although both fields have largely ignored each other. These methods are based on a sound mathematical and computationally efficient framework that implicitly embed the data of interest, respectively proteins and small molecules, in a high-dimensional feature spaces where various classification or regression tasks can be performed with linear algorithms. In this review we present the main ideas underlying these approaches, survey how both the “biological” and the “chemical” spaces have been separately constructed using the same mathematical framework and tricks, and suggest different avenues to unify both spaces for the purpose of *in silico* chemogenomics.

*Corresponding author

1 Introduction

Recent years have witnessed impressive progress in computational biology and chemistry, due in large part to the formulation of many problems as classification or regression problems and to the use of modern machine learning techniques for the analysis of biological and chemical data. In computational biology, typical problems include the prediction of various structural or functional properties of proteins from the description of their amino acid sequence, 3D structure or level of expression in various experimental conditions [31, 46, 18, 56], while in computational chemistry the objects manipulated are rather small molecules and the challenges are to predict various chemical or biological properties of these molecules related to their toxicity, pharmacokinetics or affinity for a drug target [7, 25, 37]. Both fields are confronted with similar challenges, including the difficulty to manipulate complex and structured data such as molecules, the computational and statistical burden to deal with large-dimensional representations of these data, and the need to integrate heterogeneous description of the same objects, such as the 2D and the 3D structures of molecules.

Among the variety of tools available in statistics and machine learning [19, 26, 6], a recent class of algorithms, called *kernel methods*, lends itself particularly well to the study of these aspects [69, 67, 72]. The basic philosophy of kernel methods is that with the use of a certain type of similarity measure, called a *kernel*, the data are implicitly embedded in a possibly high-dimensional vector space, in which linear methods are used for various applications such as discrimination, regression, ranking, outlier detection or low-dimensional visualization. The best-known kernel method is the support vector machine (SVM) [11, 78, 15], which has quickly gained popularity among the research communities for its efficiency and relative ease of use, and can nowadays be considered as a standard algorithm for regression and supervised classification in both computational biology and chemistry [69, 79, 30].

Choosing a kernel similarity function between proteins or small molecules is a prerequisite to the use of kernel methods, akin to choosing a set of descriptors to represent them. In fact, the inner product or Tanimoto coefficient between fingerprints of descriptors, widely used in chemoinformatics to assess the similarity of small molecules, are valid kernels that can be directly used by kernel methods. In that case, the SVM can be considered as just another multivariate method for regression or classification, that can be used in place of other methods such as partial least squares (PLS) or artificial neural networks (ANN). Interestingly, kernels can also be extended to more general measures of similarity, involving for example an infinite number of descrip-

tors. In such cases, computational tricks are needed to compute the kernel without explicitly computing nor storing the infinitely many descriptors. This extension allows kernel methods to perform multivariate analysis in infinite dimension, often at no additional computational cost compared to the finite-dimensional setting. This interesting property has been investigated by a number of researchers recently, who have proposed ingenious kernel similarity functions to encode various properties of proteins and small molecules with a large or infinite number of descriptors [31, 77, 50, 24, 65, 40, 53, 59]. These kernels have been successfully applied to various problems in computational biology and chemistry, such as gene function classification or toxicity prediction for small molecules.

Although progresses in computational biology and chemistry have largely ignored each other, the new field of chemogenomics is likely to benefit from both fields simultaneously [45, 35, 36, 63, 42]. In particular, a central challenge of *in silico* chemogenomics is to develop a coherent framework and computationally efficient algorithms to manipulate data from both the biological and chemical spaces simultaneously. For example, an *in silico* screening of a series of potential ligands against a family of proteins requires to learn a function $f(p, l)$ that predicts whether the ligand l can bind to the target protein p . In that case the basic objects to be manipulated by machine learning methods are target-ligand pairs (p, l) , instead of individual targets and ligand candidates. We argue in this review that the parallel development of kernel methods in computational biology and chemistry represents a timely and interesting opportunity to provide such a framework. Indeed, both fields have concentrated a lot of energy, through the design of descriptors and kernels, to construct various representations of the biological and chemical spaces, and kernel methods provide systematic tools to combine these different representations into a joint chemical and biological space where *in silico* chemogenomics can take place. Although the idea of creating *ad hoc* representations of target-ligand pairs for the purpose of chemogenomics is not new [47, 8], the use of kernels can provide additional flexibility and computational advantages over the explicit representation of target-ligand pairs as vectors of descriptors, as illustrated by several recent studies [8, 21, 34, 33, 32].

We hope that this short review will contribute to the development of further research on the application of kernel methods in chemogenomics. For that purpose, we start by a quick description of the main ideas and definitions underlying kernels and kernel methods in Section 2, before reviewing the recent attempts to design specific kernels for proteins (Section 3) and small molecules (Section 4). We finally present recent attempts to

2 Kernels and kernel methods

Many widely-used statistical and machine learning algorithms, including for example PLS or ANN, are designed to manipulate vector data. Using these tools to manipulate and analyze proteins or small molecules therefore poses the problem of representing these data as vectors or, equivalently, defining a set of binary or real-valued descriptors for these data and stacking them to form a vector. The design of molecular descriptors to describe various features of proteins or small molecules has indeed been much investigated over the last decades, and many such descriptors are nowadays routinely used in combination with statistical methods to correlate the structures of molecules with their physicochemical or biological properties. The explicit computation of a finite number p of molecular descriptors to represent a molecule x by a vector $\Phi(x) = (\Phi_1(x), \dots, \Phi_p(x))$ nevertheless raises several challenges, including the problem of choosing a set of descriptors sufficiently large to capture the relevant features for a given problem and sufficiently small to allow fast computation and efficient storage.

Kernel methods, including SVM, are a class of algorithms that follow a slightly different strategy to solve the problem of data representation [69, 67, 72]. Data do not need to be represented individually as vectors, they need instead to be compared to each other. More formally, instead of converting each protein or small molecule x into a p -dimensional vector $\Phi(x)$ for further processing, kernel methods can manipulate data only through a function $k(x, x')$ that compares any two proteins (or small molecules) x and x' and returns a real number. The function k is called a *kernel*, hence the name kernel methods. As a result, when a set of n proteins (or of n small molecules) x_1, \dots, x_n is given as input to a kernel method, the algorithm can only manipulate the data through the *Gram matrix*, which is the square $n \times n$ matrix K of all pairwise similarities, whose entry $K_{i,j}$ is equal to $k(x_i, x_j)$.

