

Principles of Multi-Kernel Data Mining

Vadim Mottl¹, Olga Krasotkina¹, Oleg Seredin¹, and Ilya Muchnik²

¹ Computing Center of the Russian Academy of Sciences
Vavilov St., 40, 117968 Moscow, Russia
vmottl@yahoo.com

² DIMACS, Rutgers University, P.O. Box 8018, Piscataway, NJ 08855, USA
muchnik@dimacs.rutgers.edu

Abstract. The scientific community has accumulated an immense experience in processing data represented in finite-dimensional linear spaces of numerical features of entities, whereas the kit of mathematical instruments for dissimilarity-based processing of data in metric spaces representing distances between entities, for which sufficiently informative features cannot be found, is much poorer. In this work, the problem of embedding the given set of entities into a linear space with inner product by choosing an appropriate kernel function is considered as the major challenge in the featureless approach to estimating dependences in data sets of arbitrary kind. As a rule, several kernels may be heuristically suggested within the bounds of the same data analysis problem. We treat several kernels on a set of entities as Cartesian product of the respective number of linear spaces, each supplied with a specific kernel function as a specific inner product. The main requirement here is to avoid discrete selection in eliminating redundant kernels with the purpose of achieving acceptable computational complexity of the fusion algorithm.

1 Introduction

The problem of finding empirical dependences $y(\omega) : \Omega \rightarrow Y$ in a set of entities of arbitrary kind $\omega \in \Omega$ is one of the glowing problems of modern data mining. Let a given data set be the set of experimentally measured values of a characteristic $y_j = y(\omega_j) \in Y$ within an accessible subset of entities $\Omega^* = \{\omega_1, \dots, \omega_N\} \subset \Omega$. It is required to continue this function onto the entire set Ω for it would be possible to estimate this characteristic $\hat{y}(\omega)$ for entities $\omega \in \Omega \setminus \Omega^*$ not represented in the original (training) data set [1,2]. In particular, if $y(\omega)$ takes values from a finite set, for instance, $Y = \{-1, 1\}$, the problem is usually called the pattern recognition problem, and in the case of a real-valued characteristic $Y = \mathbb{R}$ it is referred to as the problem of regression estimation.

This work is supported by the Russian Foundation for Basic Research (Grants 02-01-00107 and 05-01-00679), Grant of the President of the Russian Federation for young scientists No. MK-3173.2004.09 (O. Seredin), INTAS Grant 04-77-7347, and NSF Grant CCR 0325398 (I. Muchnik).

It is clear that the problem of function continuation is meaningless until some assumptions are taken about the relations between the values $y(\omega)$ and other characteristics of entities $\omega \in \Omega$ that are more accessible to observation than the goal characteristic. There exist many practical problems of data analysis, including pattern recognition and regression estimation, in which it is relatively easy to evaluate the degree of dissimilarity of any pair of entities. The modern machine learning theory is based on the so-called compactness hypothesis, which consists in the assumption that if two entities are close to each other in the sense of an appropriate metric then so are also, in most cases, the respective values of the goal characteristic. This fact, actually, underlies the featureless (relational, similarity-based) approach to data analysis proposed by R. Duin and his colleagues [3,4,5]. In the featureless situation, a natural mathematical model of the general set of entities is a metric space, in which the compactness hypothesis can be expressed directly in terms of the given metric.

At the same time, the mathematically most advanced methods of machine learning essentially exploit the assumption that the universe of entities can be represented as a linear space. As the simplest instrument of introducing linear operations in the set of entities $\omega \in \Omega$, the vector of some observable numerical features $\mathbf{x}(\omega) \in \mathbb{R}^n$ was traditionally considered, and the Euclidean metric produced by it $\rho(\omega', \omega'') = \|\mathbf{x}(\omega') - \mathbf{x}(\omega'')\|$ served as the basis of function continuation in respective machine learning techniques.

