



HAL
open science

The pharmacophore kernel for virtual screening with support vector machines

Pierre Mahé, Liva Ralaivola, Véronique Stoven, Jean-Philippe Vert

► **To cite this version:**

Pierre Mahé, Liva Ralaivola, Véronique Stoven, Jean-Philippe Vert. The pharmacophore kernel for virtual screening with support vector machines. 2006. hal-00020066

HAL Id: hal-00020066

<https://hal.science/hal-00020066>

Preprint submitted on 3 Mar 2006

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The pharmacophore kernel for virtual screening with support vector machines

Pierre Mahé

Center for Computational Biology
Ecole des Mines de Paris
Pierre.Mahe@ensmp.fr

Liva Ralaivola

Laboratoire d'Informatique Fondamentale
University Provence/Aix-Marseille 1
Liva.Ralaivola@cmi.univ-mrs.fr

Véronique Stoven

Center for Computational Biology
Ecole des Mines de Paris
Veronique.Stoven@ensmp.fr

Jean-Philippe Vert

Center for Computational Biology
Ecole des Mines de Paris
Jean-Philippe.Vert@ensmp.fr

March 3, 2006

Abstract

We introduce a family of positive definite kernels specifically optimized for the manipulation of 3D structures of molecules with kernel methods. The kernels are based on the comparison of the three-points pharmacophores present in the 3D structures of molecules, a set of molecular features known to be particularly relevant for virtual screening applications. We present a computationally demanding exact implementation of these kernels, as well as fast approximations related to the classical fingerprint-based approaches. Experimental results suggest that this new approach outperforms state-of-the-art algorithms based on the 2D structure of molecules for the detection of inhibitors of several drug targets.

1 Introduction

Virtual screening refers to the process of inferring biological properties of molecules *in silico*, and plays an increasingly important role at the early stages of the drug discovery process to select candidate molecules with promising drug-likeness, including good toxicity and pharmacokinetics properties, as well as the potential to bind and inhibit a target protein of interest (1). In this context, structure-activity relationship (SAR) analysis is commonly used to build predictive models for the property of interest from a description of the molecules, using statistical procedures to build these models from the analysis of molecules with known properties (2).

It is widely accepted that several drug-like properties can be efficiently deduced from the 2D structure of the molecule, that is, the description of a molecule as a set of atoms and their covalent bonds. For example, Lipinski's "rule of five" remains a widely used standard for the prediction of intestinal absorption (3), and the prediction of mutagenicity from 2D molecular fragments is an accurate state-of-the-art approach (4). In the case of target binding prediction, however, the molecular mechanisms responsible for the binding are known to depend on a precise 3D complementarity between the drug and the target, from both the steric and electrostatic perspectives (5). For this reason, there has been a long history of research on the prediction of these interactions from the 3D representation of molecules, that is, their spatial conformation in the 3D space. If the 3D structure of the target is known, the strength of the interaction can be directly evaluated by docking techniques, that quantify the complementarity of the molecule to the target in terms of energy (6). In the general case where the 3D structure of the target is unknown, however, the docking approach is not possible anymore and the modeler must resort to creating a predictive model from available data, typically a pool of molecules with known affinity to the target; this approach is usually referred to as the *ligand-based* approach to virtual screening.

Most approaches to ligand-based virtual screening require to represent and compare 3D structures of molecules. The comparison of 3D structures can for example rely on optimal alignments in the 3D space (7), or on the comparison




Figure 1: The kernel trick: instead of looking for, e.g., a separating hyperplane directly in the input space \mathcal{X} , training patterns (white and black disks) are mapped into a *feature space* \mathcal{H} through a function ϕ in which a hyperplane is computed; this hyperplane might correspond to a complex surface in the input space. Using a proper positive definite kernel k when carrying out the computations to derive the separating hyperplane is equivalent to directly working with the images of the training samples by some mapping ϕ from \mathcal{X} to some feature space \mathcal{H} ; the existence of $\phi : \mathcal{X} \rightarrow \mathcal{H}$ is guaranteed by (3).

of features extracted from the structures (8). Features of particular importance in this context are subsets of two to four atoms together with their relative spatial organization, also called *pharmacophores*. Discovering pharmacophores common to a set of known inhibitors to a drug target can be a powerful approach to the screening of other candidate molecules containing the pharmacophores, as well as a first step towards the understanding of the biological phenomenon involved (9; 10). Alternatively, pharmacophore fingerprints, that is, bitstrings representing a molecule by the pharmacophores it contains, has emerged as a potential approach to apply statistical learning methods for SAR, although sometimes with mixed results (11; 12; 13).

We focus in this paper on an extension of the fingerprint representation of molecules for building SAR models with support vector machines (SVM). SVM is an algorithm for learning a classification or regression rule from labeled examples (14; 15), that has recently been subject to much investigations for SAR applications in chemoinformatics (16; 17; 18; 19). Although SVM can be trained from a vector or bitstring representation of molecules, they can also take advantage of a mathematical trick to only rely on a measure of similarity between molecules, known as *kernel*. This trick, common to other algorithms called kernel methods (20), was for example used in (21; 18) to build SAR models from a 2D fingerprint of molecules of virtually infinite length. Here we investigate the possibility to use this trick in the context of 3D SAR modeling. We propose a measure of similarity between 3D structures, which we call the *pharmacophore kernel*, based on the comparison of pharmacophores present in the structures. It satisfies the mathematical properties required to be a valid kernel and it therefore allows the use of SVM for model building. This kernel bears some similarity with pharmacophore fingerprint approaches, although it produces more general models. In fact, we show that a fast approximation of this kernel, based on pharmacophore fingerprints, leads to significantly lower performance on a benchmark dataset. We also show that competitive performance can be obtained by a fast fingerprint-based approximation using the Tanimoto coefficient as a valid kernel for SVM (19). The overall good performance of the approach on this benchmark supports its relevance as a potentially effective tool for 3D SAR modeling.

This paper is organized as follows. A light introduction to SVM and kernel methods is provided in Section 2, followed by the definition of the pharmacophore kernel (Section 3). The exact computation of this kernel is presented in Section 4, followed by a discussion about the connection between the pharmacophore kernel and recently introduced graph kernels (Section 5) and the presentation of a fast approximation (Section 6). Experimental results on a benchmark dataset for inhibitor prediction are then presented in Section 7, followed by a short discussion.

2 Support vector machines

In this section we briefly review the basics of support vector machines (14; 15). The interested reader is invited to refer to (22; 20; 23) for further details. In its simplest form SVM is an algorithm to learn a binary classification rule

from a set of labeled examples. More formally, suppose one is given a set of examples with a binary label attached to each example, that is, a set $\mathcal{S} = \{(x_1, y_1), \dots, (x_\ell, y_\ell)\}$ where $(x_i, y_i) \in \mathcal{X} \times \{-1, +1\}$ for $i = 1, \dots, \ell$. Here \mathcal{X} is an inner-product space (e.g. \mathbb{R}^d), equipped with inner product $\langle \cdot, \cdot \rangle$, that represents the space of data to be analyzed, typically molecules represented by d -dimensional fingerprints, and the labels $+1$ and -1 are meant to represent two classes of objects, such as inhibitors or non-inhibitors of a target of interest. The purpose of SVM is to learn from \mathcal{S} a classification function $f : \mathcal{X} \rightarrow \{-1, +1\}$ that can be used to predict the class of new unlabeled examples x as $f(x)$.

In the case of SVM, the classification function is simply of the form $f(x) = \text{sign}(\langle w, x \rangle + b)$, where $\text{sign}(\cdot)$ is the function returning the sign, $+1$ or -1 , of its argument. Geometrically speaking, this means that f outputs a prediction for a pattern x depending upon which side of the hyperplane $\langle w, x \rangle + b = 0$ it falls in. More precisely, SVM learn a separating hyperplane from \mathcal{S} defined by a vector w that is a linear combination of the training vectors $w = \sum_{i=1}^{\ell} \alpha_i x_i$, for some $\alpha_i \in \mathbb{R}, i = 1, \dots, \ell$, obtained by solving a linearly constrained quadratic problem meant to optimize a trade-off between finding a hyperplane that correctly separates all the points, while being as far as possible from each point. The linear classifier f can consequently be rewritten as

$$f(x) = \text{sign} \left(\sum_{i=1}^{\ell} \alpha_i \langle x_i, x \rangle + b \right). \quad (1)$$

However, when dealing with nonlinearly separable problems, such as the one depicted on Figure 1 (left), the set of linear classifiers may not be rich enough to provide a good classification function, no matter what the values of the parameters $w \in \mathcal{X}$ and $b \in \mathbb{R}$ are. The purpose of the *kernel trick* (24; 14), is precisely to overcome this limitation by applying a linear approach to the transformed data $\phi(x_1), \dots, \phi(x_\ell)$ rather than the original data, where ϕ is an embedding from the input space \mathcal{X} to the *feature space* \mathcal{H} , usually, but not necessarily, a high-dimensional space, equipped with dot product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$. Thus, according to (1), the separating function f writes as

$$f(x) = \text{sign} \left(\sum_{i=1}^{\ell} \alpha_i \langle \phi(x_i), \phi(x) \rangle_{\mathcal{H}} + b \right). \quad (2)$$

The key ingredient in the kernel approach is to replace the dot product in \mathcal{H} with a kernel, using the definition of positive definite kernels.