Only a certain class of functions k , however, can be used in combination with kernel methods. These kernels are often called *positive definite kernels* or more simply *valid kernel*. The technical conditions that a function $k(x, x')$ must fulfill to be a valid kernel are (i) to be symmetric, i.e., $k(x, x') = k(x', x)$ for any pair of data x and x' , and (ii) to produce only Gram matrices with no negative eigenvalue. Although the second condition can sometimes be difficult to assess for a newly defined function k , mathematics textbook abound on examples of

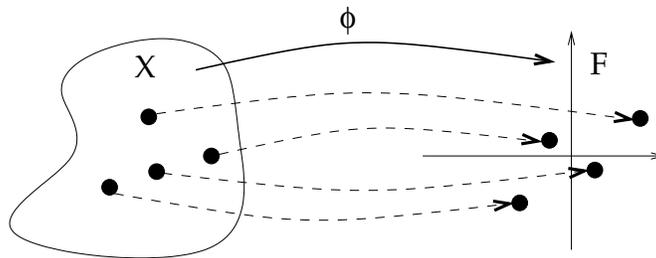


Figure 1: Defining a kernel over a space X , such as the space of all small molecules or the space of all proteins, is equivalent to embedding X in a vector space F of finite or infinite dimension through a mapping $\Phi : X \mapsto F$. The kernel between two points in X is equal to the inner products of their images in F , as shown in (1)

valid kernels and on systematic techniques to create them [1, 4, 5].

In fact, many valid kernels have been widely used for a long time in chemoinformatics. For example, given any representation of a molecule x by a vector of p descriptors $\Phi(x)$, the inner product between two vectors $\Phi(x)$ and $\Phi(x')$ representing two molecules x and x' is a valid kernel:

$$k(x, x') = \langle \Phi(x), \Phi(x') \rangle = \sum_{i=1}^p \Phi_i(x) \Phi_i(x'). \quad (1)$$

Another valid kernel is the Tanimoto coefficient between vectors of binary descriptors [60], which is commonly used to assess the similarity between molecules. When such kernels are used, the vectors of descriptors $\Phi(x)$ are explicitly computed prior to the computation of inner products or Tanimoto coefficients, and kernel methods like SVM are not fundamentally different from other methods such as PLS or ANN.

Interestingly, it can be shown that, conversely, any valid kernel $k(x, x')$ can be written as an inner product (1), for some vector representation $\Phi(x)$ [1]. This apparently establishes an equivalence between the use of valid kernels, on the one hand, and the use of explicit vector representations, on the other hand. In the converse statement, however, the vector $\Phi(x)$ are not necessarily of finite dimension, they can also involve an infinite number of descriptors. In that case, there is obviously no hope to compute the infinitely many descriptors explicitly and store them in a computer, and a computational trick must be found to compute directly the kernel $k(x, x')$ without computing $\Phi(x)$ and $\Phi(x')$. We review several examples of such kernels in the next two sections. As a result, the kernel approach can be seen as a generalization of the descriptor vector approach, where the number of descriptors can be finite or infinite (Figure 1).

Valid kernels therefore always define a vector space structure over the set of molecules to be manipulated.

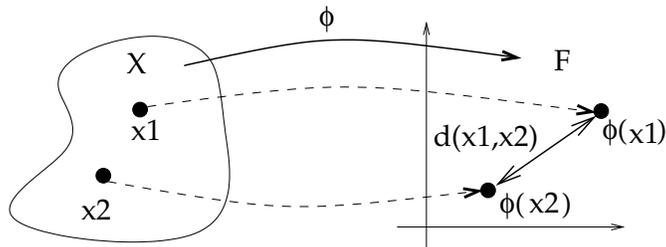


Figure 2: We can define the distance between two objects x_1 and x_2 , such as two small molecules or proteins, as the Euclidean distance between their images $\Phi(x_1)$ and $\Phi(x_2)$. If the mapping Φ is defined by a valid kernel k , then this distance can be computed easily without computing $\Phi(x_1)$ and $\Phi(x_2)$, as shown in (2). This *kernel trick* can be extended to a variety of linear algorithms that only manipulate the data through inner products.

This structure can either be defined *explicitly*, when molecular descriptors are computed in order to evaluate the kernel similarity through inner products of Tanimoto coefficients between fingerprints, or *implicitly*, when a valid kernel function $k(x, x')$ is directly computed to compare two molecules x and x' . Yet this implicit construction is sufficient to perform various data processing and manipulation in the vector space. As a simple illustration let us consider the problem of computing the distance between two feature vectors $\Phi(x)$ and $\Phi(x')$ corresponding to two molecules x and x' , as illustrated in Figure 2. A simple computation shows that:

$$\begin{aligned} \|\Phi(x) - \Phi(x')\|^2 &= \langle \Phi(x), \Phi(x) \rangle + \langle \Phi(x'), \Phi(x') \rangle - 2 \langle \Phi(x), \Phi(x') \rangle \\ &= k(x, x) + k(x', x') - 2k(x, x'), \end{aligned} \quad (2)$$

where k is the kernel associated to the vector Φ through (1). This equation shows that in order to compute the distance between points in the feature space, one does not necessarily need to first compute explicitly the vectors themselves, and can rely instead on the corresponding kernel. This trick, known as the *kernel trick*, can be applied to any algorithm for vectors that can be expressed in terms of inner products: replacing each inner product by the respective kernel evaluation allows to execute the algorithm implicitly in the feature space defined by a valid kernel. A surprising variety of methods, collectively known as *kernel methods*, can benefit from this trick. Besides the evaluation of distances using (2), one can mention for example dimensionality reduction with principal component analysis (PCA) [68], regression and pattern recognition with Gaussian process [86, 71], PLS [64], SVM [11, 78], or outlier detection with one-class SVM [66]. We refer the reader to various textbooks and survey articles for more details on these algorithms [69, 67, 72].

In summary, the definition of a positive definite kernel for certain types of data defines explicitly or implicitly

a mapping of the data to a vector space, possibly of high or infinite dimension. Yet a variety of data processing and analysis algorithm can be performed in this feature space thanks to the kernel trick, without the need to compute and store the vector representing the objects. In the next two sections, we review some recent work focusing on the definition of valid kernels for proteins and small molecules, respectively, to illustrate the possibilities offered by kernels to define implicitly “biological” and “chemical” spaces.