It became apparent soon that what immediately determines the result of training is the configuration of the training-set points, represented in \mathbb{R}^n by their pair-wise inner products $(\omega' \cdot \omega'') = \mathbf{x}^T(\omega')\mathbf{x}(\omega'')$, rather than the values of features. This observation resulted in the potential function method of machine learning [2], which later was named the kernel method [1]. The idea of a kernel $K(\omega', \omega'')$ consists in understanding it as inner product of two entities $K(\omega', \omega'') = (\omega' \cdot \omega'')$ in a linear space, maybe, a hypothetical one. If a kernel function $K(\omega', \omega'')$ is defined in an arbitrary set of entities Ω , it produces a Euclidean metric in it

$$\rho(\omega', \omega'') = [K(\omega', \omega') + K(\omega'', \omega'') - 2K(\omega', \omega'')]^{1/2} \quad (1)$$

which expresses a specific compactness hypothesis without the intervening notion of features.

There is usually much freedom in measuring similarity or dissimilarity of entities, and, thus, several heuristic kernels may be heuristically suggested within the bounds of the same data analysis problem. However, the choice of features $x_i(\omega) \in \mathbb{R}$, each of which defines, actually, a simplest kernel $K_i(\omega', \omega'') = x_i(\omega')x_i(\omega'')$, is also ever heuristic. The aim of this work is to study the ways of fusing the given set of kernels and to organize, thereby, a concurrence of several compactness hypotheses in finding empirical regularities in the given set of entities. The main requirement here is to avoid discrete selection of kernels with the purpose of achieving acceptable computational complexity of the fusion algorithm. We use here the main idea of embedding the discrete problem of choosing a subset into a continuous problem of finding optimal nonnegative weights as-

signed to the elements of the initial set. This idea was originally proposed in [6] as a means of constructing Relevance Vector Machines (RVM).

2 The linear space produced by a kernel

A kernel $K(\omega', \omega'')$ on a set of entities of arbitrary kind $\omega \in \Omega$ can be defined as a real-valued function $\Omega \times \Omega \rightarrow \mathbb{R}$ possessing two principal properties – symmetry $K(\omega', \omega'') = K(\omega'', \omega')$ and positive semi-definiteness of the matrix $[K(\omega_i, \omega_j); i, j = 1, \dots, m]$ for any finite collection of entities $\{\omega_1, \dots, \omega_m\} \subset \Omega$. The function $\rho(\omega', \omega'')$ (1) produced by a kernel is a metric [7], and, so, the set of entities Ω supplied with a kernel function becomes a metric space.

Any kernel function $K(\omega', \omega'')$ allows for mentally embedding the set Ω into a real linear space with inner product $\Omega \subseteq \tilde{\Omega}$. The null element $\phi \in \tilde{\Omega}$ and linear operations $(\omega' + \omega'' : \tilde{\Omega} \times \tilde{\Omega} \rightarrow \tilde{\Omega})$ and $(c\omega : \mathbb{R} \times \tilde{\Omega} \rightarrow \tilde{\Omega})$ are defined in $\tilde{\Omega}$ in a special way, whereas the role of inner product is played by the kernel function itself $(\omega', \omega'') = K(\omega', \omega'')$.

As the basis for introducing linear operations in the extended set $\tilde{\Omega}$, serves the notion of coaxiality of elements in a metric space [7]. Let $\langle \omega', \omega'' \rangle$ be an ordered pair of elements $\omega', \omega'' \in \tilde{\Omega}$. We shall say that the element $\omega \in \tilde{\Omega}$ is coaxial to the pair $\langle \omega', \omega'' \rangle$ with coefficient $c \in \mathbb{R}$ if $\rho(\omega', \omega) = |c| \rho(\omega', \omega'')$ and $\rho(\omega'', \omega) = |1 - c| \rho(\omega', \omega'')$. This fact will be denoted by the symbol $\omega = \text{coax}(\langle \omega', \omega'' \rangle; c)$. The triangle inequality turns into equality for any three coaxial elements ω', ω'' and ω .