Definition 1 (Positive definite kernel). *Let \mathcal{X} be a nonempty space. Let $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a symmetric function. K is said to be a positive definite kernel if and only if, for all $\ell \in \mathbb{N}$, for all $x_1, \dots, x_\ell \in \mathcal{X}$, the square $\ell \times \ell$ matrix $K = (K(x_i, x_j))_{1 \leq i, j \leq \ell}$ is positive semi-definite, that is, all its eigenvalues are nonnegative.*

For a given set $\mathcal{S}_x = \{x_1, \dots, x_\ell\}$, K is the *Gram matrix* of K with respect to \mathcal{S}_x . A fundamental property of positive definite kernels that underlies the kernel trick is the fact that each such kernel can be represented as an inner product in some space. More precisely, it can be shown (25) that for any positive definite kernel function K , there exists a space \mathcal{H} , equipped with the inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$, and a mapping $\phi : \mathcal{X} \rightarrow \mathcal{H}$ such that:

$$\forall u, v \in \mathcal{X} \quad K(u, v) = \langle \phi(u), \phi(v) \rangle_{\mathcal{H}}. \quad (3)$$

The kernel trick consists in replacing all occurrences of $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ in (2) by a positive definite kernel K such that the corresponding decision function f , for an input pattern x , is given by:

$$f(x) = \text{sign} \left(\sum_{i=1}^{\ell} \alpha_i K(x_i, x) + b \right). \quad (4)$$

For SVM as well as for other kernel methods, the knowledge of the Gram matrix suffices to obtain the coefficients α_i .

For any given positive definite kernel, applying the kernel trick turns out to be equivalent to transforming the input patterns x_1, \dots, x_ℓ into the corresponding vectors $\phi(x_1), \dots, \phi(x_\ell) \in \mathcal{H}$ and to look for hyperplanes in \mathcal{H} , as illustrated in Figure 1 (middle). The decision surface in input space \mathcal{X} corresponding to the selected separating hyperplane in \mathcal{H} might be quite complex (see Figure 1, right).

A noteworthy feature of support vector machines and more generally of kernel methods (23) is that, since ready-to-use libraries to derive separating hyperplanes are available, the only requirement for them to be applied to a specific classification problem is to have at hand a proper kernel function to assess the similarity between patterns of the input space considered. Henceforth, their use actually fit in the framework of classification problems involving structured data such as chemical compounds, provided some kernel function has been derived. The rest of the paper is devoted to the construction and analysis of such a kernel for 3D structures of molecules.




Figure 2: Left: a 3-points pharmacophore made of one hydrogen bond acceptor (topmost sphere) and two aromatic rings, with distances d_1 , d_2 and d_3 between the features. Middle: the molecule of flavone. Right: match between flavone and the pharmacophore.

3 The pharmacophore kernel

A *pharmacophore* is usually defined as a three-dimensional arrangement of atoms - or groups of atoms - responsible for the biological activity of a drug molecule (26). The present work focuses on *three-points pharmacophores*, composed of three atoms whose arrangement therefore forms a triangle in the 3D space (Figure 2). With a slight abuse we refer as pharmacophore below to *any* possible configuration of three atoms or classes of atoms arranged as a triangle and present in a molecule, representing therefore a *putative* configuration responsible for the biological property of interest.

Throughout this paper we represent the 3D structure of a molecule as a set of points in \mathbb{R}^3 . These points correspond to the 3D coordinates of the atoms of the molecule (for a given arbitrary basis of the 3D Euclidean space), and they are labeled with some information related to the atoms. More formally, we define a molecule m as

$$m = \{(x_i, l_i) \in \mathbb{R}^3 \times \mathcal{L}\}_{i=1, \dots, |m|} ,$$

where $|m|$ is the number of atoms that compose the molecule and \mathcal{L} denotes the set of atom labels. The label is meant to contain the relevant information to characterize a pharmacophore based on atoms, such as the type of atom (C, N, O, ...) and its partial charge. The three-points pharmacophores considered in this work correspond to triplets of distinct atoms of the molecules. The set of pharmacophores of the molecule m can therefore be formally defined as:

$$\mathcal{P}(m) = \{(p_1, p_2, p_3) \in m^3, p_1 \neq p_2 \neq p_3\} .$$

More generally, the set of all possible pharmacophores is naturally defined as $\mathcal{P} = (\mathbb{R}^3 \times \mathcal{L})^3$, to ensure the inclusion $\mathcal{P}(m) \subset \mathcal{P}$. We can now define a general family of kernels for molecules based on their pharmacophore content:

Definition 2. For any positive definite kernel for pharmacophores $K_{\mathcal{P}} : \mathcal{P} \times \mathcal{P} \rightarrow \mathbb{R}$, we define a corresponding pharmacophore kernel for any pair of molecules m and m' by :

$$K(m, m') := \sum_{p \in \mathcal{P}(m)} \sum_{p' \in \mathcal{P}(m')} K_{\mathcal{P}}(p, p') , \quad (5)$$

with the convention that $K(m, m') = 0$ if either $\mathcal{P}(m)$ or $\mathcal{P}(m')$ is empty.

The fact that the pharmacophore kernel defined in (5) is a valid positive definite kernel on the set of molecules, as soon as $K_{\mathcal{P}}$ is itself a valid positive definite kernel on the set of pharmacophores, is a classical result (see, e.g., (27, Lemma 1)). The problem of constructing a pharmacophore kernel for molecules therefore boils down to the simpler problem of defining a kernel between pharmacophores. A chemically relevant measure of similarity between pharmacophores should obviously quantify at least two features: first, similar pharmacophores should be made of similar atoms, and second, the atoms should have similar relative positions in the 3D space. It is therefore natural to study kernels for pharmacophores that decompose as follows:

$$K_{\mathcal{P}}(p, p') = K_I(p, p') \times K_S(p, p') , \quad (6)$$

where K_I is a kernel function assessing the similarity between the triplets of basis atoms of the pharmacophores (their so-called *intrinsic* similarity), and K_S is a kernel function introduced to quantify their *spatial* similarity.

We can furthermore investigate intrinsic and spatial kernels that factorize themselves as products of more basic kernels between atoms and pairwise distances, respectively. Triplets of atoms are indeed globally similar if the three corresponding pairs of atoms are simultaneously similar, and triangles are similar if the lengths of their edges are pairwise similar. For any pair of pharmacophores $p = ((x_i, l_i))_{i=1,2,3}$ and $p' = ((x'_i, l'_i))_{i=1,2,3}$, this suggests to define kernels as follows:

$$K_I(p, p') = \prod_{i=1}^3 K_{Feat}(l_i, l'_i), \quad (7)$$

$$K_S(p, p') = \prod_{i=1}^3 K_{Dist}(\|x_i - x_{i+1}\|, \|x'_i - x'_{i+1}\|), \quad (8)$$

where $\|\cdot\|$ denotes the Euclidean distance, the indices i are taken modulo 3, and K_{Feat} and K_{Dist} are kernels functions introduced to compare pairs of labels from \mathcal{L} , and pairs of distances, respectively. It suffices now to define the kernels K_{Feat} on $\mathcal{L} \times \mathcal{L}$ and K_{Dist} on $\mathbb{R} \times \mathbb{R}$ in order to obtain, by (5), (6), (7) and (8), a pharmacophore kernel for molecules. The first one compares the atom labels, while the second compares the distances between atoms in the pharmacophores. Intuitively they define the basic notions of similarity involved in the pharmacophore comparison, which in turns defines the overall similarity between molecules.

The kernel we use for K_{Dist} is the Gaussian radial basis function (RBF) kernel, known to be a safe default choice for SVM working on real numbers or vectors (20):

$$K_{Dist}^{RBF}(x, y) = \exp\left(-\frac{\|x - y\|^2}{2\sigma^2}\right), \quad (9)$$

where $\sigma > 0$ is the bandwidth parameter that will be optimized as part of the training of the classifier (see Section 7.2).

