

Integrating structured biological data by Kernel Maximum Mean Discrepancy

Karsten M. Borgwardt^a, Arthur Gretton^b, Malte J. Rasch,^c
Hans-Peter Kriegel,^a Bernhard Schölkopf,^b Alex J. Smola^d

^a University of Munich, ^b Max Planck Institute for Biological Cybernetics Tübingen, ^c University of Technology Graz, ^d NICTA Canberra

ABSTRACT

Motivation: Many problems in data integration in bioinformatics can be posed as one common question: Are two sets of observations generated by the same distribution? We propose a kernel-based statistical test for this problem, based on the fact that two distributions are different if and only if there exists at least one function having different expectation on the two distributions. Consequently we use the maximum discrepancy between function means as the basis of a test statistic.

The Maximum Mean Discrepancy (MMD) can take advantage of the kernel trick, which allows us to apply it not only to vectors, but strings, sequences, graphs, and other common structured data types arising in molecular biology.

Results: We study the practical feasibility of an MMD-based test on three central data integration tasks: Testing cross-platform comparability of microarray data, cancer diagnosis, and data-content based schema matching for two different protein function classification schemas. In all of these experiments, including high-dimensional ones, MMD is very accurate in finding samples that were generated from the same distribution, and outperforms or is as good as its best competitors.

Conclusions: We have defined a novel statistical test of whether two samples are from the same distribution, compatible with both multivariate and structured data, that is fast, easy to implement, and works well, as confirmed by our experiments.

Availability: <http://www.dbs.ifi.lmu.de/~borgward/MMD>

Contact: kb@dbs.ifi.lmu.de

1 INTRODUCTION

1.1 Data Integration in Bioinformatics

The ultimate need for bioinformatics is founded on the wealth of data generated by modern molecular biology. The purpose of bioinformatics is to structure and analyze this data. A central preprocessing step is the integration of datasets that were generated by different laboratories and techniques. If we know how to combine data produced in different labs, we can exploit the results jointly, not only individually. In some cases, the larger datasets thus constructed may support biologically relevant conclusions which were not possible using

the original smaller datasets, a hypothetical example being the problem of reliable gene selection from high dimensional small microarray datasets.

1.2 Distribution Testing in Data Integration

The questions arising in data integration essentially boil down to the following problem of distribution testing: Were two samples X and Y generated by the same distribution? In data integration terms, are these two samples part of the same larger dataset, or should these data be treated as originating from two different sources?

This is a fundamental question when two laboratories are studying the same biological subject. If they use identical techniques on identical subjects but yield results that are not generated by the same distribution, then this might indicate there is a difference in the way they generate data, and that their results should not be integrated directly. If the data were integrated without recalibration, differences or patterns within the joint data might be caused by experimental discrepancies between laboratories, rather than by biological processes.

As microarray data are produced by a multitude of different platforms, techniques and laboratories, they are the most prominent data structure in bioinformatics for which distribution testing is indispensable. Recently, Marshall (2004) gave an extremely negative picture of cross-platform comparability — and hence the reliability and reproducibility — of microarray technology, due to the various platforms and data analysis methods employed (Shi et al., 2005). It is therefore crucial for bioinformatics to develop computational methods that allow us to determine whether results achieved across platforms are comparable. In this article, we present a novel statistical test to tackle this problem.

What distinguishes bioinformatics is that it has produced a wealth of complex data, from protein sequences (which are strings) to protein interaction networks (which are graphs). Consequently any practically relevant distribution test needs to be *easily* applicable in all these cases. To the best of our knowledge, the statistical test proposed in our paper is the first method that can handle this wide range of different domains.

To summarize our goals, we will present a novel statistical test for differences in distribution, based on the Maximum Mean Discrepancy (MMD). We will show that it can take advantage of the kernel trick. Hence it is applicable to all data types, from high dimensional vectors to strings to graphs, arising in bioinformatics. In experiments, we will apply this test to microarray cross-platform comparability testing and cancer diagnosis. Furthermore, we will show how to perform schema matching on complex data by considering a data integration problem on two molecular graph datasets.

Outline of this article In Section 2, we present MMD and its properties. In Section 3, we first test the applicability of MMD in cross-platform microarray comparability analysis, before evaluating its usefulness on a schema matching problem. We discuss our findings in Section 4.

2 MMD AND THE TWO-SAMPLE PROBLEM

In statistics, the central question of data integration described above is often referred to as the *two-sample* or *homogeneity problem*. The principle underlying the maximum mean discrepancy is that we want to find a function that assumes different expectations on two different distributions. The hope then is that if we evaluate this function on empirical samples from the distributions, it will tell us whether the distributions they have been drawn from are likely to differ. This leads to the following statistic, which is closely related to a proposal by Fortet and Mourier (1953). Here and below, \mathcal{X} denotes our input domain and is assumed to be a nonempty compact set.

Definition 2.1. *Let \mathcal{F} be a class of functions $f : \mathcal{X} \rightarrow \mathbb{R}$. Let p and q be Borel probability distributions, and let $X = (x_1, \dots, x_m)$ and $Y = (y_1, \dots, y_n)$ be samples composed of independent and identically distributed observations drawn from p and q , respectively. We define the maximum mean discrepancy (MMD) and its empirical estimate as*

$$\text{MMD}[\mathcal{F}, p, q] := \sup_{f \in \mathcal{F}} (\mathbf{E}_p[f(x)] - \mathbf{E}_q[f(y)])$$

$$\text{MMD}[\mathcal{F}, X, Y] := \sup_{f \in \mathcal{F}} \left(\frac{1}{m} \sum_{i=1}^m f(x_i) - \frac{1}{n} \sum_{i=1}^n f(y_i) \right)$$

Intuitively it is clear that if \mathcal{F} is “rich enough”, $\text{MMD}[\mathcal{F}, p, q]$ will vanish if and only if $p = q$. Too rich an \mathcal{F} , however, will result in a statistic that differs significantly from zero for most finite samples X, Y . For instance, if \mathcal{F} is the class of *all* real valued functions on \mathcal{X} , and if X and Y are disjoint, then it is trivial to construct arbitrarily large values of $\text{MMD}[\mathcal{F}, X, Y]$, for instance by ensuring that $f|_X$ is large and $f|_Y = 0$. This phenomenon of *overfitting* can be avoided by placing restrictions on the function class. That said, these restrictions ought not to prevent the MMD from detecting differences between p and q when these are legitimately to be found. As we shall see, one way to accomplish this tradeoff is by choosing \mathcal{F} to be the unit ball in a universal reproducing kernel Hilbert space.