3 The biological space: kernels for proteins

Bioinformatics has historically been one of the first application domain for SVM and kernel methods [69], and has triggered a lot of research focused on the design of valid kernels for non-vectorial object, such as proteins. The simplest representation of a protein is the sequence of amino acids it contains, which mathematically is a string in an alphabet of 20 letters. Alternatively, when available or predicted, one can represent a protein by its 3D structure, which is likely to contain more relevant information related to physical interactions with other proteins or ligands. As summarized in Table 1, three main strategies have been followed to define kernels between proteins: (i) computing an inner product with descriptors explicitly defined, (ii) deriving a kernel from a probabilistic model, and (iii) adapting widely used measures of similarity between biological sequences or 3D structures. We now review in more details these different strategies, starting with kernels defined for amino acid or nucleotide sequences.

Table 1: A typology of kernels for proteins.

Strategy	Input data	Examples
Define a list of descriptors	Sequence	Physico-chemical kernels [84, 87] Spectrum, mismatch kernels [48, 49, 43, 44] Pairwise, motif kernel [52, 3, 51]
	3D Structure	Kernel based on 3D descriptors [18]
Derive a kernel from a generative model	Sequence	Fisher, TOP kernel [31, 76] Mutual information kernels [16] Marginalized kernels [77, 82, 41]
	3D Structure	Random walk kernel[10]
Derive a kernel from a measure of similarity	Sequence	Local alignment kernels [27, 85, 81, 65, 62]
	3D Structure	Structure alignment kernel [59]

The first strategy to make a kernel is to define a set of descriptors to characterize various features of protein

sequences, and to compute the inner products between the resulting vectors to obtain a kernel. As an example, [48] uses as descriptors how many times each sequence of n letters occurs consecutively in the string, for a fixed integer n (a sequence of n contiguous letters is called a n -mer). These descriptors could be relevant to detect homologous proteins, which are likely to contain similar proportions of the various n -mers, or to predict biological properties that depend on short motifs of amino acids, such as structural or binding motifs. Taking for instance $n = 2$, the DNA sequence $x = AATCGCAACT$ is represented by the 16-dimensional vector $\Phi(x) = (2, 1, 0, 1, 1, 0, 1, 1, 0, 1, 0, 0, 1, 0, 0)$, where the numbers are the counts of occurrences of each 2-mer AA, AC, \dots, TG, TT lexicographically ordered. The dimension of $\Phi(x)$ is then 4^n for nucleotide sequences, and 20^n for amino acid sequences, which can be prohibitively large for, e.g., $n = 5$. Fortunately, [48] show that a computational trick allows to compute the kernel between two sequences with a complexity in time and memory linear with respect to the sum of the length of the sequences, independently from the dimension of $\Phi(x)$ [48, 83]. Several variants to this kernel have also been proposed, including kernels based on counts of n -mers appearing with up to a few mismatches in the sequences [50], matching of n -mers with the possibility of gaps or substitution [49], or counts of n -mers derived from a profile instead of a single sequence [43, 44]. Alternatively one can first replace each amino acid by one or several numerical features, such as physico-chemical properties, and then extract features from the resulting variable-length numerical time series using classical signal processing techniques such as Fourier transforms [84] or autocorrelation analysis [87]. The resulting features can be explicitly computed to form a vector, and any valid kernel for vector can then be used. These descriptors are interesting to encode information about the variations of physico-chemical properties along the sequence, e.g., to detect elements of the secondary structure. Finally, another popular approach to design features and therefore kernels for biological sequences is to “project” them onto a fixed dictionary of sequences or motifs, using classical similarity measures, and to use the resulting vector of similarities as the feature vector. For example, [52] represent each sequence by a 10,000-dimensional vector indicating the presence of 10,000 motifs of the BLOCKS database; similarly, [3] use a vector that indicates the presence or absence of about 500,000 motifs in the eMOTIF database, requiring the use of a tree structure to compute efficiently the kernel without explicitly storing the 500,000 features; and [51] represent each sequence by a vector of sequence similarities with a fixed set of sequences. The choice of sequences or motifs to be included in the dictionary is crucial and may be problem dependent, as it allows to extract for example the occurrences

of particular functional or structural motifs in the protein sequences.

A second strategy to design kernels for amino acid sequences has been to derive them from probabilistic models. Indeed, before the interest on string kernels grew, a number of ingenious probabilistic models had been defined to represent biological sequences or families of sequences, including for example Markov and hidden Markov models for protein sequences, or stochastic context-free grammars for RNA sequences [20]. Several authors have therefore explored the possibility to use such models to make kernels, starting with the seminal work of [31] that introduced the *Fisher kernel*. This kernel uses a parametric probabilistic model to explicitly extract features from each sequence. The features for a sequence x are related to the influence of each parameter of the model on the probability of x . The resulting vector of features, known as the Fisher score vector in statistics, has a fixed dimension equal to the number of parameters in the model, and therefore provides a principled way to map sequences of different length to a vector of fixed length. The Fisher kernel was generalized by the Tangent Of Posterior (TOP) kernel [76]. Intuitively, the descriptors encoded in the Fisher and TOP kernel describe how each individual sequence differs from a model supposed to represent an “average” sequence, and the choice of the model and its parameters influence therefore a lot the kernel. A second line of thought to make a kernel out of a parametric probabilistic model is to use the concept of mutual information (MI) kernels [70]. Contrary to the Fisher kernel, MI kernels do not provide an explicit finite-dimensional representation for each sequence. Instead the dimensions of the feature space are indexed by all possible values of the model parameters, and the feature $\Phi_\theta(x)$ extracted from the sequence x for the parameter θ is the probability of x under the model P_θ , i.e., $\Phi_\theta(x) = P(x|\theta)$. The computation of this kernel involves a summation over all parameters, i.e., takes the form:

$$K(x, x') = \int P_\theta(x)P_\theta(x')d\mu(\theta),$$

where $d\mu$ is a prior distribution on the parameter space. Hence for practical applications one must chose probabilistic models that allow the computation of the above integral. This was carried out by [16] who present a family of variable-length Markov models for strings and an algorithm to perform the integral over parameters and models at the same time, resulting in a string kernel with linear complexity in time and memory with respect to the total length of the sequences. There exists an information-theoretic interpretation of mutual information kernels: they quantify how much information is shared between two sequences, in particular if the knowledge

of a sequence can be helpful to compress another one. Finally, a third strategy to derive valid kernels from probabilistic models with latent variables, such as HMM, is to build a *marginalized kernel* [77]. Latent variables in probabilistic models often represent meaningful information, such as the local structure or function of a protein sequence. The basic idea behind marginalized kernel is to first design a kernel over the latent and observed variables, as if the latent ones were observed, and then to take the expectation of the kernel with respect to the conditional distribution of the latent variable given the sequences. As for the MI kernel, this kernel can only be computed for judicious choices of random models. Several beautiful examples of such kernels for various probabilistic models have been worked out, including hidden Markov models for sequences [77, 82] or stochastic context-free grammars for RNA sequences [41].