Concerning the kernel K_{Feat} between labels, we investigate several choices. The labels belonging in principle to a finite set of possible labels, e.g., the set of atom types with their charges (C , C^+ , C^- , N, ...), the following *Dirac kernel* is a natural default choice to compare a pair of atom labels $l, l' \in \mathcal{L}$:

$$K_{Feat}^{Dirac}(l, l') = \begin{cases} 1 & \text{if } l = l', \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

Alternatively, it might be relevant for pharmacophore definition to compare atoms not only on the basis of their types and partial charges, but also in terms of other physicochemical parameters such as their size, polarity and electronegativity. Formally, a physicochemical parameter for an atom with label l is a real number $f(l)$. In that case, the Gaussian RBF kernel (9) could be applied directly to the parameter values to compare labels. A practical drawback of this kernel, however, is that it never vanishes. This induces an important computational burden compared to the Dirac kernel (see Section 4). As a result, we prefer to use the related *triangular kernel*:

$$K_{Feat}^{Tri}(l, l') = \begin{cases} \frac{C - \|f(l) - f(l')\|}{C} & \text{if } \|f(l) - f(l')\| \leq C, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

An important difference between the Gaussian RBF and triangular kernels lies in the fact that the triangular kernel has a compact support, which means that it can be equal to 0 for different atoms, resulting in important computational gain. The parameter C , to be optimized during the training phase of the algorithm, represents the range beyond which the kernel vanishes.

Note finally that the Gaussian RBF (9), triangular (11), and Dirac (10) kernels are known to be positive definite (20), and it follows from the closure properties of the family of kernel functions, that the kernel between pharmacophores $K_{\mathcal{P}}$ is valid for any choice of the kernels K_{Dist} and K_{Feat} proposed above.

4 Kernel computation

We are now left with the task of computing the pharmacophore kernel (5) for a particular choice of feature and distance kernels K_{Feat} and K_{Dist} . In this section we provide a simple analytical formula for this computation.

For any pair of molecules $m = \{(x_i, l_i) \in \mathbb{R}^3 \times \mathcal{L}\}_{i=1, \dots, |m|}$ and $m' = \{(x'_i, l'_i) \in \mathbb{R}^3 \times \mathcal{L}\}_{i=1, \dots, |m'|}$, let us define a square matrix M of size $n = |m| \times |m'|$, whose dimensions are indexed by the Cartesian product of m and m' . In other words, to each index $i \in [1, n]$ corresponds a unique couple of indices $(i_1, i_2) \in [1, |m|] \times [1, |m'|]$, and to each dimension of the matrix M corresponds a distinct pair of points taken from the molecules m and m' . Denoting by $\mathbf{1}(\cdot)$ the indicator function equal to one if its argument is true, zero otherwise, the entries of M are defined by:

$$\begin{aligned} M[i, j] &= M[(i_1, i_2), (j_1, j_2)] \\ &= K_{Feat}(l_{i_1}, l'_{i_2}) \times K_{Dist}(\|x_{i_1} - x_{j_1}\|, \|x'_{i_2} - x'_{j_2}\|) \times \mathbf{1}(i_1 \neq j_1) \times \mathbf{1}(i_2 \neq j_2). \end{aligned} \quad (12)$$

The value of the pharmacophore kernel between m and m' can now be deduced from the matrix M by the following result:

Proposition 1. *The pharmacophore kernel (5) between a pair of molecules m and m' is equal to:*

$$K(m, m') = \text{trace}(M^3),$$

where M is the square matrix of dimension $|m| \times |m'|$ constructed from m and m' by (12).

Proof. Developing the matrix products involved in the expression of M^3 we get

$$\text{trace}(M^3) = \sum_{i, j, k=1}^n M[i, j]M[j, k]M[k, i],$$

where $n = |m| \times |m'|$ is the size of M . Using the fact that the indices of M ranges over the Cartesian product of the set of indices $[1, |m|]$ and $[1, |m'|]$, we can rewrite this expression as :

$$\text{trace}(M^3) = \sum_{i_1, j_1, k_1=1}^{|m|} \sum_{i_2, j_2, k_2=1}^{|m'|} M[(i_1, i_2), (j_1, j_2)]M[(j_1, j_2), (k_1, k_2)]M[(k_1, k_2), (i_1, i_2)].$$

Substituting with the definition of M given in (12), we obtain :

$$\begin{aligned} \text{trace}(M^3) &= \sum_{i_1, j_1, k_1=1}^{|m|} \sum_{i_2, j_2, k_2=1}^{|m'|} \mathbf{1}(i_1 \neq j_1) \mathbf{1}(j_1 \neq k_1) \mathbf{1}(k_1 \neq i_1) \times \mathbf{1}(i_2 \neq j_2) \mathbf{1}(j_2 \neq k_2) \mathbf{1}(k_2 \neq i_2) \\ &\quad \times K_{Feat}(l_{i_1}, l'_{i_2}) \times K_{Dist}(\|x_{j_1} - x_{i_1}\|, \|x'_{j_2} - x'_{i_2}\|) \\ &\quad \times K_{Feat}(l_{j_1}, l'_{j_2}) \times K_{Dist}(\|x_{k_1} - x_{j_1}\|, \|x'_{k_2} - x'_{j_2}\|) \\ &\quad \times K_{Feat}(l_{k_1}, l'_{k_2}) \times K_{Dist}(\|x_{i_1} - x_{k_1}\|, \|x'_{i_2} - x'_{k_2}\|) \\ &= \sum_{i_1, j_1, k_1=1}^{|m|} \sum_{i_2, j_2, k_2=1}^{|m'|} \mathbf{1}(i_1 \neq j_1 \neq k_1) \times \mathbf{1}(i_2 \neq j_2 \neq k_2) \\ &\quad \times K_{\mathcal{P}}(((x_{i_1}, l_{i_1}), (x_{j_1}, l_{j_1}), (x_{k_1}, l_{k_1})), ((x'_{i_2}, l'_{i_2}), (x'_{j_2}, l'_{j_2}), (x'_{k_2}, l'_{k_2}))) \\ &= \sum_{\substack{i_1, j_1, k_1=1, \\ i_1 \neq j_1 \neq k_1}}^{|m|} \sum_{\substack{i_2, j_2, k_2=1, \\ i_2 \neq j_2 \neq k_2}}^{|m'|} K_{\mathcal{P}}(((x_{i_1}, l_{i_1}), (x_{j_1}, l_{j_1}), (x_{k_1}, l_{k_1})), ((x'_{i_2}, l'_{i_2}), (x'_{j_2}, l'_{j_2}), (x'_{k_2}, l'_{k_2}))) \\ &= \sum_{p \in \mathcal{P}(m)} \sum_{p' \in \mathcal{P}(m')} K_{\mathcal{P}}(p, p') \\ &= K(m, m'). \end{aligned}$$

□

If we neglect the cost of the addition and product operations, and let u be that of evaluating the basis kernels K_{Feat} and K_{Dist} , the complexity of the kernel between pharmacophores $K_{\mathcal{P}}$ is $6u$. Since the cardinality of the set of pharmacophore $\mathcal{P}(m)$ of the molecule m is $|m|^3$, the complexity of the direct computation of the pharmacophore kernel given in definition 2 is $(|m| \times |m'|)^3 \times 6u$. On the other hand, the computation given in Proposition 1 is a two step process :

- first, initialization of the matrix M : each of the $(|m| \times |m'|)^2$ entries is initialized by the product of a kernel K_{Feat} with a kernel K_{Dist} , for a complexity of $(|m| \times |m'|)^2 \times 2u$
- second, computation of the trace of M^3 , which has a complexity of $(|m| \times |m'|)^3$

The global complexity of the matrix-based computation of the kernel is therefore $(|m| \times |m'|)^3 + (|m| \times |m'|)^2 \times 2u$, or equivalently $(|m| \times |m'|)^3 \times (1 + 2u/(|m| \times |m'|))$. In comparison with the direct approach, the matrix-based implementation proposed in Proposition 1 reduces the number of basis kernels K_{Dist} and K_{Feat} to be computed and is therefore more efficient.

In any case, the complexity of the pharmacophore kernel computation is therefore $\mathcal{O}((|m| \times |m'|)^3)$. Even for relatively small molecules (of the order of 50 atoms), this complexity becomes in practice a serious issue when the size of the dataset increases to thousands or tens of thousands of molecules. However, we can note from the definition given in (12), that the lines of M corresponding to pairs of points $(x, l) \in m$ and $(x', l') \in m'$ for which $K_{Feat}(l, l') = 0$ are filled with zeros. Based on this consideration, we observe that the cost of computing the kernel can be reduced by limiting the size of the matrix M , according to the following proposition.