We will propose a test of $p = q$, based on an unbiased variant of MMD $[\mathcal{F}, X, Y]$,¹ which relies on the asymptotic Gaussianity of this test statistic and on the guaranteed rapid convergence to this asymptotic regime. Thus, the performance guarantees provided by the test apply in the case of a large sample size. The test has a computational cost of $O((m+n)^2)$, although randomization techniques could be employed to reduce the cost to essentially linear time-complexity (at the expense of a somewhat reduced sensitivity).

2.1 MMD for Kernel Function Classes

We now introduce a class of functions for which MMD may easily be computed, while retaining the ability to detect all discrepancies between p and q without making any simplifying assumptions. To this end, let \mathcal{H} be a complete inner product space (i.e., a Hilbert space) of functions $f : \mathcal{X} \rightarrow \mathbb{R}$, where \mathcal{X} is a nonempty compact set. Then \mathcal{H} is termed a reproducing kernel Hilbert space if for all $x \in \mathcal{X}$, the linear point evaluation functional mapping $f \rightarrow f(x)$ exists and is continuous. In this case, $f(x)$ can be expressed as an *inner product* via

$$f(x) = \langle f, \phi(x) \rangle_{\mathcal{H}} \quad (1)$$

where $\phi : \mathcal{X} \rightarrow \mathcal{H}$ is known as the *feature space map* from x to \mathcal{H} . Moreover, the inner product between two feature maps is called the (*positive definite*) *kernel*, $k(x, x') := \langle \phi(x), \phi(x') \rangle_{\mathcal{H}}$. Of particular interest are cases where we have an analytic expression for k that can be computed quickly, despite \mathcal{H} being high or even infinite dimensional. An example of an infinite dimensional \mathcal{H} is that corresponding to the Gaussian kernel $k(x, x') = \exp(-\|x - x'\|^2 / (2\sigma^2))$.

We will consider universal reproducing kernel Hilbert spaces in the sense defined by Steinwart (2002). Although we do not go into technical detail here, we are guaranteed that RKHSs based on Gaussian kernels are universal, as are string kernels (Section 2.3). See also (Schölkopf et al., 2004) for an extensive list of further kernels.

When \mathcal{F} is the unit ball in a universal RKHS, the following theorem (Smola et al., 2006) guarantees that $\text{MMD}[\mathcal{F}, p, q]$ will detect any discrepancy between p and q .

Theorem 2.2. *Let p, q be Borel probability measures on \mathcal{X} , a compact subset of a metric space, and let \mathcal{H} be a universal reproducing kernel Hilbert space with unit ball \mathcal{F} . Then $\text{MMD}[\mathcal{F}, p, q] = 0$ if and only if $p = q$.*

Moreover, denote by $\mu_p := \mathbf{E}_p[\phi(x)]$ the expectation of $\phi(x)$ in feature space (assuming that it exists).² Then one may rewrite MMD as

$$\text{MMD}[\mathcal{F}, p, q] = \|\mu_p - \mu_q\|_{\mathcal{H}}. \quad (2)$$

¹ Note that $\text{MMD}[\mathcal{F}, X, Y]$ as defined above is biased: even when $p = q$, it will tend to give strictly positive results for finite sample sizes.

² A sufficient condition for this is $\|\mu_p\|_{\mathcal{H}}^2 < \infty$, which is rearranged as $\mathbf{E}_p[k(x, x')] < \infty$, where x and x' are independent random variables drawn according to p .

The main ideas for the proof can be summarized as follows. It is known from probability theory (Dudley, 2002, Lemma 9.3.2) that under the stated conditions, a sufficient condition for $p = q$ is that for all continuous functions f , we have $\int f dp = \int f dq$. Such functions f , however, can be arbitrarily well approximated using functions in a universal RKHS (Steinwart, 2002). For the second part of the result, observe that due to (1), we may rewrite the MMD as

$$\begin{aligned} \text{MMD}[\mathcal{F}, p, q] &= \sup_{\|f\|_{\mathcal{H}} \leq 1} \mathbf{E}_p[f(x)] - \mathbf{E}_q[f(y)] \\ &= \sup_{\|f\|_{\mathcal{H}} \leq 1} \mathbf{E}_p[\langle \phi(x), f \rangle_{\mathcal{H}}] - \mathbf{E}_q[\langle \phi(y), f \rangle_{\mathcal{H}}] \\ &= \sup_{\|f\|_{\mathcal{H}} \leq 1} \langle \mu_p - \mu_q, f \rangle_{\mathcal{H}} = \|\mu_p - \mu_q\|_{\mathcal{H}}. \end{aligned}$$

The finite sample computation of MMD is greatly simplified by (2), as shown in the corollary below:

Corollary 2.3. *Under the assumptions of theorem 2.2 the following is an unbiased estimator of $\text{MMD}^2[\mathcal{F}, p, q]$:*

$$\begin{aligned} \text{MMD}^2[\mathcal{F}, X, Y] &= \frac{1}{m(m-1)} \sum_{i \neq j}^m k(x_i, x_j) \\ &+ \frac{1}{n(n-1)} \sum_{i \neq j}^n k(y_i, y_j) - \frac{2}{mn} \sum_{i,j=1}^{m,n} k(x_i, y_j). \end{aligned}$$

Proof We compute

$$\begin{aligned} \text{MMD}^2[\mathcal{F}, p, q] &:= \langle \mu_p - \mu_q, \mu_p - \mu_q \rangle_{\mathcal{H}} \\ &= \langle \mu_p, \mu_p \rangle_{\mathcal{H}} + \langle \mu_q, \mu_q \rangle_{\mathcal{H}} - 2 \langle \mu_p, \mu_q \rangle_{\mathcal{H}} \\ &= \mathbf{E}_p \langle \phi(x), \phi(x') \rangle_{\mathcal{H}} + \mathbf{E}_q \langle \phi(y), \phi(y') \rangle_{\mathcal{H}} \\ &\quad - 2 \mathbf{E}_{p,q} \langle \phi(x), \phi(y) \rangle_{\mathcal{H}}, \end{aligned}$$

where x' is a random variable independent of x with distribution p , and y' is a random variable independent of y with distribution q . The proof is completed by applying $\langle \phi(x), \phi(x') \rangle_{\mathcal{H}} = k(x, x')$, and replacing the expectations with their empirical counterparts. ■

We illustrate the behavior of MMD in Figure 1 using a one-dimensional example: the data X and Y are generated from distributions p and q with equal means and variances, however p is Gaussian and q is Laplacian. For the application of MMD we pick \mathcal{H} to be an RKHS using the Gaussian kernel. We observe that the function f that witnesses the MMD (in other words, the function maximizing the mean discrepancy) is smooth, positive where the Laplace density exceeds the Gaussian density (at the center and tails), and negative where the Gaussian density is larger. Moreover, the magnitude of f is