A third strategy to define a kernel is to go back to the interpretation of kernels as “measure of similarity”, and try to adapt well-known measures of similarities between biological sequences to make valid kernels. This idea was pioneered by [27] introduced the concept of *convolution kernels* for structured objects that can be decomposed into subparts, such as sequences that can be decomposed into subsequences concatenated to each other (see also [81]). Convolution kernels offer the possibility to combine several kernels adapted to each subpart of the sequences into a single kernel for the whole sequence. Besides proving the validity of convolution kernels, [27, 85] give several examples of convolution kernels relevant for biological sequences. This work is extended by [81, 65] where a valid convolution kernel based on the alignment of two sequences is proposed. This kernel, named *local alignment kernel*, is a close relative of the widely used Smith-Waterman local alignment score [73], and gives excellent results on the problem of detecting remote homologs of proteins. This work was later extended to alignment kernels for sequence profiles [62]. This strategy is particularly relevant when the kernel is used by a SVM to predict a property that is conserved across “similar” sequences. In particular, the local alignment score attempts to quantify a measure of evolutionary distance, and the local alignment kernel is therefore particularly adapted to predict biological properties conserved during evolution.

While kernels for sequences, that implicitly map proteins to a feature space through their primary structure, have by far attracted the largest attention so far, several groups have recently attempted to map protein 3D structures through the construction of kernels between structures. [18] explicitly represent each structure by a vector made of carefully chosen features, such as secondary structure content, amino acid propensities, surface properties, etc. Alternatively, [10] use a representation based upon walks defined on a graph of secondary

structural elements, while [59] show that a kernel derived from a measure of structure superpositions is more efficient to relate the structure of a protein to its function.

These kernels for proteins have been widely applied, often in combination with SVM, to various classification tasks in computational biology, including for example the prediction of structural or functional classes [17, 31, 81, 38, 14, 18, 10, 59] or the prediction of the subcellular localization of proteins [29, 58, 56]. The performance reported in these studies are often state-of-the-art, which might be in large part due to the efficacy of algorithms like SVM to estimate classification or regression function. While each kernel for proteins corresponds to a particular embedding of the space of proteins in a vector space, it has been observed that the choice of the kernel, hence of the embedding, can have an important effect on the final performance of the algorithm. For example, in the context of remote protein homology detection, [81] compared different kernels and observed that the local alignment kernel was particularly efficient for this application. Besides the performance criterion, different kernels can have different computational complexities which might become prohibitive if large datasets are to be processed. Hence in practical applications the choice of a particular kernel is often a trade-off between computational consideration and performance.

4 The chemical space: kernels for small molecules

Kernel methods are also increasingly used in chemoinformatics for various analysis, regression or classification tasks with small molecules. As for proteins, the design of kernels for small molecules can follow different strategies. We now review the main recent contributions in this field, as summarized in Table 2

Table 2: A typology of kernels for small molecules

Strategy	Input data	Examples
Use classical fingerprints of molecular descriptors	1D to 4D structure	Tanimoto or inner products between fingerprints [13, 60, 2]
Use an infinite number of descriptors and a computational trick	2D structure	Walk kernels [39, 40, 23, 54] Shortest-path fragment kernel [9] Subtree kernel [61, 55]
	3D structure	Cyclic fragment kernel [28] Pharmacophore kernel [53]
Use a measure of similarity	2D structure	Optimal assignment kernel [22]

The problem of explicitly representing and storing small molecules as finite-dimensional vectors has a long

history in chemoinformatics, and a multitude of molecular descriptors have been proposed [75]. These descriptors include in particular physicochemical properties of the molecules, such as its solubility or logP, descriptors derived from the 2D structure of the molecule, such as fragment counts or structural fingerprints, or descriptors extracted from the 3D structure. All classical vector fingerprint and vector representations of molecules define an explicit “chemical space” where each molecule is represented by a finite-dimensional vector, and these vector representations can obviously be used as such to define kernels between molecules [13, 2].

Alternatively, several groups have investigated different strategies to build implicit chemical spaces by defining kernels between molecules that do not require the explicit computation of vector representations. These attempts were pioneered simultaneously and independently by [39, 40] and [23] who proposed to represent the 2D structure of a molecule by an infinite-dimensional vector of linear fragment counts and showed how SVM can handle this representation with the kernel trick. [54] extended these works by showing how irrelevant fragments can be filtered out and proposing a trick to increase the dimension of the feature space to make the fragments more specific while simultaneously increasing the speed of the kernel computation. [60] also tested several variants of these kernels, and showed in particular that the Tanimoto index, widely used in chemoinformatics, is a valid kernel. [9] investigated the possibility to restrict the fragment counts to shortest-path fragments. Several groups have also tried to extend the substructures extracted from the molecular graphs beyond linear fragments, and therefore to trade some increase in expressiveness against loss in computational efficiency [61]. For example, [28] considers kernels based on cyclic fragments, while [61] suggests to consider tree fragments instead of linear fragments, an idea that was later extended and validated in [55]. Finally, [22] defines a kernel between molecular graphs by scoring an optimal matching between the atoms of two molecules to be compared; this kernel, however, is not a valid one [80].

A few attempts to define kernels based on the 3D structure of molecules have also been proposed. [53] design a kernel focused on the detection of 3D pharmacophores, while [74] considers similarity measures between histograms of pairwise distances between atom classes and [2] use Delaunay tetrahedrization and other techniques from computational geometry to characterize the 3D structures of small molecules and make kernels. Finally, [2] shows how kernels can also handle multiple 3D conformations and demonstrates the relevance of this idea on several QSAR experiments.