Proposition 2. *If we let M_2 be the reduced version of a square matrix M_1 , where the null lines and the corresponding columns are removed, then $\text{trace}(M_2^3) = \text{trace}(M_1^3)$.*

Proof. Let n_1 (resp. n_2) be the size of M_1 (resp. M_2), and define P (resp. N) as the subset of the set of indices $[1, n_1]$ that corresponds to the non-null (resp. null) lines of M_1 . By definition, we have

$$\begin{aligned} \text{trace}(M_1^3) &= \sum_{i=1}^{n_1} M_1^3[i, i] \\ &= \sum_{i,j,k=1}^{n_1} M_1[i, j]M_1[j, k]M_1[k, i]. \end{aligned} \quad (13)$$

Moreover, if $i \in N$, then $M_1[i, j] = 0 \forall j \in [1, n_1]$. As a consequence, the term $M_1[i, j]M_1[j, k]M_1[k, i]$ in the summations over i, j , and k in (13) is zero as soon as at least one index i, j or k is in the set N . It follows that

$$\begin{aligned} \text{trace}(M_1^3) &= \sum_{i,j,k \in P} M_1[i, j]M_1[j, k]M_1[k, i] \\ &= \sum_{i,j,k=1}^{n_2} M_2[i, j]M_2[j, k]M_2[k, i] \\ &= \text{trace}(M_2^3). \end{aligned}$$

□

Proposition 2 implies that the Cartesian product of m and m' involved in the matrix M defined in (12) can be restricted to the pairs of points for which the label kernel K_{Feat} is non-zero. In the case of the Dirac kernel (10) for discrete labels, this boils down to introducing a dimension in M for any pair of atoms having the same label. This result can have important consequences in practice. Consider for example the case where the atoms of the molecules m and m' are uniformly distributed in k classes of atom labels. In this case, the size of the matrix M is equal to $k(|m|/k \times |m'|/k) = |m| \times |m'|/k$. The complexity of the kernel computation is therefore $\mathcal{O}((|m| \times |m'|/k)^3) = \mathcal{O}((1/k^3)(|m| \times |m'|)^3)$. It is therefore reduced by a factor k^3 in comparison with the original implementation. More generally this shows that important gains in memory and computation can be expected when the set of labels is increased. Section 7.3 discusses such a case in more details when the partial charges of atoms are included or not in the labels. This also justifies why the triangular kernel (11), with its compact support, can lead to much faster implementations than the Gaussian RBF kernel (9) when applied to the comparison of physicochemical properties of atoms. Finally, in a similar way, the kernel K_{Dist} to compare distances can be set to a compactly supported kernel, such as the triangular kernel (11). This has the effect of introducing sparsity in the matrix M , allowing the kernel computation to benefit from sparse matrix algorithms. This possibility was not further explored in this work.

5 Relation with graph kernels

In this section we show that the pharmacophore kernel can be seen as an extension of the walk-count graph kernels (28) to the 3D representation of molecules. The walk-count graph kernel is based on the representation of a molecule m as a labeled graph $m = (\mathcal{V}, \mathcal{E})$, defined by a set of vertices \mathcal{V} , a set of edges $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$ connecting pairs of vertices, and a labeling function $l : \mathcal{V} \cup \mathcal{E} \rightarrow \mathcal{A}$, assigning a label $l(x)$ in an alphabet \mathcal{A} to any vertex or edge x . In the case of molecules, the set of vertices \mathcal{V} corresponds to the atoms of the molecule, and the edges of the graph are usually defined as the covalent bonds between the atoms of the molecules (28; 21; 18). In order to extend this 2D representation to a graph structure capturing 3D information, we propose to introduce an edge between any pair of vertices of the graph. Molecules are therefore seen as complete, atom-based graphs. If we now define a walk of length n as a succession of $n + 1$ connected vertices, it is easy to see that there is a one-to-one correspondence between the set of pharmacophores $\mathcal{P}(m)$ of a molecule m , and its set of self-returning walks of length-three, that we call $\mathcal{W}_3(m)$. We can therefore write the pharmacophore kernel (5) as a walk-based graph kernel :

$$K(m, m') = \sum_{p \in \mathcal{P}(m)} \sum_{p' \in \mathcal{P}(m')} K_{\mathcal{P}}(p, p') = \sum_{w \in \mathcal{W}_3(m)} \sum_{w' \in \mathcal{W}_3(m')} K_{Walk}(w, w'),$$

where $K_{Walk}(w, w') = K_{\mathcal{P}}(p, p')$ for the pair of walks (w, w') corresponding to the pair of pharmacophores (p, p') . More precisely, consider a pair of pharmacophores $p = ((x_i, l_i))_{i=1,2,3}$ and $p' = ((x'_i, l'_i))_{i=1,2,3}$, and a corresponding pair of walks $w = (w_1, w_2, w_3, w_1)$ and $w' = (w'_1, w'_2, w'_3, w'_1)$. There is a direct equivalence between $K_{\mathcal{P}}$ and K_{Walk} if we choose to label the vertices of the graphs by the atom labels involved in the pharmacophore characterization, and to label the edges by the Euclidian distance between the atoms they connect. Indeed, in this case we can write :

$$\begin{aligned} K_{\mathcal{P}}(p, p') &= \prod_{i=1}^3 K_{Feat}(l_i, l'_i) K_{Dist}(\|x_i - x_{i+1}\|, \|x'_i - x'_{i+1}\|) \\ &= \prod_{i=1}^3 K_{Feat}(l(w_i), l(w'_i)) K_{Dist}(l((w_i, w_{i+1})), l((w'_i, w'_{i+1}))) \\ &= K_{Walk}(w, w') \end{aligned}$$

A striking point of this kernel between walks is that it can be factorized along the edges of the walks :

$$\begin{aligned} K_{Walk}(w, w') &= \prod_{i=1}^3 K_{Feat}(l(w_i), l(w'_i)) K_{Dist}(l((w_i, w_{i+1})), l((w'_i, w'_{i+1}))) \\ &= \prod_{i=1}^3 K_{Step}((w_i, w_{i+1}), (w'_i, w'_{i+1})) \end{aligned} \quad (14)$$

The pharmacophore kernel therefore formulates as a walk-based graph kernel, with a walk kernel factorizing along the edges of the walks. It follows from (21) that it can be computed by the formalism based on product-graphs and powers of the adjacency matrix proposed in (28), if the adjacency matrix of the product-graph is weighted by the walk-step kernels K_{Step} (14). Consequently, the matrix M defined in (12) and upon which is based the kernel computation of Proposition 1, can be seen as a weighted adjacency matrix of a product-graph defined on complete, atom-based, molecular factor graphs.

6 Fast approximation

As an alternative to the costly computation presented in Section 4, we introduce in this section several fast approximations to the pharmacophore kernel based on a discretization of the pharmacophore space.

Our definition of pharmacophores is based on the atoms 3D coordinates, but they can equivalently be characterized by the pairwise distances between atoms. In order to define discrete pharmacophores, we restrict ourselves to discrete sets of atom labels (e.g., the set of atom types), and we discretize the range of distances between atoms into a predefined number of bins. Each distance is then mapped to the index of the bin it falls in, and a discrete pharmacophore is defined by a triplet of atom labels together with a triplet of bin indices. More formally, if the distance range is discretized into

p bins, the set of discrete pharmacophores is a finite set defined as $\mathcal{T}_3 = \mathcal{L}^3 \times [1, p]^3$, where \mathcal{L} is the set of atom labels. Consider the mappings ϕ_0^{3pt} and ϕ_1^{3pt} from the set of molecules to the set of discrete pharmacophores \mathcal{T}_3 , defined for the molecule m as :

- $\phi_0^{3pt}(m) = (\phi_{t,0}(m))_{t \in \mathcal{T}_3}$, where $\phi_{t,0}(m)$ is the number of times the pharmacophore t is found in the molecule m ,
- $\phi_1^{3pt}(m) = (\phi_{t,1}(m))_{t \in \mathcal{T}_3}$, where $\phi_{t,1}(m)$ equals 1 if the pharmacophore t is found in the molecule m , and 0 otherwise.

With these new definitions at hand we propose three discrete kernels,

Definition 3 (Three-points spectrum kernel). For a pair of molecules m and m' , we define the three-points spectrum kernel K_{Spec}^{3pt} as

$$K_{Spec}^{3pt}(m, m') = \langle \phi_0^{3pt}(m), \phi_0^{3pt}(m') \rangle = \sum_{t \in \mathcal{T}_3} \phi_{t,0}(m) \phi_{t,0}(m') . \quad (15)$$

Note that if we define the mapping $d : \mathcal{P} \mapsto \mathcal{T}_3$, such that $d(p)$ is the discretized version of the pharmacophore $p \in \mathcal{P}$, we can explicitly write the three-points spectrum kernel as a particular pharmacophore kernel (5):

$$K_{Spec}^{3pt}(m, m') = \sum_{p \in \mathcal{P}(m)} \sum_{p' \in \mathcal{P}(m')} \mathbf{1}(d(p) = d(p')) .$$

This equation shows that this is a crude pharmacophore kernel, based on a kernel for pharmacophores that simply checks if two given pharmacophores have identical discretized versions or not.