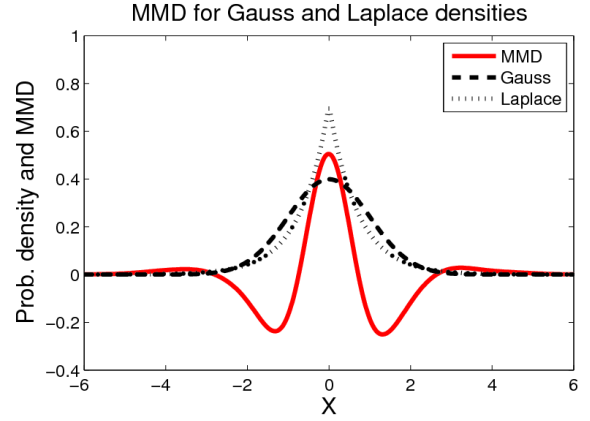


Fig. 1. Illustration of the function maximizing the mean discrepancy in the case where a Gaussian is being compared with a Laplace distribution. Both distributions have zero mean and unit variance. The maximizer of the MMD has been scaled for plotting purposes, and was computed empirically on the basis of 2×10^4 samples, using a Gaussian kernel with $\sigma = 0.5$.

a direct reflection of the amount by which one density exceeds the other, insofar as the smoothness constraint permits it.³

Although the expression of $\text{MMD}^2(\mathcal{F}, X, Y)$ in Corollary 2.3 is the minimum variance unbiased estimate (Serfling, 1980), a more tractable unbiased expression can be found in the case where $m = n$, with a slightly higher variance (the distinction is in practice irrelevant, since the terms that differ decay much faster than the variance). It is obtained by dropping the cross-terms $i = j$ from the sum over $k(x_i, y_j)$:

Lemma 2.4. *Assuming the samples X and Y both have size m , define $z_i = (x_i, y_i)$, and let*

$$h(z_i, z_j) := k(x_i, x_j) + k(y_i, y_j) - k(x_i, y_j) - k(x_j, y_i).$$

An unbiased estimate of $\text{MMD}^2[\mathcal{F}, p, q]$ is given by

$$\text{MMD}^2[\mathcal{F}, X, Y] := \frac{1}{m(m-1)} \sum_{i \neq j}^m h(z_i, z_j).$$

Note that with some abuse of notation we used the *same* symbol as in Corollary 2.3 for a slightly different estimator. However there should be no ambiguity in that we use only the present version for the remainder of the paper.

An important property of the new statistic is that its kernel $h(z_i, z_j)$ is a positive definite kernel in its own right, since

$$h(z_i, z_j) = \langle \phi(x_i) - \phi(y_i), \phi(x_j) - \phi(y_j) \rangle.$$

³ One may show that the maximizer of $\text{MMD}[\mathcal{F}, p, q]$ is given by $f(x) = \langle \mu_p - \mu_q, \phi(x) \rangle$. The same holds true for the maximizer of the empirical quantity, however with the means being replaced by empirical means. See (Smola et al., 2006) for further details and a proof.

Thus $z = (x, y) \mapsto \phi(x) - \phi(y)$ is a valid feature map for h . This gives another interpretation of MMD: it is the expected inner product between vectors obtained by connecting a point from the one distribution to a point of the other one.

As a consequence of the above, $\text{MMD}^2[\mathcal{F}, X, Y]$ is a positive definite kernel in its own right. For detailed discussions of the problem of defining kernels between distributions and sets, see Cuturi et al. (2005); Hein and Bousquet (2005).

2.2 MMD Tests

We now propose a two-sample test based on the asymptotic distribution of an unbiased estimate of MMD^2 , which applies in the case where \mathcal{F} is a unit ball in a RKHS, and $m = n$. This uses the following theorem, due to Hoeffding (1948). See also Serfling (1980, Section 5.5.1). For a proof and further details see Smola et al. (2006).

Theorem 2.5. *Let z_i and $h(z_i, z_j)$ as in Definition 2.4 and assume that $\mathbf{E}_{p,q}[\text{MMD}^4[\mathcal{F}, X, Y]] < \infty$. Then for $m \rightarrow \infty$, the statistic $\text{MMD}^2(\mathcal{F}, X, Y)$ converges in distribution to a Gaussian with mean $\text{MMD}^2[\mathcal{F}, p, q]$ and variance*

$$\sigma_{\text{MMD}}^2 = \frac{2^2}{m} \left(\mathbf{E}_z [(\mathbf{E}_{z'} h(z, z'))^2] - [\mathbf{E}_{z,z'} (h(z, z'))]^2 \right)$$

The convergence to the normal occurs rapidly: according to Serfling (1980, Theorem B, p. 193), the CDF of the U-statistic converges uniformly to the asymptotic CDF at rate $1/\sqrt{m}$.

Our goal is to test whether the above normal distribution has zero mean (the null hypothesis), as opposed to a mean that is positive. Since we need not concern ourselves with negative deviations from the mean ($\text{MMD}[\mathcal{F}, p, q] \geq 0$ may never become negative), it suffices to test whether $\text{MMD}^2[\mathcal{F}, X, Y] \leq \epsilon$ for some threshold ϵ . Thus, we obtain the two-sample test below as a corollary to Theorem 2.5, following the principles outlined by Casella and Berger (2002, Section 10.3.2).

Corollary 2.6. *A test of the null hypothesis $p = q$ with asymptotic size⁴ α , and asymptotic Type II error zero, has the acceptance region*

$$\begin{aligned} \text{MMD}^2[\mathcal{F}, X, Y] &\leq \hat{\sigma}_{\text{MMD}} z_\alpha \\ \text{where } \hat{\sigma}_{\text{MMD}}^2 &= \frac{4}{m^2(m-1)^2} \sum_{i=1}^m \left(\sum_{j \neq i}^m h(z_i, z_j) \right)^2 \\ &\quad - \frac{4}{m} \text{MMD}^4[\mathcal{F}, X, Y] \end{aligned}$$

or any empirical estimate of σ_{MMD} that converges in probability. Here z_α satisfies $\Pr(z > z_\alpha) = \alpha$ when $z \sim N(0, 1)$.

It is also of interest to estimate the p -value of the test. We first describe a sample-based heuristic. We draw randomly

without replacement from the aggregated data $Z = \{X, Y\}$ to get two new m -samples X^* and Y^* , and compute the test statistic $\text{MMD}_*^2(\mathcal{F}, X^*, Y^*)$ between these new samples (bear in mind that under the null hypothesis $p = q$, this aggregation is over data drawn from a single distribution). We repeat this procedure t times to get a set of test statistics under the null hypothesis (conditioned on the observations). We then add the original statistic $\text{MMD}^2(\mathcal{F}, X, Y)$ to this set, and sort the set in ascending order. Define as r the rank of the original test statistic within this ordering. Then our estimated p -value is $p = (t + 1 - r)/(t + 1)$. Alternatively, we can find an upper bound on p using the distribution-free large deviation result of Hoeffding (1963, p. 25) (see Smola et al., 2006, Section 6), which is exact at finite sample sizes. This bound is only tight when m is large, however, and may be too conservative at small sample sizes.