A metric space will be said to be unboundedly convex if for any ordered pair $\langle \alpha', \alpha'' \rangle$ and any $c \in \mathbb{R}$ it contains at least one element coaxial to this pair with coefficient c . It is proved in [7] that the coaxial element is unique if the metric forming an unboundedly convex metric space is produced by a kernel function (1). Such metric spaces are called Euclidean metric spaces. It is assumed that the given set of entities Ω is embedded into a greater set $\Omega \subseteq \tilde{\Omega}$ in which a kernel function is defined and which is, so, a Euclidean metric space.

It is possible to define linear operations in the Euclidean metric space $\tilde{\Omega}$ in the following way (see [7] for details):

- the null element is a hypothetical element $\phi \in \tilde{\Omega}$ for which $K(\phi, \phi) = 0$;
- multiplying by real coefficient $c\omega = \text{coax}(\langle \phi, \omega \rangle; c)$;
- summation $\omega' + \omega'' = 2\text{coax}(\langle \omega', \omega'' \rangle; 1/2)$;

- inner product and norm $(\omega' \cdot \omega'') = K(\omega', \omega'')$, $\|\omega\| = [K(\omega, \omega)]^{1/2}$.

It is just this system of linear operations which is produced in the extended set $\tilde{\Omega}$ by a kernel function defined in the original set of entities $\Omega \subseteq \tilde{\Omega}$.

The dimensionality of the linear space $\tilde{\Omega}$ is the maximum number of elements $\{\omega_1, \dots, \omega_m\} \subset \tilde{\Omega}$ for which the matrix $[K(\omega_i, \omega_j); i, j = 1, \dots, m]$ can be positive definite. We do not study here the question of the dimensionality of this space, which may be finite or infinite, but this issue is extremely important for the generalization performance of the decision rules inferred from a training set.

3 The class of linear decision rules in the linear space produced by a kernel function

The convenience of a kernel function as a means of measuring dissimilarity of any two entities by the respective Euclidean metric (1) consists in that it involves the notion of a linear function $y(\omega): \Omega \rightarrow \mathbb{R}$ in the set of entities of any kind. This circumstance makes it possible to develop very simple algorithms of estimating dependencies between, generally speaking, arbitrary entities by exploiting, in the featureless situation, practically all known methods which had been worked up for linear spaces.

In this Section, we consider the commonly adopted class of kernel-based decision rules as a class of linear functions in the extended set of entities $\tilde{\Omega}$ supplied with linear operations and inner product produced by a continuation of the given kernel function. The class of linear functions in $\tilde{\Omega}$ is defined by two parameters $\vartheta \in \tilde{\Omega}$ and $b \in \mathbb{R}$

$$y(\omega | \vartheta, b) = K(\vartheta, \omega) + b, \omega \in \Omega. \quad (2)$$

We shall call parameter ϑ the direction element of the linear function.

If the real value of the linear function is immediately treated as the goal characteristic of an entity, the choice of parameters $\vartheta \in \tilde{\Omega}$ and $b \in \mathbb{R}$ determines a regression dependence. If the sign of the linear function is understood as the goal characteristic, the parameters specify a classification of the set of entities into two classes: $y(\omega) = K(\vartheta, \omega) + b > 0 \rightarrow$ class 1, $y(\omega) \leq 0 \rightarrow$ class 2.