Although the construction of valid kernels for molecules is a young discipline, it has witnessed impressive

progresses in just a few years, triggered by potential application in chemoinformatics and drug design. Large avenues that could be relevant for kernel design remain however largely unexplored, such as the modeling of 3D surfaces, their electrostaticity and polarity, or the dynamics of the structures. We expect fast progresses in this field in the coming years.

5 Towards *in silico* chemogenomics with kernels

Chemogenomics intends to screen the chemical universe, or at least subsets of this universe, against the therapeutic target universe [12, 45, 57, 35, 36, 63, 42]. While quantitative structure-activity relationship (QSAR) analysis in chemoinformatics attempt to model the affinity of a family of molecules for a single target protein, chemogenomics can be thought of as a generalization where several targets are screened simultaneously. The screening process for each target can indeed be expected to benefit from the known data for other targets. In particular it is possible to make accurate prediction for targets with few known ligands or even no known ligand if more data is available for targets that are *similar* to the target of interest [42].

The development of *in silico* chemoinformatics also needs tools to go beyond classical QSAR and classification methods. The problem of QSAR can be formulated as the problem of learning a function $f(m)$ to predict the affinity of a molecule m to a given target. The function f can for example be a linear function when m is represented by a vector $\Phi(m)$. In the chemogenomics era, however, multiple targets are considered simultaneously and the problem is rather to infer a function $f(m, p)$ to predict the affinity of a molecule m for a target protein p . In order to apply classical classification or regression tools, a vector representation $\Phi(m, p)$ of the pair (m, p) must therefore be chosen. In other words, instead of simply working with the chemical or biological spaces alone, we must now consider their *product space*, i.e., the set of (m, p) pairs, and design machine learning algorithms in this product space. Several methods, referred to as *target-ligand approaches* in [63], have started to emerge to formally combine the compounds and target universes.

We argue that kernel methods offer an attractive framework for *in silico* chemogenomics, seen as the problem of inferring functions in the product chemical and biological space. Indeed, as surveyed in Sections 3 and 4, a variety of kernels capturing various features or proteins and small molecules have already been defined. Each of these kernels is equivalent to an explicit or implicit definition of the chemical or biological space as a possibly high-dimensional vector space. Now let us suppose that in these spaces, a protein p is represented by a vector

$\Phi_P(p)$ corresponding to a kernel k_P , while a small molecule m is represented by a vector $\Phi_M(m)$ corresponding to a kernel k_M . Perhaps the most natural way to represent the pair (p, m) for the purpose of chemogenomics is to concatenate both vectors $\Phi_P(p)$ and $\Phi_M(m)$ into a single vector $\Phi(m, p)$ [8]. Interestingly, in terms of kernels, it is easy to show that the inner product in the joint space obtained by concatenating the vectors, which is usually called the tensor sum kernel and denoted $k_P \oplus k_M$, is simply the sum of the inner products in both spaces taken separately, and therefore that the kernel for (p, m) is simply the sum of the protein kernel and of the molecule kernel [1]:

$$k_P \oplus k_M((p, m), (p', m')) = k_P(p, p') + k_M(m, m').$$

A drawback of this sum is that if a linear algorithm is run on the concatenation of two vectors, it results in a linear function $f(p, m)$ that decomposes as the sum of a protein-specific function and a molecule-specific function, i.e., $f(p, m) = f_P(p) + f_M(m)$. In particular, if the function $f(p, m)$ is then used to rank molecules m for a fixed protein p , then the ranking only depends on the molecule-specific term $f_M(m)$. A undesired consequence is that the ranking of molecules is then the same for all proteins.

More interestingly, correlations between the protein and the molecule descriptors can be introduced by considering, instead of the simple concatenation of vectors, the tensor product between vectors whose descriptors are all possible products between a protein descriptor and a molecule descriptor. As a result, if a protein is characterized by d_P descriptors and a small molecule by d_M descriptors, then their tensor product is characterized by $d_P \times d_M$ descriptors. In the case of binary descriptors that describe the presence or absence of particular features in the protein and the small molecule, a descriptor in the tensor product will be 1 if and only if the corresponding descriptors in the protein and in the small molecule are both equal to 1, hence this representation allows to encode the simultaneous presence of particular features in both the protein and the ligand. Unfortunately, this rich representation can hardly be computed for classical vectors of descriptors due to the explosion in the size of the tensor product vector: for example, using a vector of molecular descriptors of size 1024 for small molecules, and representing a protein by the vector of counts of all 2-mers of amino acids in its sequence ($d_P = 20 \times 20 = 400$) results in more than 400 thousands descriptors in the tensor product. The kernel formulation has a very strong advantage here: indeed, it can be shown with classical algebra that the kernel associated with a tensor product, which is usually called the tensor product kernel and denoted $k_P \otimes k_M$, is simply the

product of individual kernels [1]:

$$k_P \otimes k_M((p, m), (p', m')) = k_P(p, p')k_M(m, m').$$

This shows that just multiplying any protein kernel with any small molecule kernel results in a valid kernel for protein-ligand pairs, which encodes a rich information about features jointly present in the protein and the molecule.

This useful property has been investigated by a few authors recently. [21] applied the tensor product kernel to two vector of descriptors for proteins and molecules, in order to estimate a model to predict interactions between drug targets and small molecules from a collection of disparate HTS campaigns aimed at screening 24 biological targets. The molecular descriptors ranged from atom frequencies and topological indexes to 3D surface area descriptors, while the protein descriptors encode information about the amino acids in the binding pocket. The study concludes that the tensor product combined with a SVM allows to improve the prediction accuracy over models trained for each target independently from each other. [34] followed a similar approach in the context of epitope prediction for multiple MHC-I alleles, *i.e.* interaction between short peptides and MHC-I molecules. In that case the candidate epitopes (short peptides of 9 amino acids) were described by the amino acids present at each position, while different kernels for MHC-I molecules were tested, including kernels based on the primary sequence of the proteins, and kernels based on the presence of particular amino acids at key residues of the interaction pocket. It was observed that simultaneously learning predictors for the binding of short peptides to different alleles of the MHC-I, using a SVM in the product space defined by a tensor product kernel, resulted in significant improvement over the learning of predictors for different alleles independently. More recently, some attempts to systematically apply this strategy to large classes of therapeutic targets, including enzymes, G-protein coupled receptors (GPCR), and ion channels have been published [33, 32]. Combining again various descriptors for small molecules and protein targets with the tensor product kernel, the authors demonstrate that very accurate models for target-ligand interactions can be inferred by sharing data across targets. In particular, they show that their model is able to infer correct ligands even for targets with no known ligand used during the training of SVM. This suggests that this ligand-based procedure could in principle contribute to the prediction of ligands for orphan drug targets, which is of tremendous interest for the pharmaceutical industry.