We give the complete pseudocode for the above MMD-based test in Algorithm 1. We emphasize that the computational cost is $O(m^2)$, and that the method is easily parallelized (the kernel matrix can be broken up into submatrices, and the relevant sums computed independently before being combined). In addition, the kernel matrix need never be stored in memory, but only a running sum must be kept, which makes feasible the analysis of very large data sets. Randomized methods could also be used to speed up the double-loop required for evaluating Algorithm 1, by only computing parts of the sum. This procedure would reduce the quality of the test, however.

Algorithm 1 MMD test using asymptotic normality

Input: positive definite kernel k , level of test $\alpha \in (0, 1)$, samples X and Y of size m drawn from p and q respectively
 $\text{MMD}^2 \leftarrow 0$ and $\sigma^2 \leftarrow 0$

for $i = 1$ to m **do**

$t \leftarrow 0$

for $j = 1$ to m **do**

if $j \neq i$ **then**

$t \leftarrow t + k(x_i, x_j) + k(y_i, y_j) - k(x_i, y_j) - k(x_j, y_i)$

end if

end for

$\text{MMD}^2 \leftarrow \text{MMD}^2 + \frac{1}{m(m-1)} t$ and $\sigma^2 \leftarrow \sigma^2 + t^2$

end for

$\sigma^2 \leftarrow \frac{4}{(m^2(m-1)^2)} \sigma^2 - \frac{4}{m} (\text{MMD}^2)^2$

$\epsilon \leftarrow \sqrt{2\sigma^2} \text{erfinv}(1 - 2\alpha)$

Output: If $\text{MMD}^2 \leq \epsilon$ return $p = q$ accepted. Otherwise return $p = q$ rejected.

Finally, we note that other approaches are also possible in determining the acceptance region of the test. For instance, Smola et al. (2006) describe two tests based on large deviation bounds: the first uses Rademacher averages to obtain a bound that explicitly accounts for the variation in the test

⁴ Size and level are defined following Casella and Berger (2002, Section 8.3).

statistic, the second uses a distribution-independent upper bound on the test statistic variation due to Hoeffding (1963, p. 25). These approaches have the advantage of giving an exact, distribution-free test level α that holds for finite samples, and not just in the asymptotic regime. In addition, they provide a finite sample upper bound on the p -value, which is again distribution-free. A disadvantage of these approaches is that they require a larger sample size than the test in Corollary 2.6 before they can detect a given disparity between the distributions p and q , i.e. they have a higher Type II error. For this reason, we do not use these tests in Section 3.

2.3 Universal Kernels for Discrete Data

While many examples of universal kernels on compact subsets of \mathbb{R}^d are known (Steinwart, 2002), little attention has been given to finite domains. It turns out that the issue is considerably easier in that case: the weaker notion of *strict positive definiteness* (kernels inducing nonsingular Gram matrices $(k(x_i, x_j))_{i,j}$ for arbitrary sets of distinct points x_i) ensures that every function on a discrete domain $\mathcal{X} = \{x_1, \dots, x_n\}$ lies in the corresponding RKHS (and hence that the kernel is universal). To see this, let $f \in \mathbb{R}^n$ be an arbitrary function on \mathcal{X} . Then $\alpha = K^{-1}f$ ensures that the function $f = \sum_j \alpha_j k(\cdot, x_j)$ satisfies $f(x_i) = f_i$ for all i .

It turns out that string kernels fall in this class:

Theorem 2.7. *Let \mathcal{X} be a finite set of strings, and denote by $\#_s(x)$ the number of times substring s occurs in x . Then any string kernel of the form $k(x, x') = \sum_{s \in \mathcal{X}} w_s \#_s(x) \#_s(x')$ with $w_s > 0$ for all $s \in \mathcal{X}$ is strictly positive definite.*

Proof We will show that the vectors $\{\phi(x) | x \in \mathcal{X}\}$ obtained by the feature map are linearly independent, implying that all Gram matrices are nonsingular. The feature map is given by $\phi(x) = (\sqrt{w_s} \#_s(x), \sqrt{w_{s'}} \#_{s'}(x), \dots)$ where we assume for the purpose of the proof that all substrings s are ordered by nondecreasing length. Now for a given set X of size m consider the matrix with columns $\phi(x_1), \dots, \phi(x_m)$ where the entries in X are assumed to be ordered in the same manner as the substrings (i.e. by nondecreasing length). By construction the upper triangle of this matrix is zero, with the highest nonzero entry of each row being $\sqrt{w_x}$, which implies linear independence of its rows. ■

For graphs unfortunately no strictly positive definite kernels exist which are efficiently computable. Note first that it is necessary for strict positive definiteness that $\phi(x)$ be injective, for otherwise we would have $\phi(x) = \phi(x')$ for some $x \neq x'$, implying that the kernel matrix obtained from $X = \{x, x'\}$ is singular. However, as Gärtner et al. (2003) show, an injective $\phi(x)$ allows one to match graphs by computing $\|\phi(x) - \phi(x')\|^2 = k(x, x) + k(x', x') - 2k(x, x')$. Graph matching, however, is NP-hard, hence no such kernel

can exist. That said, there exists a number of useful graph kernels. See e.g. (Borgwardt et al., 2005) for further details.

2.4 Kernel Choice

So far, we have focused on the case of universal kernels. These kernels have various favorable properties, including the facts that

- universal kernels are strictly positive definite, making the kernel matrix invertible and avoiding non-uniqueness in the dual solutions of SVMs, and
- continuous functions on \mathcal{X} can be arbitrarily well approximated (in the $\|\cdot\|_\infty$ -norm) using an expansion in terms of universal kernels, and SVMs using universal kernels are consistent in the sense that (subject to certain conditions) their solutions converge to the Bayes optimal solution (Steinwart, 2002),
- MMD using universal kernels is a test for identity of arbitrary Borel probability distributions.

However, note that for instance in pattern recognition, there might well be situations where the best kernel for a given problem is not universal. In fact, the kernel corresponds to the choice of a prior, and thus using a kernel which does *not* afford approximations of arbitrary continuous function can be very useful — provided that the functions it does not approximate are known not to be solutions of the given problem.