Such a way of specifying a linear function may appear nonconstructive because it involves a hypothetical element of a linear space $\vartheta \in \tilde{\Omega}$ as direction element in (2), which is nothing else than product of our imagination. But when solving the problem of inferring the regression dependence or decision rule of pattern recognition from a training set $\{(y_j, \omega_j); j = 1, \dots, N\}$ by the principles of Support Vector Machines [1] or Relevance Vector Machines [6], the only reasonable choice of ϑ will be a linear combination of really existing objects $\hat{\vartheta} = \sum_{j=1}^N a_j \omega_j$ in accordance with the linear operations induced in the extended set $\tilde{\Omega}$ by the kernel function $K(\omega', \omega'')$ [8]. As inner product in $\tilde{\Omega}$, the

kernel function is linear with respect to its arguments, hence, the linear function resulting from training will include the values of the kernel function only for objects existing in reality $\hat{y}(\omega) = \sum_{j=1}^N a_j K(\omega_j, \omega)$.

4 Cartesian product of linear spaces produced by several kernel functions

It is natural to expect that different experts skilled in the specific knowledge area will propose different kernel function. The main idea of this work is to shift the burden of the final choice onto the training algorithm by concurrently fusing the given set of heuristically chosen kernels.

Let $K_i(\omega', \omega'')$, $i = 1, \dots, n$, be the kernel functions defined on the same set of entities $\omega \in \Omega$ by different experts. These kernel functions embed the set Ω into different linear spaces $\Omega \subset \tilde{\Omega}_i$, $i = 1, \dots, n$, with different inner products and, respectively, different linear operations. It is convenient to treat the n linear spaces jointly as Cartesian product

$$\tilde{\tilde{\Omega}} = \Omega_1 \times \dots \times \Omega_n = \{ \bar{\omega} = \langle \omega_1, \dots, \omega_n \rangle : \omega_i \in \tilde{\Omega}_i \} \quad (3)$$

formed by ordered n -tuples of elements from $\tilde{\Omega}_1, \dots, \tilde{\Omega}_n$. The kernel function (i.e. inner product) in this linear space can be defined as the sum of the kernel functions (inner products) of the corresponding components in any two n -tuples $\bar{\omega}' = \langle \omega'_1, \dots, \omega'_n \rangle$ and $\bar{\omega}'' = \langle \omega''_1, \dots, \omega''_n \rangle$:

$$K(\bar{\omega}', \bar{\omega}'') = \sum_{i=1}^n K_i(\omega'_i, \omega''_i), \bar{\omega}', \bar{\omega}'' \in \tilde{\tilde{\Omega}}. \quad (4)$$

The dimensionality of the combined linear space $\tilde{\tilde{\Omega}}$ (3) will not exceed the sum of dimensionalities of the particular linear spaces.

A really existing entity $\omega \in \Omega$ will be represented by its n -fold repetition $\bar{\omega} = \langle \omega, \dots, \omega \rangle \in \tilde{\tilde{\Omega}}$. Then any real-valued linear function $\Omega \rightarrow \mathbb{R}$ is specified by the choice of parameters $\bar{\vartheta} \in \tilde{\tilde{\Omega}}$ and $b \in \mathbb{R}$

$$y(\omega) = K(\bar{\vartheta}, \bar{\omega}) + b = \sum_{i=1}^n K_i(\vartheta_i, \omega) + b, \quad (5)$$

where $\bar{\vartheta}$ is a combination of hypothetical elements of particular linear spaces $\bar{\vartheta} = \langle \vartheta_1, \dots, \vartheta_n \rangle$, $\vartheta_i \in \tilde{\Omega}_i$, produced by particular kernel functions $K_i(\omega', \omega'')$ in $\tilde{\Omega}_i$.

Thus, to define a numerical dependence over a set of entities of any kind by combining several kernel functions $K_i(\omega', \omega'')$, we have, first of all, to choose, as parameters, one element in each of linear spaces $\vartheta_i \in \tilde{\Omega}_i$ into which the kernel functions embed the original set $\Omega \subseteq \tilde{\Omega}_i$. It should be marked that the less the norm of the i th parameter in

its linear space $\|\vartheta_i\|^2 = K_i(\vartheta_i, \vartheta_i)$, the less the influence of the respective summand on the value of the function (5). If $K(\vartheta_i, \vartheta_i) \rightarrow 0$, i.e. $\vartheta_i \cong \phi_i \in \tilde{\Omega}_i$, the i th kernel will practically not affect the goal function.