6 Conclusion

Kernel methods have become prominent in both bioinformatics and chemoinformatics. In particular much effort has been spent on the construction of various kernel functions for proteins and small molecules, which capture various structural, evolutionary or physico-chemical properties of these molecules. Kernel methods provide state-of-the-art performance on many protein or small molecule classification or regression tasks, and offer a promising framework for the new field of *in silico* chemogenomics. In particular various operations on kernel functions allow the principled construction of kernels for target-ligand pairs, resulting in the systematic and implicit construction of a combined target-ligand space amenable to various data analysis by kernel methods. Given the large choice of existing protein and molecule kernels, operations on kernels, and kernel methods, we believe the burgeoning field of kernel methods for *in silico* chemogenomics offers a lot of opportunities to accompany the current *in vitro* and *in vivo* chemogenomics wave.

The first attempts to apply these state-of-the-art machine learning algorithms to *in silico* chemogenomics are very encouraging, although it is still a bit early to assess their real impact on the drug discovery process. It is likely that new issues related to the particular data manipulated will arise, such as the fact that the set of known interactions involves only a very small number of molecules and proteins compared to the huge size of the chemical and biological space, and that the negative results are usually not recorded in HTS campaigns although they may be useful to train *in silico* chemogenomics models. These and other challenges now need to be clearly identified and taken into account for the development and improvement of *in silico* models.

References

- [1] N. Aronszajn. Theory of reproducing kernels. *Trans. Am. Math. Soc.*, 68:337 – 404, 1950.
- [2] C.-A. Azencott, A. Ksikes, S. J. Swamidass, J. H. Chen, L. Ralaivola, and P. Baldi. One- to four-dimensional kernels for virtual screening and the prediction of physical, chemical, and biological properties. *J Chem Inf Model*, 47(3):965–974, 2007.
- [3] A. Ben-Hur and D. Brutlag. Remote homology detection: a motif based approach. *Bioinformatics*, 19(Suppl. 1):i26–i33, 2003.

- [4] C. Berg, J. P. R. Christensen, and P. Ressel. *Harmonic analysis on semigroups*. Springer-Verlag, New York, 1984.
- [5] A. Berlinet and C. Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer, 2003.
- [6] C.M. Bishop. *Pattern recognition and machine learning*. Springer, 2006.
- [7] J. F. Blake. Chemoinformatics - predicting the physicochemical properties of 'drug-like' molecules. *Curr Opin Biotechnol*, 11(1):104–107, Feb 2000.
- [8] Joel R Bock and David A Gough. Virtual screen for ligands of orphan g protein-coupled receptors. *J Chem Inf Model*, 45(5):1402–1414, 2005.
- [9] K. M. Borgwardt and H.-P. Kriegel. Shortest-path kernels on graphs. *Int. Conf on Data Mining*, 0:74–81, 2005.
- [10] K.M. Borgwardt, C.S. Ong, S. Schönauer, S.V.N. Vishwanathan, A.J. Smola, and H.-P. Kriegel. Protein function prediction via graph kernels. *Bioinformatics*, 21(Suppl. 1):i47–i56, Jun 2005.
- [11] B. E. Boser, I. M. Guyon, and V. N. Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the 5th annual ACM workshop on Computational Learning Theory*, pages 144–152. ACM Press, 1992.
- [12] M. Bredel and E. Jacoby. Chemogenomics: an emerging strategy for rapid target and drug discovery. *Nat. Rev. Genet.*, 5(4):262–275, Apr 2004.
- [13] R. Burbidge, M. Trotter, B. Buxton, and S. Holden. Drug design by machine learning: support vector machines for pharmaceutical data analysis. *Comput. Chem.*, 26(1):4–15, December 2001.
- [14] C.Z. Cai, W.L. Wang, L.Z. Sun, and Y.Z. Chen. Protein function classification via support vector machine approach. *Math. Biosci.*, 185(2):111–122, 2003.
- [15] N. Cristianini and J. Shawe-Taylor. *An introduction to Support Vector Machines and other kernel-based learning methods*. Cambridge University Press, 2000.
- [16] M. Cuturi and J.-P. Vert. The context-tree kernel for strings. *Neural Network.*, 18(4):1111–1123, 2005.
- [17] C.H.Q. Ding and I. Dubchak. Multi-class protein fold recognition using support vector machines and neural networks. *Bioinformatics*, 17:349–358, 2001.

- [18] P.D. Dobson and A.J. Doig. Predicting enzyme class from protein structure without alignments. *J. Mol. Biol.*, 345(1):187–199, Jan 2005.
- [19] R. O. Duda, P. E. Hart, and D. G. Stork. *Pattern Classification*. Wiley-Interscience, 2001.
- [20] R. Durbin, S. Eddy, A. Krogh, and G. Mitchison. *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge University Press, 1998.
- [21] Dumitru Erhan, Pierre-Jean L’heureux, Shi Yi Yue, and Yoshua Bengio. Collaborative filtering on a family of biological targets. *J Chem Inf Model*, 46(2):626–635, 2006.
- [22] H. Fröhlich, J. K. Wegner, F. Sieker, and A. Zell. Optimal assignment kernels for attributed molecular graphs. In *Proceedings of the 22nd international conference on Machine learning*, pages 225 – 232, New York, NY, USA, 2005. ACM Press.
- [23] T. Gärtner, P. Flach, and S. Wrobel. On graph kernels: hardness results and efficient alternatives. In B. Schölkopf and M. Warmuth, editors, *Proceedings of the Sixteenth Annual Conference on Computational Learning Theory and the Seventh Annual Workshop on Kernel Machines*, volume 2777 of *Lecture Notes in Computer Science*, pages 129–143, Heidelberg, 2003.
- [24] T. Gärtner, J.W. Lloyd, and P.A. Flach. Kernels and distances for structured data. *Mach. Learn.*, 57(3):205–232, 2004.
- [25] J. Gasteiger and T. Engel, editors. *Cheminformatics : a Textbook*. Wiley, 2003.
- [26] T. Hastie, R. Tibshirani, and J. Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer, 2001.
- [27] D. Haussler. Convolution Kernels on Discrete Structures. Technical Report UCSC-CRL-99-10, UC Santa Cruz, 1999.
- [28] T. Horváth, T. Gärtner, and S. Wrobel. Cyclic pattern kernels for predictive graph mining. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 158–167, New York, NY, USA, 2004. ACM Press.
- [29] S. Hua and Z. Sun. Support vector machine approach for protein subcellular localization prediction. *Bioinformatics*, 17(8):721–728, 2001.