The situation is similar for MMD. Consider the following example: suppose we knew that the two distributions we are testing are both Gaussians (with unknown mean vectors and covariance matrices). Since the empirical means of products of input variables up to order two are sufficient statistics for the family of Gaussians, we should thus work in an RKHS spanned by products of order up to two — any higher order products contain no information about the underlying Gaussians and can therefore mislead us. It is straightforward to see that for $c > 0$, the polynomial kernel $k(x, x') = (\langle x, x' \rangle + c)^2$, with $c > 0$, does the job: it equals

$$\sum_{i,j=1}^d x_i x_j x'_i x'_j + 2c \sum_{i=1}^d x_i x'_i + c^2 = \langle \phi(x), \phi(x') \rangle,$$

where $\phi(x) = (c, \sqrt{2c}x_1, \dots, \sqrt{2c}x_d, x_i x_j | i, j = 1, \dots, d)^\top$. If we want to test for differences in higher order moments, we use a higher order kernel $k(x, x') = (\langle x, x' \rangle + c)^p$.⁵

Note, however, that this does not tell us how to choose c . With additional prior knowledge, we could further improve the odds of our test working well on small sample sizes. E.g., if we knew that the Gaussians differ mainly in their covariance structures, then we could incorporate this by choosing a small c . If the available prior knowledge is less specific, we could also sum up several MMDs by using summed kernels.

⁵ Kernels with infinite dimensional RKHS can be viewed as a nonparametric generalization where we have infinitely many sufficient statistics.

2.5 Related Methods

Various empirical methods have been proposed to determine whether two distributions are different. The first test we consider, and the simplest, is a multivariate generalization of the t-test (Hotelling, 1951), which assumes both distributions are multivariate Gaussian with unknown, identical covariance structure. This test is not model-free in the sense of MMD (and the tests described below) — indeed, it is easy to construct examples in which it fails completely (Figure 1).

Two well-established model-free univariate tests are the Kolmogorov-Smirnov statistic and the Wald-Wolfowitz runs test. Both tests are powerful in that the distribution of the test statistic is known independently of p and q for finite sample sizes, under the null hypothesis $p = q$. A generalization of the Wald-Wolfowitz runs test to the multivariate domain was proposed by Friedman and Rafsky (1979). It involves counting the number of edges in the minimum spanning tree over the aggregated data that connect points in X to points in Y . The resulting test relies on the asymptotic normality of the test statistic. The computational cost of this method using Kruskal’s algorithm is $O((m+n)^2 \log(m+n))$, although more modern methods improve on the $\log(m+n)$ term. Two possible generalizations of the Kolmogorov-Smirnov test to the multivariate case were studied by Bickel (1969); Friedman and Rafsky (1979). The approach of Friedman and Rafsky in this case again requires a minimal spanning tree, and thus has a similar cost to their multivariate runs test.

Hall and Tajvidi (2002) propose to aggregate the data as $Z = \{X, Y\}$, find the j points in Z closest to each point in X for all $j \in \{1, \dots, m\}$, count how many of these are from Y , and compare this with the number of points expected under the null hypothesis (the procedure is repeated for each point in Y wrt points in X). The test statistic is costly to compute; Hall and Tajvidi (2002) consider only tens of points in their experiments.

Another approach is to use some distance (e.g. L_1 or L_2) between estimates of the densities as a test statistic (Anderson et al., 1994; Biau and Györfi, 2005), based on the asymptotic distribution of this distance given $p = q$. One problem with the approach of Biau and Györfi (2005), however, is that it requires the space to be partitioned into a grid of bins, which becomes difficult or impossible for high dimensional problems (such as those in Section 3).

We now illustrate these tests with a simple example. In Figure 2, we compare several alternatives to the MMD-based test in distinguishing 100 samples taken from each of two normal distribution with unit variance. Results are averaged over a series of Euclidean distances between the means of both distributions, and plotted as a function of increasing dimensionality. The t-test has the highest chance of correctly rejecting the null hypothesis for low dimensions. However, for high dimensions the estimation of the sample covariance matrices is poor due to the limited sample sizes. Note that we do *not* apply the

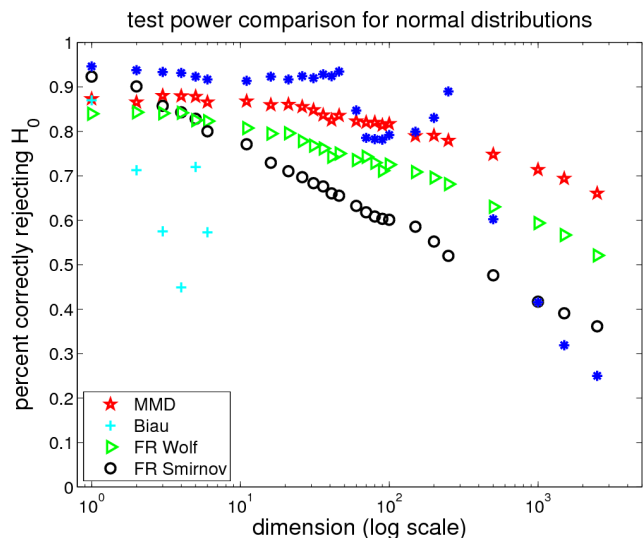


Fig. 2. Test between samples from two normal distribution with different means and unit variance based on a significance level $\alpha = 0.05$. The cumulative percentage of correctly rejecting the null hypothesis out for 80 runs and a set of Euclidean distances in means (0.25,0.3,0.4,0.5,0.6,0.7,0.8,1,2,5,10,15) is shown in dependence of the dimensionality of the normal distributions. Sample size was 100 each. MMD with Gaussian kernel was employed. Kernel size σ were obtained by maximizing MMD (for σ values within 0.25 and 20) to get the most conservative test. In case of the t-test, a ridge was added to the covariance estimate, to avoid singularity (repeatedly in steps of 0.01 until the 2-norm condition number was below 10). For the Biau test equal partitions per dimension were used, which is intractable for high dimensions. When samples from two standard normal distribution were compared, tests reject null hypothesis in the following number of trials out of 2080 (summed over all dimensions in plot, 80 runs each): 33 (MMD), 195 (t-test), 112 (FR Wolf), 83 (FR Smirnov). For Biau test: 0 out of 400 trials.

Biau & Györfi test for high dimensionalities when memory requirements force the number of partitions per dimension to be too low.

MMD performs very well and outperforms all other model-free approaches, namely the multivariate Kolmogorov-Smirnov test (FR Smirnov), the multivariate Wald-Wolfowitz runs test (FR Wolf), and the Biau & Györfi test (Biau). The comparison becomes harder for increasing dimensionality, since the sample size is fixed to 100 random vectors per distribution for all dimensions. Moreover, MMD also yields very low rejection rate of the null hypothesis, if the null hypothesis is actually true in case of testing samples from two standard normal distributions (see figure legend).