This means that the parametric family of numerical functions (5) implies also an instrument of emphasizing “adequate” kernels with respect to the available observations and suppressing “inadequate” ones. Which kernels should be considered as adequate is the key question for providing a good generalization performance of the decision rule when it is applied to entities not represented in the training set.

5 Fusion of kernel functions

If the total dimensionality of the combined extended linear space $\tilde{\Omega}$ (3) is greater than the number of entities in the training set $\{(\omega_j, y_j); y_j \in \mathbb{R}, j = 1, \dots, N\}$ or $\{(\omega_j, g_j); g_j \in \{-1, 1\}, j = 1, \dots, N\}$ there always exist linear functions (5) that exactly reproduce the trainer’s data. Following the widely adopted principle [1], we shall prefer the function with the minimum norm of the direction element under the constraints of the training set:

$$\begin{cases} \|\bar{\vartheta}\|^2 \rightarrow \min, \bar{\vartheta} = \langle \vartheta_1, \dots, \vartheta_n \rangle \in \tilde{\Omega}, \\ \sum_{i=1}^n K_i(\vartheta_i, \omega_j) + b = y_j \\ \text{or } g_j \sum_{i=1}^n K_i(\vartheta_i, \omega_j) + b \geq \text{const.} \end{cases} \quad (6)$$

However, the norm in $\tilde{\Omega}$ can be measured in several ways. The simplest version of norm follows from (4)

$$\|\bar{\vartheta}\|^2 = \sum_{i=1}^n K_i(\vartheta_i, \vartheta_i), \quad (7)$$

but any linear combination of kernel functions with nonnegative coefficients also possesses all the properties of norm $\|\bar{\vartheta}\|^2 = \sum_{i=1}^n (1/r_i) K_i(\vartheta_i, \vartheta_i)$. In this case, the criterion (6) will try to avoid kernels with small r_i . If $r_i = 0$, the respective kernel does not participate in forming the goal function.

The idea of adaptive training consists in jointly inferring the direction elements ϑ_i and the weights r_i from the training set by additionally penalizing large weights [6]:

$$\begin{cases} \sum_{i=1}^n [(1/r_i) K_i(\vartheta_i, \vartheta_i) + \log r_i] \rightarrow \min(\vartheta_i, r_i), \\ \sum_{i=1}^n K_i(\vartheta_i, \omega_j) + b = y_j \\ \text{or } g_j \sum_{i=1}^n K_i(\vartheta_i, \omega_j) + b \geq \text{const}, j = 1, \dots, N. \end{cases} \quad (8)$$

This adaptive training criterion displays a pronounced tendency to emphasize the kernel functions which are “adequate” to the trainer’s data and to suppress up to negligibly small values the weights r_i at “redundant” ones.

The reasoning for the adaptive training criterion (8) is the view on the unknown direction elements $\mathfrak{G}_i \in \tilde{\Omega}_i$ in each of the linear spaces $\tilde{\Omega}_i$ as hidden independent random variables whose mathematical expectations coincide with the respective null elements $M(\mathfrak{G}_i) = \phi_i \in \tilde{\Omega}_i$. The parameter r_i has the sense of the unknown mean-square distance of the random direction element from the null element in the sense of metric (1). Then (8) is equivalent to finding the joint maximum-likelihood estimate of the variables $\mathfrak{G}_1, \dots, \mathfrak{G}_n$ and their variances r_1, \dots, r_n under the additional assumption that the dimensionality of each of linear spaces $\tilde{\Omega}_i$ is, maybe, very large but finite, and the respective direction element is normally distributed in it.