Definition 4 (Three-points binary kernel). For a pair of molecules m and m' , we define the three-points binary kernel K_{Bin}^{3pt} as

$$K_{Bin}^{3pt}(m, m') = \langle \phi_1^{3pt}(m), \phi_1^{3pt}(m') \rangle = \sum_{t \in \mathcal{T}_3} \phi_{t,1}(m) \phi_{t,1}(m') . \quad (16)$$

Definition 5 (Three-points Tanimoto kernel). For a pair of molecules m and m' , we define the three-points Tanimoto kernel K_{Tani}^{3pt} as

$$K_{Tani}^{3pt}(m, m') = \frac{K_{Bin}^{3pt}(m, m')}{K_{Bin}^{3pt}(m, m) + K_{Bin}^{3pt}(m', m') - K_{Bin}^{3pt}(m, m')} . \quad (17)$$

Note that the mapping $\phi_1^{3pt}(m)$ corresponds to a classical *pharmacophore fingerprint* representation of the molecule m , where the bitstring is indexed by the pharmacophores of \mathcal{T}_3 . As a result, the three-points Tanimoto kernel is the equivalent of the Tanimoto coefficient for pharmacophore fingerprints, and constitutes a standard pharmacophore-based similarity measure (29; 30; 12). Note that the dimensionality of the feature spaces associated to these kernels corresponds to the cardinality of \mathcal{T}_3 , which is by definition $(np)^3$ for a label set \mathcal{L} of cardinality n and p distance bins.

Finally, we consider additional ‘‘two-points pharmacophore’’ versions of the kernels (15), (16) and (17), based on pairs, instead of triplets, of atoms (19). Letting \mathcal{T}_2 be the set of all possible two-points pharmacophores, that is, pairs of atom types together with the bin index of the edge connecting them, and $\phi_0^{2pt}(m) = (\phi_{t,0}(m))_{t \in \mathcal{T}_2}$ and $\phi_1^{2pt}(m) = (\phi_{t,1}(m))_{t \in \mathcal{T}_2}$ be the mappings of the molecule m to \mathcal{T}_2 , corresponding to $\phi_0^{3pt}(m)$ and $\phi_1^{3pt}(m)$, we define the three following kernels.

Definition 6 (Two-points spectrum kernel). For a pair of molecules m and m' , we define the two-points spectrum kernel K_{Spec}^{2pt} as :

$$K_{Spec}^{2pt}(m, m') = \langle \phi_0^{2pt}(m), \phi_0^{2pt}(m') \rangle = \sum_{t \in \mathcal{T}_2} \phi_{t,0}(m) \phi_{t,0}(m') . \quad (18)$$

Definition 7 (Two-points binary kernel). For a pair of molecules m and m' , we define the two-points binary kernel K_{Bin}^{2pt} as :

$$K_{Bin}^{2pt}(m, m') = \langle \phi_1^{2pt}(m), \phi_1^{2pt}(m') \rangle = \sum_{t \in \mathcal{T}_2} \phi_{t,1}(m) \phi_{t,1}(m') . \quad (19)$$

Definition 8 (Two-points Tanimoto kernel). For a pair of molecules m and m' , we define the two-points Tanimoto kernel K_{Tani}^{2pt} as :

$$K_{Tani}^{2pt}(m, m') = \frac{K_{Bin}^{2pt}(m, m')}{K_{Bin}^{2pt}(m, m) + K_{Bin}^{2pt}(m', m') - K_{Bin}^{2pt}(m, m')}. \quad (20)$$

The following proposition justifies the use of these fast kernels with SVM:

Proposition 3. The kernels (15), (16), (17), (18), (19) and (20) are positive definite.

Proof. Kernels (15), (16), (18), (19) are directly expressed as dot-products, and are consequently positive definite. Kernels (17) and (20) follow the definition of the Tanimoto kernel which is known to positive definite (19). \square

The kernels (15), (16), (17), (18), (19) and (20) can be computed efficiently using an algorithm derived from that used in the implementation of spectrum string kernels (31). We describe this algorithm in the case of the three-points kernels (15), (16) and (17), its extension to the two-points kernels (18), (19) and (20) being straightforward. Following the notation of Section 5, we represent molecules by complete, atom-based labeled graphs, with the difference that the set of atom labels \mathcal{L} defining the vertices labels is considered to be discrete (e.g., the atom types), and the edges are now labeled by the bin index of the corresponding inter-atomic distance. We consider the problem of computing the Gram matrix K associated to such a set of molecular graphs $\{G_i = (\mathcal{V}_{G_i}, \mathcal{E}_{G_i})\}_{i=1, \dots, n}$ for the kernels (15), (16) and (17). The alphabet \mathcal{A} , involved in the graph labeling function l of section 5, is defined as $\mathcal{A} = L_V \cup L_E$, where L_V is the set of vertices labels, corresponding to the set of atom labels \mathcal{L} , and L_E is the set of edges labels, corresponding to the set of distance bins indices.

The algorithm is based on the manipulation of sets of walk pointers within each graph, according to a tree transversal process. If we let n and p be the cardinalities of L_V and L_E respectively, we define a rooted, depth-four tree structuring the space of pharmacophores \mathcal{T}_3 as follows :

- the root node has n sons, corresponding to the n possible vertex labels
- the depth-one and depth-two nodes have $n \times p$ sons, corresponding to the $n \times p$ possible pairs of edge and vertex labels
- the depth-three nodes have p sons, corresponding to the p possible edges labels, a leaf node being implicitly associated the vertex label of its depth-one ancestor.

A path from the root to a leaf node therefore corresponds to a triplet of distinct vertex labels, together with a triplet of distinct edge labels. There is therefore a one-to-one correspondence between the leaf nodes and the pharmacophores of \mathcal{T}_3 . The principle of the algorithm is to recursively transverse this tree until each leaf node (i.e., each potential pharmacophore) is visited. During this process, a set of walk pointers is maintained within each molecule. The pointers are recursively updated such that the pointed walks correspond to the pharmacophores under construction in the tree-transversal process. When reaching a leaf node, the pointed walks correspond to the occurrences of a particular pharmacophore t in the molecules. The mappings $\phi_{t,0}(G_i)$ and $\phi_{t,1}(G_i)$ can therefore be computed for the molecular graphs $\{G_i\}_{i=1, \dots, n}$, and the kernel matrix can be updated.

A pseudo code of the algorithm is given in Algorithms 1, 2, 3 and 4. Algorithm 1 is the main program in charge of the tree-transversal process, and Algorithms 2, 3 and 4 are subroutines, introduced to initialize the walks pointers, extend the pointed walks, and update the Gram matrix respectively. This pseudo-code relies on the abstract types *Pointer* and *Label*, to represent the walk pointers involved in the algorithm, and the generic vertices and edges labels, belonging to L_V and L_N respectively. Formally, a *Pointer* object consists of two graph vertices: a *start* and *current* vertex, representing the first and the current vertices of the pointed walk being extended. To maintain walks pointers within each molecule, we introduce a matrix of pointers $walkPointers = Pointer[][]$: this matrix is initially empty, and during the walk extension process, $walkPointers[i][j]$ corresponds to the j th pointer of the molecular graph G_i . The stopping criterion of the recursion is controlled by an integer variable *depth* corresponding to the depth in the tree during the transversal process. It is initialized to zero and incremented at each recursive call. When *depth* is three, a depth-three node was reached in the tree, which corresponds to pointers on length-two walks in the graphs. In the subsequent recursive step, *depth* is four, and the pointers are updated to ensure that the extended walks correspond to self-returning ones. A leaf node is then reached and the recursion terminates, leading to an update of the Gram matrix. Note however that the recursion is aborted whenever the set of walk pointers becomes empty for all graphs, since we only need to reach the leaf nodes corresponding to the pharmacophores truly present in the set of graphs.

The Gram matrix is updated according to the spectrum (15), binary (16) or Tanimoto (17) definition of the kernel, and we introduce a $n \times n$ Gram matrix K , initialized to zero, together with a binary variable *kernelType* that can take the values 'spectrum', 'binary' or 'Tanimoto'.