Finally, we mention that the connection between means in RKHSs and distributions has, in a less general setting, been observed before in the kernel machines field. Schölkopf and Smola (2002) point out that the empirical mean of a set of

points in an RKHS can be viewed as a Parzen windows estimate of the density underlying the data; Shawe-Taylor and Cristianini (2004) propose to use the distance to the mean as a novelty detection criterion, and provide a statistical analysis.

3 EXPERIMENTS

In this section, we present applications of MMD in data integration for bioinformatics, from microarray cross-platform comparability via cancer (subtype) diagnosis to schema matching.

3.1 Microarray Cross-Platform Comparability

Experimental scenario Microarrays as a large-scale gene expression observation tool offer a unique possibility for molecular biologists to study gene activity at a cellular level. In recent years, there have been a great number of developments in different microarray platforms, techniques and protocols, advances in these techniques, and biological and medical studies making use of these approaches. Even steps in data processing such as missing value imputation may differ. As a result, microarray data for a given problem, and the results derived from it (e.g. marker genes for a certain subtype of cancer), may vary greatly (Carter et al., 2005), both between labs and platforms. Consequently, despite an avalanche of microarray data being generated nowadays, it remains to be determined if and how to combine microarray data from different studies on the same biological subject.

Hence it is necessary to establish a statistical test of whether two microarray measurements on the same subject, obtained by two different labs or on two different platforms, can be regarded as comparable and can be used for joint data analysis. We define such a test using MMD as a statistic: If an MMD-based test rejects the null hypothesis that the microarray measurements are generated from the same distribution, then we deem them not comparable.

We test this approach on published microarray datasets from two different platforms. If our criterion is useful in practice and able to detect the limited cross-platform comparability of microarray data, then MMD should judge microarray data achieved on different platforms as being less often comparable than those found on the same platform.

Data For our first experiment, we obtained datasets from Hu et al. (2005). These data comprise gene expression levels of 6661 genes from 39 datasets from several published studies. Each dataset contains a minimum of eight microarray measurements, observed either on cDNA platforms or oligonucleotide microarray platforms.

We tried to find out via MMD if there is any statistically significant difference between the microarray results achieved on these different platforms. We compared the MMD results to the multivariate t-test and the Friedman-Rafsky multivariate Kolmogorov-Smirnov and Wald-Wolfowitz tests

Table 1. Cross-platform comparability of microarray profiles

Platforms	H_0	MMD	t-test	FR	FR
				Wolf	Smirnov
Same	accepted	100	100	80	96
Same	rejected	0	0	20	4
Different	accepted	0	100	94	70
Different	rejected	100	0	6	30

Cross-platform comparability tests on microarray level for cDNA and oligonucleotide platforms. Per sample: 50 vectors with 6661 dimensions, 100 repetitions

(denoted Smirnov and Wolf, respectively). The high dimensionality of this problem, as well as of the experiments below, prevents a comparison with the Biau-Györfi test.

We chose $\alpha = 0.05$ as the level of significance for all tests. A Gaussian kernel was employed for MMD, with $\sigma = 20$. We obtained an average performance over 100 distribution tests using 100 microarray measurements from different platforms (X being 50 cDNA measurements and Y being 50 oligonucleotide measurements), and 100 distribution tests with data from 100 microarray measurements taken from only one of the two platforms. For each test, the studies were randomly selected without replacement from the relevant measurement pools. We repeated this experiment for MMD and each of the competing methods.

Results Results are reported in Table 1, showing the number of times MMD and the other three methods deemed two samples as originating from the same distribution, on data from both identical and dissimilar platforms. While all tests usually recognize samples from the same platform as being generated by the same distribution, MMD is the only test that does not consider microarray data from different platforms as originating from the same distribution. As inter-platform comparability of microarray data is reported to be modest in many recent publications (van Ruissen et al., 2005; Carter et al., 2005; Stec et al., 2005), MMD is very successful in detecting these differences in our experiments.

3.2 Cancer and Tumor Subtype Diagnosis

Experimental scenario Besides microarray cross-platform comparability, it is interesting to examine if MMD can distinguish between the gene expression profiles of groups of people who are respectively healthy or ill, or who suffer from different subtypes of a particular cancer. Alternatively, as in the previous experiment, MMD can be employed to determine whether we should integrate two sets of observations (which might arise from different subtypes of a cancer) into one joint set, or if we should treat them as distinct classes.

When using MMD for cancer diagnosis, we test whether the microarray data at hand contain a significant level of difference

between ill and healthy patients. Conversely, when looking at cancer (or tumor) subtypes, MMD indicates whether two subtypes of cancer should be considered independently when designing a computational predictor of cancer, or if they can be assigned to one common super-class. In terms of classification methods, MMD can be used to choose whether binary (cancer/healthy) or multi-class (healthy, cancer subtype 1, . . . , cancer subtype n) classification will be more accurate when developing a diagnosis tool.

Data For our second microarray experiment, we obtained two datasets from Warnat et al. (2005). The first one is a preprocessed dataset of microarray data from prostate cancer patients, with gene expression profiles from 33 normal and tumor tissues (Welsh et al., 2001). The second one is a preprocessed dataset of 58 gene expression profiles of different tumor subtypes from Gruvberger et al. (2001). In particular, we are interested in cancer diagnosis on the data of Welsh et al. (2001) and tumor subtype diagnosis on the data of Gruvberger et al. (2001).

Cancer diagnosis

We examine whether MMD can distinguish between normal and tumor tissues, using the microarray data from a prostate cancer study by Welsh et al. (2001). Again, α was set to 0.05. Randomly choosing 100 pairs of 5 healthy and 5 cancer patients' gene expression profiles, we used MMD to test the null hypothesis that both samples were generated by the same distribution. We then did the same test for 100 randomly chosen pairs of samples of size 5, both drawn from the same tissue type (healthy or tumor). For all 200 pairs of samples, we compared our results to those of the multivariate t-test and both Friedman-Rafsky tests (Wolf and Smirnov).

Results We emphasise at the outset that for such small sample sizes, we are unlikely to find ourselves in the asymptotic regime, either in our test or in the Friedman-Rafsky tests. Indeed, any test for such a small sample size will likely rely almost exclusively on the non-overlap of the distributions generating X and Y , rather than on any higher order statistics of these distributions. That said, the classes have very little overlap for these data, as evidenced by the 99.09% classification accuracy achieved by Warnat et al. (2005) with a support vector machine.

Results are reported in Table 2. Both MMD and the Friedman-Rafsky tests are in agreement that there is a large difference between samples from cancer patients and healthy patients, and little difference within a particular class. It is notable that the t-test performs so poorly in this case, despite the difference in means being the likely distinguishing feature of the two classes. This is consistent with our Figure 2 example, however, for high dimensional spaces.