- [30] O. Ivanciuc. Applications of support vector machines in chemistry. In K. B. Lipkowitz and T. R. Cundari, editors, *Reviews in Computational Chemistry*, volume 23, pages 291–400, Weinheim, 2007. Wiley-VCH.
- [31] T. Jaakkola, M. Diekhans, and D. Haussler. A Discriminative Framework for Detecting Remote Protein Homologies. *J. Comput. Biol.*, 7(1,2):95–114, 2000.
- [32] L. Jacob, B. Hoffmann, B. Stoven, and J.-P. Vert. Virtual screening of GPCRs: an *in silico* chemogenomics approach. Technical Report 0801.4301, Arxiv, 2008.
- [33] L. Jacob and J.-P. Vert. Kernel methods for in silico chemogenomics. Technical Report 0709.3931v1, arXiv, 2007.
- [34] L. Jacob and J.-P. Vert. Efficient peptide-MHC-I binding prediction for alleles with few known binders. *Bioinformatics*, 2008. To appear.
- [35] E. Jacoby, editor. *Chemogenomics: Knowledge-based Approaches to Drug Discovery*. Imperial College Press, 2006.
- [36] S. E. Jaroch and H. Weinmann, editors. *Chemical Genomics: Small Molecule Probes to Study Cellular Function*. Ernst Schering Research Foundation Workshop. Springer, 2006.
- [37] Svava Osk Jnsdttir, Flemming Steen Jrgensen, and Sren Brunak. Prediction methods and databases within chemoinformatics: emphasis on drugs and drug candidates. *Bioinformatics*, 21(10):2145–2160, May 2005.
- [38] R. Karchin, K. Karplus, and D. Haussler. Classifying G-protein coupled receptors with support vector machines. *Bioinformatics*, 18:147–159, 2002.
- [39] H. Kashima, K. Tsuda, and A. Inokuchi. Marginalized Kernels between Labeled Graphs. In T. Faucett and N. Mishra, editors, *Proceedings of the Twentieth International Conference on Machine Learning*, pages 321–328. AAAI Press, 2003.
- [40] H. Kashima, K. Tsuda, and A. Inokuchi. Kernels for graphs. In B. Schölkopf, K. Tsuda, and J.P. Vert, editors, *Kernel Methods in Computational Biology*, pages 155–170. MIT Press, 2004.
- [41] T. Kin, K. Tsuda, and K. Asai. Marginalized kernels for RNA sequence data analysis. In R.H. Lathrop, K. Nakai, S. Miyano, T. Takagi, and M. Kanehisa, editors, *Genome Informatics 2002*, pages 112–122. Universal Academic Press, 2002.

- [42] T. Klabunde. Chemogenomic approaches to drug discovery: similar receptors bind similar ligands. *Br J Pharmacol*, 152:5–7, May 2007.
- [43] R. Kuang, E. Ie, K. Wang, K. Wang, M. Siddiqi, Y. Freund, and C. Leslie. Profile-based string kernels for remote homology detection and motif extraction. *Proc IEEE Comput Syst Bioinform Conf*, pages 152–160, 2004.
- [44] R. Kuang, E. Ie, K. Wang, K. Wang, M. Siddiqi, Y. Freund, and C. Leslie. Profile-based string kernels for remote homology detection and motif extraction. *J Bioinform Comput Biol*, 3(3):527–550, Jun 2005.
- [45] H. Kubinyi, G. Müller, R. Mannhold, and G. Folkers, editors. *Chemogenomics in Drug Discovery: A Medicinal Chemistry Perspective*. Methods and Principles in Medicinal Chemistry. Wiley-VCH, 2004.
- [46] G. R. G. Lanckriet, T. De Bie, N. Cristianini, M. I. Jordan, and W. S. Noble. A statistical framework for genomic data fusion. *Bioinformatics*, 20(16):2626–2635, 2004.
- [47] M. Lapinsh, P. Prusis, A. Gutcaits, T Lundstedt, and J. E. S. Wikberg. Development of proteo-chemometrics: A novel technology of use for analysis of drug-receptor interactions. *Biochem. Biophys. Acta*, 1525:180–190, 2001.
- [48] C. Leslie, E. Eskin, and W.S. Noble. The spectrum kernel: a string kernel for SVM protein classification. In Russ B. Altman, A. Keith Dunker, Lawrence Hunter, Kevin Lauerdale, and Teri E. Klein, editors, *Proceedings of the Pacific Symposium on Biocomputing 2002*, pages 564–575. World Scientific, 2002.
- [49] C. Leslie and R. Kuang. Fast string kernels using inexact matching for protein sequences. *J. Mach. Learn. Res.*, 5:1435–1455, 2004.
- [50] C. S. Leslie, E. Eskin, A. Cohen, J. Weston, and W. S. Noble. Mismatch string kernels for discriminative protein classification. *Bioinformatics*, 20(4):467–476, 2004.
- [51] L. Liao and W.S. Noble. Combining Pairwise Sequence Similarity and Support Vector Machines for Detecting Remote Protein Evolutionary and Structural Relationships. *J. Comput. Biol.*, 10(6):857–868, 2003.
- [52] B. Logan, P. Moreno, B. Suzek, Z. Weng, and S. Kasif. A Study of Remote Homology Detection. Technical Report CRL 2001/05, Compaq Cambridge Research laboratory, June 2001.