Computing the Gram matrix K simply requires a call to the *COMPUTE* function of Algorithm 1 with these initialized data : *COMPUTE*(*walkPointers*, *depth*, K , *kernelType*), for a specified *kernelType*, and where the Pointer array *walkPointers* is empty, *depth* equals zero and the Gram matrix K is filled with zeros. Note however that in the case of the Tanimoto kernel type, this procedure computes the 'raw' kernel that actually corresponds to the binary kernel (16). The matrix K must be further normalized according to definition 5, that is : $K[i][j] = \frac{K[i][j]}{K[i][i]+K[j][j]-K[i][j]}$.

The cost of this algorithm depends on the number of leaf nodes visited, and is therefore bounded by the total number of leaves of the tree, that is $(np)^3$ if the number of distinct vertex labels is n and the number of distance bins is p . However, the maximum number of distinct pharmacophores that can be found in the molecule m is $|m|^3$, and we do not need to exhaustively transverse the tree. This means that to compute the kernel between the molecules m and m' , at most $\min(|m|^3, |m'|^3)$ leaves, corresponding to the common pharmacophores of m and m' , need to be visited. The complexity of the algorithm is therefore $\mathcal{O}(\min((np)^3, \min(|m|^3, |m'|^3)))^1$. For small molecules, the cost of the kernel will therefore depend on their number of atoms, while it will depend on the size of the discrete pharmacophores space for large molecules.

Note finally that although we omit the details, the previous algorithm and complexity analysis hold for the two-points versions of the kernels : the tree involved in the recursive transversal process is smaller (a depth-two tree, with n^2p leaf nodes), and the complexity is reduced to $\mathcal{O}(\min(n^2p, \min(|m|^2, |m'|^2)))$.

Algorithm 1 main program

COMPUTE(Pointer[][] *walkPointers*, Integer *depth*, Float[][] K , String *kernelType*)

```

depth = depth + 1
if depth = 1 then
  for label  $\in L_V$  do
    walkPointers = initPointers(label)
    compute(walkPointers, depth)
  end for
else
  for label1  $\in L_V$  do
    for label2  $\in L_E$  do
      walkPointers = extendPointers(walkPointers, depth, label1, label2)
      if walkPointers  $\neq$  [] then
        if depth = 4 then
          updateGram(walkPointers,  $K$ , kernelType)
        else
          compute(walkPointers, depth)
        end if
      end if
    end for
  end for
end if

```

7 Experiments

We now turn to the experimental section. The problem considered here consists in building predictive models to distinguish *active* from *inactive* molecules on several protein targets. This problem is naturally formulated as a supervised binary classification problem that can be solved by SVM.

¹Note however that in the case of the Tanimoto kernel (17), the self kernels have to be computed, and the worst case complexity of the algorithm is $\mathcal{O}(\min((np)^3, |m|^3 + |m'|^3))$.

Algorithm 2 Sub-routine 1 : initialize walks pointers

INITPOINTERS(Label *label*)

```
walkPointers = Pointer[][]
for  $i = 1, \dots, n$  do
  for  $v \in \mathcal{V}_{G_i}$  do
    if  $l(v) = \textit{label}$  then
      walkPointers[ $i$ ].addPointer(start =  $v$ , current =  $v$ )
    end if
  end for
end for
return walkPointers
```

Algorithm 3 Sub-routine 2 : extend walks pointers

EXTENDPOINTERS(Pointer[][] *walkPointers_{in}*, Integer *depth*, Label *label₁*, Label *label₂*)

```
walkPointersout = Pointer[][]
for  $i = 1, \dots, n$  do
  for  $ptr \in \textit{walkPointers}_{in}[i]$  do
    for  $(ptr.current, v) \in \mathcal{E}_{G_i}$  do
      if  $l(v) = \textit{label}_1 \wedge l(ptr.current) = \textit{label}_2$  then
        if not(  $depth = 4 \wedge v \neq ptr.start$  ) then
          walkPointersout[ $i$ ].addPointer(start =  $ptr.start$ , current =  $v$ )
        end if
      end if
    end for
  end for
end for
return walkPointersout
```

Algorithm 4 Sub-routine 3 : update Gram matrix

UPDATEGRAM(Pointer[][] *walkPointers*, Float [][] *K*, String *kernelType*)

```
for  $i = 1, \dots, n$  do
  for  $j = 1, \dots, n$  do
    if  $\textit{walkPointers}[i] \neq [] \wedge \textit{walkPointers}[j] \neq []$  then
      if kernelType = 'spectrum' then
         $update = \textit{walkPointers}[i].size() \times \textit{walkPointers}[j].size()$ 
      else
         $update = 1$ 
      end if
       $K[i][j] = K[i][j] + update$ 
      if  $i \neq j$  then
         $K[j][i] = K[j][i] + update$ 
      end if
    end if
  end for
end for
```

	TRAIN		TEST	
	Pos	Neg	Pos	Neg
BZR	94	87	63	62
COX	87	91	61	64
DHFR	84	149	42	118
ER	110	156	70	110

Table 1: Basic informations about the datasets considered.

7.1 Datasets

We tested the pharmacophore kernel on several datasets used in a recent SAR study (32). More precisely, we considered the following four publicly available datasets ²:

- the *BZR* dataset, a set of 405 ligands for the benzodiazepine receptor,
- the *COX* dataset, a set of 467 cyclooxygenase-2 inhibitors,
- the *DHFR* dataset, a set of 756 inhibitors of dihydrofolate reductase,
- the *ER* dataset, a set of 1009 estrogen receptor ligands.

These datasets contain the 3D structures of the molecules, together with a quantitative measure of their ability to inhibit a biological mechanism. Datasets were filtered and split into training and test sets according to a particular data preparation scheme detailed in (32). Table 1 gathers basic informations about the datasets involved in the study.

7.2 Experimental setup

We investigated in this study a simple labeling scheme to describe each atom (hydrogen atoms were systematically removed), and therefore the potential pharmacophores: the label of an atom is composed of its type (e.g., *C*, *O*, *N*...) and the sign of its partial charge (+, - or 0). Hence the set of labels can be expanded as $\mathcal{L} = \{C^+, C^0, C^-, O^+, O^0, O^-, \dots\}$. The partial charges account for the contribution of each atom to the total charge of the molecule, and were computed with the QuacPAC software developed by OpenEye ³. It is important to note that, contrary to the physicochemical properties of atoms, partial charges depend on the molecule and describe the spatial distribution of charges. Although the partial charges take continuous values, we simply kept their signs for the labeling as basic indicators of charges in the description of pharmacophores. We call *categorical kernel* the kernel resulting from this labeling, where the kernel between labels K_{Feat} is the Dirac kernel (10) and the kernel between distances K_{Dist} is the Gaussian RBF kernel (9).

Alternatively, we tested several variants of this basic categorical kernel. First, we tested the effect of the partial charges by removing them from the labels, and keeping the same Dirac and Gaussian RBF kernels for the labels and distances, respectively. In this case the label of an atom reduces to its type. Second, we tested alternatives to the Dirac kernel between labels, by taking into account similarities between physicochemical properties of atoms with different labels. We considered the four following properties, taken from (33) : the *Van der Waals radius*, which represents the radius of an imaginary sphere enclosing the atom, the *covalent radius*, corresponding to half of the distance between two identical covalently bonded atomic nuclei, the *first ionization energy*, the energy required to strip it of an electron from the atom, and the *electronegativity*, a measure of the ability of an atom or molecule to attract electrons in the context of a chemical bond. The Van der Waals and covalent radii account for the steric property of atoms, while the two latter properties encode their electrostatic behavior. In these cases, for computational reasons, a triangular kernel was used to compare different atoms with respect to these properties. Third, we tested the six fast approximations mentioned in Section 6 with our original labeling scheme (3- and 2-points spectrum, binary and Tanimoto kernels).

In addition, we tested the state-of-the-art Tanimoto kernel based on the 2D structure of molecules (19) to evaluate the potential gain obtained by including 3D information. This kernel is defined as the Tanimoto coefficient between fingerprints indicating the presence or absence of all possible molecular fragments of length up to 8 in the 2D structure

²Available as supporting information of the original study at <http://pubs.acs.org/journals/jcisd8/>

³<http://www.eyesopen.com/products/applications/quacpac.html>

of the molecule, where a fragment refers to a sequence of atoms connected by covalent bonds. We note that this fingerprint is similar to classical 2D-fingerprints such as the Daylight representation⁴, with the difference that our implementation does not require to fold the fingerprint into a small-size vector (18).