Table 2. Cancer diagnosis

Tissue	H_0	MMD	t-test Wolf	FR Smirnov	FR
Same	accepted	90	50	91	94
Same	rejected	10	50	9	6
Different	accepted	3	0	5	14
Different	rejected	97	100	95	86

Comparing samples from normal and tumor tissue (Welsh et al., 2001). H_0 is hypothesis that $p = q$. Repetitions 100, size (each) 5, dimension of sample vectors: 4344

Table 3. Tumor subtype tests

Tissue	H_0	MMD	t-test	FR Wolf	FR Smirnov
Same	accepted	92	0	94	86
Same	rejected	8	100	6	14
Different	accepted	75	0	82	87
Different	rejected	25	100	18	13

Comparing samples from different and identical tumor subtypes (Gruvberger et al., 2001). H_0 is hypothesis that $p = q$. Repetitions 100, Sample size (each) 5, dimension of sample vectors: 2166

Tumor subtype diagnosis

We performed the same experiment as described before for tumor subtype diagnosis on data from Gruvberger et al. (2001). Classes are now estrogen receptor positive tumors and estrogen receptor negative tumors, but otherwise we followed the procedure outlined above.

Results We report results in Table 3. Unlike the previous experiment, both MMD and the Friedman-Rafsky tests prefer to accept the null hypothesis that both samples are generated by the same distribution, when the tissue types are different. In other words, all three tests seem to have difficulties in finding discrepancies between samples from different tumor subtypes in this case. This is consistent with previous results by Warnat et al. (2005) showing that separation of these tissue types via SVM classification on this dataset has an unusually high error rate (approximately 20%), while many other cancer studies have error rates of 5% and less. In other words, the two classes have some overlap, and tests based on our very small sample size are not able to distinguish the distributions. It remains an open question, and topic of current research, as to whether our test can distinguish tissue types from such hard-to-classify distributions when the sample size is larger.

3.3 Schema Matching on Molecular Graph Data

Experimental scenario Classifying biological data into ontologies or taxonomies is the central step in structuring and

Table 4. Data-content based schema matching

Test	EC 1	EC 2	EC 3	EC 4	EC 5	EC 6
EC 1	0	50	45	50	50	50
EC 2	50	0	50	50	50	50
EC 3	48	50	0	50	50	50
EC 4	50	50	50	0	50	50
EC 5	50	50	50	50	0	50
EC 6	50	50	50	50	50	0

Data-content based schema matching for $\alpha = 0.01$. Numbers indicate how often null hypothesis($p=q$) was rejected.

organization of these data. However, different studies may use different ontologies, resulting in the need to find correspondences between two ontologies. Here we employ MMD to discover matching terms in two ontologies using the data entries associated with these terms.

We study the following scenario: Two researchers have each dealt with 300 enzyme protein structures. These two sets of 300 proteins are disjoint, i.e. there is no protein studied by both researchers. They have assigned the proteins to six different classes according to their enzyme activity. However, both have used different protein function classification schemas for these six classes and are not sure which of these classes correspond to each other.

To find corresponding classes, MMD can be employed. We obtained 600 proteins modeled as graphs from Borgwardt et al. (2005), and randomly split these into two subsets A and B of 300 proteins each, such that 50 enzymes in each subset belong to each one of the six EC top level classes. We then computed MMD for all pairs of EC classes from subset A and subset B to check if the null hypothesis is rejected or accepted. To compute MMD, we employed the protein random walk kernel function for protein graphs, following Borgwardt et al. (2005). We compared all pairs of classes via MMD, and repeated the experiment 50 times.

Results For a significance level of $\alpha = 0.05$, MMD rejected the null hypothesis that both samples are from the same distribution whenever enzymes from two different EC classes were compared. When enzymes from the same EC classes were compared, MMD accepted the null hypothesis. MMD thus achieves error-free data-based schema matching here.

We checked whether the same good results were found even for a higher significance level of $\alpha = 0.01$. We report results in Table 4. This time, in 7 comparisons out of 1800 comparisons the null hypothesis is incorrectly accepted, whereas in all other cases, the correct decision is taken. Hence even for the high significance level of $\alpha = 0.01$ MMD is very accurate.

In addition to these promising results, note that although we consider the basic case of 1:1 correspondence between classes in our experiment, the fact that MMD uses the kernel trick

allows for even more powerful approaches to data-content based schema matching. As kernels are closed under addition and pointwise multiplication, we can test complex correspondences between different classes as well, where one class in schema A corresponds to a combination of classes in schema B. Schema matching for complex correspondences via MMD is a topic of current research.

4 DISCUSSION AND CONCLUSIONS

In this paper, we have presented, to the best of our knowledge, the first principled statistical test for distribution testing and data integration of structured objects, using the Maximum Mean Discrepancy (MMD) as a test statistic. MMD makes use of kernels and is hence not limited to vectorized data. As a consequence, MMD is not only applicable to a wide range of problems in molecular biology, but also to all common data types in bioinformatics, such as strings and graphs. All kernels for biological data, defined usually for classification tasks, can now be employed for distribution testing. Amongst others, these include kernels on protein sequences, protein structures, and microarray time series (Schölkopf et al., 2004).

MMD is easy to implement, memory-efficient, and fast to compute. In our experiments, it often outperformed competing methods (provided the latter were applicable at all, i.e., on vectorial data). We applied our MMD-based test to microarray cross-platform comparability, cancer diagnosis, and data-content based schema matching.

We believe that MMD could also be employed to validate computational simulations of biological processes. If wetlab experiments and simulations generate results and predictions that MMD deems comparable, it is likely that the simulator has produced realistic predictions. This validation procedure will become increasingly relevant as more model-based simulations of microarray data become available (den Bulcke et al., 2006). MMD could also be used for keyplayer gene selection from microarray data. For example, this type of feature selection could be employed to find genes that are involved in cancer outbreak when looking at gene expression profiles from healthy and cancer patients. MMD would be applied to subsets of genes from two classes of microarrays to find the subset that maximizes the probability that the two classes arise from different distributions. These genes should be studied experimentally in more detail. If, however, MMD cannot find any subset of genes that results in significant differences between healthy and cancer patients, then this might serve as an indicator that the microarrays did not contain the essential genes involved in cancer progress.

ACKNOWLEDGEMENTS

The authors are deeply grateful to Patrick Warnat (DKFZ, Heidelberg) and Haiyan Hu (USC) for providing the datasets for the experiments. This work was supported in part by National ICT Australia and by the German Ministry for Education,

Science, Research and Technology (BMBF) under grant no. 031U112F within the BFAM (Bioinformatics for the Functional Analysis of Mammalian Genomes) project which is part of the German Genome Analysis Network (NGFN). National ICT Australia is funded through the Australian Government's *Backing Australia's Ability* initiative, in part through the Australian Research Council.