- [53] P. Mahé, L. Ralaivola, V. Stoven, and J.-P. Vert. The pharmacophore kernel for virtual screening with support vector machines. *J Chem Inf Model*, 46(5):2003–2014, 2006.
- [54] P. Mahé, N. Ueda, T. Akutsu, J.-L. Perret, and J.-P. Vert. Graph kernels for molecular structure-activity relationship analysis with support vector machines. *J Chem Inf Model*, 45(4):939–51, 2005.
- [55] P. Mahé and J.-P. Vert. Graph kernels based on tree patterns for molecules. Technical Report ccscd-00095488, HAL, September 2006.
- [56] A. Matsuda, J.-P. Vert, H. Saigo, N. Ueda, H. Toh, and T. Akutsu. A novel representation of protein sequences for prediction of subcellular location using support vector machines. *Protein Sci.*, 14(11):2804–2813, 2005.
- [57] Jordi Mestres. Computational chemogenomics approaches to systematic knowledge-based drug discovery. *Curr Opin Drug Discov Devel*, 7(3):304–313, May 2004.
- [58] K.-J. Park and M. Kanehisa. Prediction of protein subcellular locations by support vector machines using compositions of amino acids and amino acid pairs. *Bioinformatics*, 19(13):1656–1663, 2003.
- [59] Jian Qiu, Martial Hue, Asa Ben-Hur, Jean-Philippe Vert, and William Stafford Noble. A structural alignment kernel for protein structures. *Bioinformatics*, 23(9):1090–1098, May 2007.
- [60] L. Ralaivola, S. J. Swamidass, H. Saigo, and P. Baldi. Graph kernels for chemical informatics. *Neural Netw.*, 18(8):1093–1110, Sep 2005.
- [61] J. Ramon and T. Gärtner. Expressivity versus efficiency of graph kernels. In T. Washio and L. De Raedt, editors, *Proceedings of the First International Workshop on Mining Graphs, Trees and Sequences*, pages 65–74, 2003.
- [62] H. Rangwala and G. Karypis. Profile-based direct kernels for remote homology detection and fold recognition. *Bioinformatics*, 21(23):4239–4247, Dec 2005.
- [63] D. Rognan. Chemogenomic approaches to rational drug design. *Br J Pharmacol*, 152:38–52, May 2007.
- [64] R. Rosipal and L. J. Trejo. Kernel partial least squares regression in reproducing kernel hilbert space. *J. Mach. Learn. Res.*, 2:97–123, 2001.
- [65] H. Saigo, J.-P. Vert, N. Ueda, and T. Akutsu. Protein homology detection using string alignment kernels. *Bioinformatics*, 20(11):1682–1689, 2004.

- [66] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the Support of a High-Dimensional Distribution. *Neural Comput.*, 13:1443–1471, 2001.
- [67] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, 2002.
- [68] B. Schölkopf, A.J. Smola, and K.-R. Müller. Kernel principal component analysis. In B. Schölkopf, C. Burges, and A. Smola, editors, *Advances in Kernel Methods - Support Vector Learning*, pages 327–352. MIT Press, 1999.
- [69] B. Schölkopf, K. Tsuda, and J.-P. Vert. *Kernel Methods in Computational Biology*. MIT Press, 2004.
- [70] M. Seeger. Covariance Kernels from Bayesian Generative Models. In *Adv. Neural Inform. Process. Syst.*, volume 14, pages 905–912, 2002.
- [71] Matthias Seeger. Gaussian processes for machine learning. *Int J Neural Syst*, 14(2):69–106, Apr 2004.
- [72] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- [73] T. Smith and M. Waterman. Identification of common molecular subsequences. *J. Mol. Biol.*, 147:195–197, 1981.
- [74] S. J. Swamidass, J. Chen, J. Bruand, P. Phung, L. Ralaivola, and P. Baldi. Kernels for small molecules and the prediction of mutagenicity, toxicity and anti-cancer activity. *Bioinformatics*, 21(Suppl. 1):i359–i368, Jun 2005.
- [75] R. Todeschini and V. Consonni. *Handbook of Molecular Descriptors*. Wiley-VCH, 2002.
- [76] K. Tsuda, M. Kawanabe, G. Rätsch, S. Sonnenburg, and K.-R. Müller. A new discriminative kernel from probabilistic models. *Neural Computation*, 14(10):2397–2414, 2002.
- [77] K. Tsuda, T. Kin, and K. Asai. Marginalized Kernels for Biological Sequences. *Bioinformatics*, 18:S268–S275, 2002.
- [78] V. N. Vapnik. *Statistical Learning Theory*. Wiley, New-York, 1998.
- [79] J.-P. Vert. Kernel methods in genomics and computational biology. In G. Camps-Valls, J.-L. Rojo-Alvarez, and M. Martinez-Ramon, editors, *Kernel Methods in Bioengineering, Signal and Image Processing*. IDEA Group, 2007.

- [80] J.-P. Vert. The optimal assignment kernel is not positive definite. Technical Report 0801.4061, Arxiv, 2008.
- [81] J.-P. Vert, H. Saigo, and T. Akutsu. Local alignment kernels for biological sequences. In B. Schölkopf, K. Tsuda, and J.P. Vert, editors, *Kernel Methods in Computational Biology*, pages 131–154. MIT Press, 2004.
- [82] J.-P. Vert, R. Thurman, and W. S. Noble. Kernels for gene regulatory regions. In Y. Weiss, B. Schölkopf, and J. Platt, editors, *Adv. Neural. Inform. Process Syst.*, volume 18, pages 1401–1408, Cambridge, MA, 2006. MIT Press.
- [83] S. V. N. Vishwanathan and A. J. Smola. Fast kernels for string and tree matching. In B. Schölkopf, K. Tsuda, and J.-P. Vert, editors, *Kernel methods in computational biology*, pages 113–130. MIT Press, 2004.
- [84] M. Wang, J. Yang, G.-P. Liu, Z.-J. Xu, and K.-C. Chou. Weighted-support vector machines for predicting membrane protein types based on pseudo-amino acid composition. *Protein Eng. Des. Sel.*, 17(6):509–516, 2004.
- [85] C. Watkins. Dynamic alignment kernels. In A.J. Smola, P.L. Bartlett, B. Schölkopf, and D. Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 39–50. MIT Press, Cambridge, MA, 2000.
- [86] C.K.I. Williams. Prediction with Gaussian Processes: From Linear Regression to Linear Prediction and Beyond. In M.I. Jordan, editor, *Learning and Inference in Graphical Models*. Kluwer Academic Press, 1998.
- [87] S.-W. Zhang, Q. Pan, H.-C. Zhang, Y-L. Zhang, and H.-Y. Wang. Classification of protein quaternary structure with support vector machine. *Bioinformatics*, 19(18):2390–2396, 2003.