The different kernels were implemented in C++, and the SVM experiment was conducted with the freely available Python machine learning package PyML⁵. For each experiment, all parameters of the kernel and the SVM were optimized over a grid of possible choices on the training set only, to maximize the mean area under the ROC curve (AUC) (34) over an internal 10-fold cross-validation. The results on the test set correspond to the performance of the SVM with the selected parameters only. The optimized parameters include the width $\sigma \in \{0.1, 1, 10\}$ (in angstroms) of the Gaussian RBF kernel used to compare distances, the soft-margin parameter of the SVM over the grid $\{0.1, 0.5, 1, 1.5, \dots, 20\}$, the number of bins used to discretize the distances for the fast approximations over the grid $\{4, 6, 8, \dots, 30\}$ and the cut-off parameter C of the triangular kernel when physicochemical properties are used. This later parameter was chosen among 3 values chosen such that 10, 25 or 50% of all atom pairs in the training set have a non-zero value for the kernel. The larger C , the more atoms of distinct types are matched by the kernel, but the longer the kernel computation.

7.3 Results

Table 3 shows the results of classification for the different kernel variants. Each line corresponds to a kernel, and reports several statistics : the accuracy (fraction of correctly classified compounds), sensitivity (fraction of positive compounds that were correctly classified), specificity (fraction of negative compounds that were correctly classified), and AUC. The first line corresponds to the basic categorical kernel. The following five lines show the results of the five variants of the categorical kernel obtained by modifying the kernel between labels (Dirac kernel for labels without partial charges, and triangular kernel for 4 physicochemical properties). The results obtained by the six fast approximations follow. Finally, we added the performance obtained by the state-of-the-art 2D Tanimoto kernel based on the 2D structure of the molecules and the best results reported in the reference publication (32).

The results of parameters optimization on the training set often led to similar choices for different kernels. For example, the width of the Gaussian RBF kernel to compare distances was usually selected at 0.1 angstrom, which corresponds to a very strong constraint on the pharmacophore matching. The cut-off parameter for the triangular kernel was usually chosen to allow 10 or 25% of matches between atoms on the training set. Finally, the number of bins selected by the fast approximations to discretize the distances was usually between 20 and 30 bins.

The results show that in general, the five variants of the categorical kernel obtained by modifying the kernel between labels (lines 2-6) lead to similar or slightly worse performances than the categorical kernel. Removing the partial charges from atom labels decreases the accuracy by 3 to 5% on all datasets except COX, confirming that the partial charge information is important for the definition of pharmacophores. The variants based on the four physicochemical properties of atoms lead to results globally similar to those obtained with atom type labels without partial charges, from which they are deduced. This shows that, in the context of this study, subtle pharmacophoric features based on physicochemical parameters instead of simply the type of atoms could not be detected.

The fast pharmacophore kernel obtained by applying a Dirac kernel to check when pairs of candidate pharmacophores fall in the same bin of the discretized space (3pt-spectrum) systematically degrades the accuracy of 1 to 5% over all four datasets compared to the categorical kernel. This suggests that the gain in computation time obtained by discretizing the space and computing a 3D-fingerprint-like representation of molecules has a cost in terms of accuracy of the final model. A particular limitation of the fingerprint-based method is that two pharmacophores could remain unmatched in they fall into two different bins, although they might be very similar but close to the bins boundaries. In the case of the pharmacophore kernel, such pairs of similar pharmacophores would always be matched.

Interestingly, however, performances competitive with the categorical kernel are obtained by the fast 3pt-binary and 3pt-Tanimoto kernels. On the BZR dataset, the 3pt-binary kernel even gives the best performance. Contrary to the 3pt-spectrum kernel, these kernels are not pharmacophore kernels in the sense of Definition 2; however they are based on the same representation as the 3pt-spectrum kernel, the only difference being the way to obtain the kernel value from the fingerprint description. Note finally that these two kernels give overall similar results.

We observe moreover that except for the COX dataset, the discrete kernels based on two-points pharmacophores lead to significantly worse results than their three-points counterparts.

⁴<http://www.daylight.com/dayhtml/doc/theory/theory.toc.html>

⁵Available at <http://pyml.sourceforge.net>

	Exact	Discrete
With charges	20'	6'
Without charges	249'	7'

Table 2: Computation times in minutes needed to compute the different kernel matrices on the BZR training set. The first column refers to the computation of the exact kernel (6), and the second one to the approximate kernels (15) and (17).

For each dataset, the results obtained with the 2D-Tanimoto kernel are significantly worse than those of the categorical kernel, with a decrease ranging from 3 to 7% on the different datasets. This confirms the relevance of 3D information for drug activity prediction, that motivated this work. Finally we note that on all but the COX dataset, the categorical kernel outperforms the best results of (32), confirming the competitiveness of our method compared to state-of-the-art methods.

Regarding the computational complexity of the different methods, Table 2 shows the time required to compute the kernel matrices on the BZR training set for different kernels, on a desktop computer, equipped with a Pentium 4 - 3.6 GHz processor, and 1 GB RAM. In the discrete version, the distance range was split into 24 bins, and as expected, the kernels based on the discretization of the pharmacophore space are faster than their counterparts by a factor of 4 to 35, depending on the type of labels used (with or without the partial charge information). In the exact kernel computation, the effect of removing the partial charges from the labels is to induce more matches between atoms and therefore, as discussed in Section 4, to drastically slow the computation by a factor of 12, consistent with the theoretical estimate that dividing the size of the label classes by k increases the speed by a factor k^3 .

8 Discussion and conclusion

This paper presents an attempt to extend the application of recent machine learning algorithms for classification to the manipulation of 3D structures of molecules. This attempt is mainly motivated by applications in drug activity prediction, for which 3D pharmacophores are known to play important roles. Although previous attempts to define kernels for 3D structures (similar in fact to the 2pt-spectrum kernel we tested) led to mixed results (35), we obtained performance competitive with state-of-the-art algorithms for the categorical kernel based on the comparison of pharmacophores contained in the two molecules to be compared. This kernel is not an inner product between fingerprints, and therefore fully exploits the mathematical trick that allows SVM to manipulate measures of similarities rather than explicit vector representations of molecules, as opposed to other methods such as neural networks. We even observed that for the closest fingerprint-based approximation obtained by discretizing the space of possible pharmacophores (3pt-spectrum kernel), the performance significantly decreases. This highlights the benefits that can be gained from the use of kernels, which provide a satisfactory answer to the common issue of choosing a "good" discretization of the pharmacophore space to make fingerprints: once discretized, pharmacophores falling on different sides of bins edges do not match although they might be very close. We notice that approaches based on fuzzy fingerprints (36), for example, aim at correcting this effect by matching pharmacophores based on different distance bins.

Although the best overall method in this study is the categorical kernel, it is interesting to notice that very competitive results are obtained by the binary and Tanimoto kernels applied to the discretized pharmacophore representation. Compared to the 3pt-spectrum, the better performance of the binary kernel suggests that the choice of the functional form of the kernel given a representation of molecules can play a critical role in terms of performance. Representing pharmacophores by indicators (bits) of their presence rather than by their precise counts can be interpreted as a trivial way to emphasize rare versus frequent pharmacophores in the kernels. Alternatively, it might be possible for example, to adopt more flexible schemes to weight the pharmacophores depending on their probability of appearance in the molecules, and to modify in a similar way the functional form of the pharmacophore kernels (Definition 2) to improve performance over the categorical kernel.

Concerning the practical use of our approach for screening of large datasets, Table 2 shows that, even for the fastest variants, the approach based on kernel methods can be computationally demanding even for relatively small datasets. In practice, however, the time to train the SVM can be smaller than the times presented in Table 2 because not all entries of the matrix are required. Speeding up SVM and kernel methods for large datasets is currently a topic of interest in the machine learning community, and applications in virtual screening on large databases of molecules will

certainly benefit from the advances in this field.

Among the possible extensions to our work, a first direction would be to test and validate different definitions and labeling for the vertices of the pharmacophores. We limited ourselves to the simplest possible 3-points pharmacophores based on single atoms annotated by their types and partial charges. The method could be improved by testing other schemes known to be relevant features as basic components of pharmacophores. Instead of single atoms, it is for example possible to consider groups of atoms forming functional units instead of single atoms to form pharmacophores. A second possible extension is to generalize this work to pharmacophores with more points, e.g., 4 or 5. Although several results will not remain valid in this case, such as the expression of the kernel as the trace of a matrix, this could lead to more accurate models in cases where the binding mechanism is well characterized by such pharmacophores. Finally, a third promising direction that is likely to be relevant for many real-world applications is to take into accounts different conformers of each molecule. Indeed, it is well-known that the biological activity to be predicted is often due to one out of several conformers for a given molecule, which suggests to represent a molecule not as a single 3D structure but as a set of structures. The kernel approach lends itself particularly well to this extension, thanks to the possibility to define kernels between sets of structures from a kernel between structures, just like we defined a kernel between sets of pharmacophores from a kernel between single pharmacophores.