REFERENCES

- Anderson, N., Hall, P., and Titterton, D. (1994). Two-sample test statistics for measuring discrepancies between two multivariate probability density functions using kernel-based density estimates. *Journal of Multivariate Analysis*, 50:41–54.
- Biau, G. and Györfi, L. (2005). On the asymptotic properties of a nonparametric l_1 -test statistic of homogeneity. *IEEE Transactions on Information Theory*, 51(11):3965–3973.
- Bickel, P. (1969). A distribution free version of the smirnov two sample test in the p-variate case. *The Annals of Mathematical Statistics*, 40(1):1–23.
- Borgwardt, K. M., Ong, C. S., Schonauer, S., Vishwanathan, S. V., Smola, A. J., and Kriegel, H. P. (2005). Protein function prediction via graph kernels. *Bioinformatics*, 21(Suppl 1):i47–i56.
- Carter, S. L., Eklund, A. C., Mecham, B. H., Kohane, I. S., and Szallasi, Z. (2005). Redefinition of affymetrix probe sets by sequence overlap with cDNA microarray probes reduces cross-platform inconsistencies in cancer-associated gene expression measurements. *BMC Bioinformatics*, 6(1):107.
- Casella, G. and Berger, R. (2002). *Statistical Inference*. Duxbury, Pacific Grove, CA, 2nd edition.
- Cuturi, M., Fukumizu, K., and Vert, J.-P. (2005). Semigroup kernels on measures. *JMLR*, 6:1169–1198.
- den Bulcke, T. V., Leemput, K. V., Naudts, B., van Remortel, P., Ma, H., Verschoren, A., Moor, B. D., and Marchal, K. (2006). Syntren: a generator of synthetic gene expression data for design and analysis of structure learning algorithms. *BMC Bioinformatics*, 7(1):43.
- Dudley, R. M. (2002). *Real analysis and probability*. Cambridge University Press, Cambridge, UK.
- Fortet, R. and Mourier, E. (1953). Convergence de la réparation empirique vers la réparation théorique. *Ann. Scient. École Norm. Sup.*, 70:266–285.
- Friedman, J. and Rafsky, L. (1979). Multivariate generalizations of the wald-wolfowitz and smirnov two-sample tests. *The Annals of Statistics*, 7(4):697–717.
- Gärtner, T., Flach, P., and Wrobel, S. (2003). On graph kernels: Hardness results and efficient alternatives. In Schölkopf, B. and Warmuth, M. K., editors, *Sixteenth Annual Conference on Computational Learning Theory and Seventh Kernel Workshop, COLT*. Springer.
- Gruvberger, S., Ringner, M., Chen, Y., Panavally, S., Saal, L. H., Borg, A., Ferno, M., Peterson, C., and Meltzer, P. S. (2001). Estrogen receptor status in breast cancer is associated with remarkably distinct gene expression patterns. *Cancer Res*, 61(16):5979–5984.
- Hall, P. and Tajvidi, N. (2002). Permutation tests for equality of distributions in high-dimensional settings. *Biometrika*, 89(2):359–374.
- Hein, M. and Bousquet, O. (2005). Hilbertian metrics and positive definite kernels on probability measures. In Ghahramani, Z. and Cowell, R., editors, *Proc. of AI & Statistics*, volume 10.
- Hoeffding, W. (1948). A class of statistics with asymptotically normal distribution. *The Annals of Mathematical Statistics*, 19(3):293–325.
- Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30.
- Hotelling, H. (1951). A generalized t test and measure of multivariate dispersion. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 23–41.
- Hu, H., Yan, X., Huang, Y., Han, J., and Zhou, X. J. (2005). Mining coherent dense subgraphs across massive biological networks for functional discovery. *Bioinformatics*, 21 Suppl 1:i213–i221.
- Marshall, E. (2004). Getting the noise out of gene arrays. *Science*, 306(5696):630–631.
- Schölkopf, B. and Smola, A. (2002). *Learning with Kernels*. MIT Press, Cambridge, MA.
- Schölkopf, B., Tsuda, K., and Vert, J.-P. (2004). *Kernel Methods in Computational Biology*. MIT Press, Cambridge, MA.
- Serfling, R. (1980). *Approximation Theorems of Mathematical Statistics*. Wiley, New York.
- Shawe-Taylor, J. and Cristianini, N. (2004). *Kernel Methods for Pattern Analysis*. Cambridge University Press, Cambridge, UK.
- Shi, L., Tong, W., Fang, H., Scherf, U., Han, J., Puri, R. K., Frueh, F. W., Goodsaid, F. M., Guo, L., Su, Z., Han, T., Fuscoe, J. C., Xu, Z. A., Patterson, T. A., Hong, H., Xie, Q., Perkins, R. G., Chen, J. J., and Casciano, D. A. (2005). Cross-platform comparability of microarray technology: intra-platform consistency and appropriate data analysis procedures are essential. *BMC Bioinformatics*, 6 Suppl 2:S12.
- Smola, A., Gretton, A., and Borgwardt, K. (2006). Maximum mean discrepancy. Technical Report NICTA-SML-06-001, National ICT Australia.
- Stec, J., Wang, J., Coombes, K., Ayers, M., Hoersch, S., Gold, D. L., Ross, J. S., Hess, K. R., Tirrell, S., Linette, G., Hortobagyi, G. N., Symmans, W. F., and Pusztai, L. (2005). Comparison of the predictive accuracy of dna array-based multigene classifiers across cDNA arrays and affymetrix genechips. *J Mol Diagn*, 7(3):357–367.
- Steinwart, I. (2002). On the influence of the kernel on the consistency of support vector machines. *Journal of Machine Learning Research*, 2:67–93.
- van Ruissen, F., Ruijter, J. M., Schaaf, G. J., Asgharnegad, L., Zwijnenburg, D. A., Kool, M., and Baas, F. (2005). Evaluation of the similarity of gene expression data estimated with sage and affymetrix genechips. *BMC Genomics*, 6:91.
- Wernat, P., Eils, R., and Brors, B. (2005). Cross-platform analysis of cancer microarray data improves gene expression based classification of phenotypes. *BMC Bioinformatics*, 6:265.
- Welsh, J. B., Sapinoso, L. M., Su, A. I., Kern, S. G., Wang-Rodriguez, J., Moskaluk, C. A., Frierson HF, J., and Hampton, G. M. (2001). Analysis of gene expression identifies candidate markers and pharmacological targets in prostate cancer. *Cancer Res*, 61(16):5974–5978.