References

- [1] Charles Manly, Shirley Louise-May, and Jack Hammer. The impact of informatics and computational chemistry on synthesis and screening. *Drug Discov. Today*, 6(21):1101–1110, Nov 2001.
- [2] Darko Butina, Matthew D Segall, and Katrina Frankcombe. Predicting ADME properties in silico: methods and models. *Drug Discov. Today*, 7(11 Suppl):S83–S88, Jun 2002.
- [3] C. A. Lipinski, F. Lombardo, B. W. Dominy, and P. J. Feeney. Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv. Drug. Deliv. Rev.*, 46(1-3):3–26, Mar 2001.
- [4] R. D. King, S. H. Muggleton, A. Srinivasan, and M. J. E. Sternberg. Structure-activity relationships derived by machine learning: The use of atoms and their bond connectivities to predict mutagenicity by inductive logic programming. *Proc. Natl. Acad. Sci. USA*, 93:438–442, 1996.
- [5] Hans-Joachim Böhm, Gisbert Schneider, Raimund Mannhold, Hugo Kubinyi, and Gerd Folkers. *Protein-ligand interactions*. Wiley, 2003.
- [6] Inbal Halperin, Buyong Ma, Haim Wolfson, and Ruth Nussinov. Principles of docking: An overview of search algorithms and a guide to scoring functions. *Proteins*, 47(4):409–443, Jun 2002.
- [7] C. Lemmen and T. Lengauer. Computational methods for the structural alignment of molecules. *J. Comput. Aided. Mol. Des.*, 14(3):215–232, Mar 2000.
- [8] L. Xue and J. Bajorath. Molecular descriptors in chemoinformatics, computational combinatorial chemistry, and virtual screening. *Comb. Chem. High. Throughput Screen.*, 3(5):363–372, Oct 2000.
- [9] J. D. Holliday and P. Willett. Using a genetic algorithm to identify common structural features in sets of ligands. *J. Mol. Graph. Model.*, 15(4):221–232, Aug 1997.
- [10] Paul Finn, Stephen Muggleton, David Page, and Ashwin Srinivasan. Pharmacophore discovery using the inductive logic programming language Progol. *Machine Learning*, 30:241–270, 1998.
- [11] Hans Matter and Thorsten Pötter. Comparing 3D pharmacophore triplets and 2D fingerprints for selecting diverse compound subsets. *J. Chem. Inf. Model.*, 39:1211–1225, 1999.
- [12] R. D. Brown and Y. C. Martin. The information content of 2D and 3D structural descriptors relevant to ligand-receptor binding. *J. Chem. Inf. Comput. Sci.*, 37:1–9, 1997.
- [13] J. Bajorath. Selected concepts and investigations in compound classification, molecular descriptor analysis, and virtual screening. *J. Chem. Inf. Comput. Sci.*, 41(2):233–245, 2001.

	BZR				COX				DHFR				ER			
	Acc.	Sens.	Spec.	AUC	Acc.	Sens.	Spec.	AUC	Acc.	Sens.	Spec.	AUC	Acc.	Sens.	Spec.	AUC
Categorical	76.4	74.0	78.9	82.1	69.8	69.8	69.8	75.1	81.9	63.3	88.8	84.8	79.8	72.0	84.7	86.8
Categorical, no partial charges	74.3	73.6	75.0	81.5	70.0	68.5	70.9	74.6	78.1	65.2	82.7	82.2	77.6	71.7	81.4	87.2
Van der Waals	74.6	74.0	75.3	81.7	67.1	67.2	67.0	72.1	77.8	68.6	81.1	83.1	77.7	72.5	80.9	87.1
Covalent	74.3	73.6	75.0	80.8	70.0	68.5	71.2	74.3	76.9	68.8	79.7	82.6	77.8	72.0	81.5	87.3
Ionization	74.1	73.8	74.7	81.8	68.9	68.4	69.4	74.0	77.9	66.7	82.0	82.8	77.6	72.3	81.5	87.1
Electroneg.	74.9	74.4	75.3	81.6	70.2	67.9	72.5	73.4	77.9	67.4	81.6	83.0	78.2	70.4	83.1	87.7
3pt-spectrum	75.4	74.4	76.3	81.3	67.0	64.4	69.5	75.9	76.9	70.9	79.0	81.9	78.6	78.3	78.8	87.4
3pt-binary	78.5	74.4	82.6	81.5	68.2	70.5	65.9	74.8	80.8	66.2	85.9	81.1	79.3	74.7	82.2	87.5
3pt-Tanimoto	78.3	74.6	82.1	84.7	68.0	68.0	68.0	74.2	81.6	69.8	85.6	83.1	79.0	67.9	86.4	88.7
2pt-spectrum	71.4	61.3	81.6	80.3	68.9	70.2	67.7	74.7	67.7	67.4	67.9	72.3	78.7	75.9	80.4	84.5
2pt-binary	72.3	66.5	78.2	77.2	71.3	71.0	71.6	76.5	66.5	78.3	62.3	76.2	75.6	87.8	67.8	84.8
2pt-Tanimoto	75.0	69.7	80.5	80.3	69.8	67.0	72.3	74.2	72.4	71.9	72.5	80.6	74.3	85.6	67.1	85.1
2D-Tanimoto	71.2	71.9	70.5	80.8	63.0	67.5	58.6	69.8	76.9	73.8	78.0	83.0	77.1	69.3	82.1	83.6
Sutherland	75.2	70.0	81.0	XXX	73.6	75.0	72.0	XXX	71.9	74.0	71.0	XXX	78.9	77.0	80.0	XXX

Table 3: Classification of the test sets, after model selection on the training set.

- [14] 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.
- [15] V. N. Vapnik. *Statistical Learning Theory*. Wiley, New-York, 1998.
- [16] 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.
- [17] E. Byvatov, U. Fechner, J. Sadowski, and G. Schneider. Comparison of support vector machine and artificial neural network systems for drug/nondrug classification. *J Chem Inf Comput Sci*, 43(6):1882–9, 2003.
- [18] 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.
- [19] L. Ralaivola, S. J. Swamidass, H. Saigo, and P. Baldi. Graph kernels for chemical informatics. *Neural Netw.*, 18(8):1093–1110, Sep 2005.
- [20] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, Cambridge, MA, 2002.
- [21] 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.
- [22] C. J. C. Burges. A Tutorial on Support Vector Machines for Pattern Recognition. *Data Min. Knowl. Discov.*, 2(2):121–167, 1998.
- [23] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge University Press, 2004.
- [24] M. A. Aizerman, E. M. Braverman, and L. I. Rozonoér. Theoretical foundations of the potential function method in pattern recognition learning. *Automation and Remote Control*, 25:821–837, 1964.
- [25] N. Aronszajn. Theory of reproducing kernels. *Trans. Am. Math. Soc.*, 68:337 – 404, 1950.
- [26] O. F. Güner. *Pharmacophore Perception, Development, and Use in Drug Design*, volume 2 of *IUL Biotechnology Series*. International University Line, 2000.
- [27] D. Haussler. Convolution Kernels on Discrete Structures. Technical Report UCSC-CRL-99-10, UC Santa Cruz, 1999.
- [28] T. Gärtner, P. Flach, and S. Wrobel. On graph kernels: hardness results and efficient alternatives. In B. Schölkopf and M. Warmuth, editors, *Proc. 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.
- [29] S. D. Pickett, J. S. Mason, and I. M. McLay. Diversity profiling and design using 3D pharmacophores : pharmacophores-derived queries (PQD). *J. Chem. Inf. Comput. Sci.*, 36:1214–1223, 1996.
- [30] J. Saeh, P. Lyne, B. Takasaki, and D. Cosgrove. Lead hopping using SVM and 3D pharmacophore fingerprints. *J. Chem. Inf. Model.*, 45(4):1122–1133, Jul 2005.
- [31] 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.
- [32] J. J. Sutherland, L. A. O’Brien, and D. F. Weaver. Spline-fitting with a genetic algorithm : a method for developing classification structure-activity relationships. *J. Chem. Inf. Comput. Sci.*, 43:1906–1915, 2003.
- [33] John Emsley. *The Elements (third edition)*. Oxford University Press, 1998.
- [34] T. Fawcett. ROC graphs: notes and practical considerations for data mining researchers. Technical Report 2003-4, HP Laboratories, Palo Alto, CA, USA, 2003.

- [35] 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.
- [36] Dragos Horvath and Catherine Jeandenans. Neighborhood behavior of in silico structural spaces with respect to in vitro activity spaces—a novel understanding of the molecular similarity principle in the context of multiple receptor binding profiles. *J. Chem. Inf. Comput. Sci.*, 43:680–690, 2003.