Since \mathfrak{G} is element of an abstract linear space but not a vector, for minimizing the Lagrangian of the respective constrained optimization problem (8) we have to use the notion of Frechet differential instead of that of gradient [9]. The Frechet differential of a real-valued function over a linear space is element of this space: $\nabla_{\mathfrak{G}} K(\mathfrak{G}, \omega) = \omega$, $\nabla_{\mathfrak{G}} K(\mathfrak{G}, \mathfrak{G}) = 2\mathfrak{G}$. It can be shown that the following iterative procedure solves both regression estimation and pattern recognition problem:

$$\mathfrak{G}_i^{(k)} = r_i^{(k-1)} \sum_{j=1}^N \lambda_j^{(k)} \omega_j \quad \text{or} \quad \mathfrak{G}_i^{(k)} = r_i^{(k-1)} \sum_{j=1}^N \lambda_j^{(k)} \mathbf{g}_j \omega_j, \quad (9)$$

$$r_i^{(k)} = (r_i^{(k-1)})^2 \sum_{j=1}^N \sum_{l=1}^N K_i(\omega_j, \omega_l) \lambda_j^{(k)} \lambda_l^{(k)}, \quad (10)$$

where the real numbers $\lambda_1^{(k)}, \dots, \lambda_N^{(k)}$ are the Lagrange multipliers (nonnegative in the case of pattern recognition) found as solutions of the respective dual problem. Updating the constant $b^{(k)}$ doesn’t offer any difficulty.

As we see, the abstract variables $\mathfrak{G}_i^{(k)} \in \tilde{\Omega}_i$ (9) are linear combinations of the entities of the training set in the sense of linear operations induced by the kernel functions as inner products in the respective linear spaces. Substitution of (9) and (10) into (5) eliminates $\mathfrak{G}_i^{(k)}$ and gives the completely constructive estimate of the sought function, respectively, for regression estimation and pattern recognition:

$$\hat{y}^{(k)}(\omega) = \sum_{j=1}^N \lambda_j^{(k)} \sum_{i=1}^n r_i^{(k)} K_i(\omega_j, \omega) + b^{(k)}, \quad (11)$$

$$\hat{y}^{(k)}(\omega) = \sum_{j=1}^N \lambda_j^{(k)} \mathbf{g}_j \sum_{i=1}^n r_i^{(k)} K_i(\omega_j, \omega) + b^{(k)} \begin{cases} > 0, \\ < 0. \end{cases}$$

As a rule, the process converges in 10-15 steps and displays a pronounced tendency to suppressing the weights at “redundant” kernel functions $r_i \rightarrow 0$ along with emphasizing $r_i \gg 0$ the kernel functions which are “adequate” to the trainer’s data. This fact pro-

vides a computationally effective selection of kernel functions without straightforward discrete choice of their subsets.

6 A particular case: Feature selection as kernel fusion

There is no insurmountable barrier between the featureless kernel-based way of forming parametric families of numerical functions on a set of entities of any kind and the usual parametric family of linear functions on the set of entities represented by vectors of their numerical features. The latter way is particular case of the former one.

Indeed, a numerical feature $x(\omega): \Omega \rightarrow \mathbb{R}$ is equivalent to the simplest kernel function in the form of product $K(\omega', \omega'') = x(\omega')x(\omega'')$ that embeds the set of entities into a one-dimensional linear space $\Omega \subseteq \tilde{\Omega}$. Respectively, a vector of features $\mathbf{x}(\omega) = [x_1(\omega) \cdots x_n(\omega)]$ gives n kernel functions at once $K_i(\omega', \omega'') = x_i(\omega')x_i(\omega'')$ and n versions of such an embedding $\Omega \subseteq \tilde{\Omega}_i$. The choice of one entity in each of these spaces $\mathfrak{G}_i \in \tilde{\Omega}_i$, $i=1, \dots, n$, namely, n real numbers $(x_1(\mathfrak{G}_1) \cdots x_n(\mathfrak{G}_n)) \in \mathbb{R}^n$, along with a numerical constant $b \in \mathbb{R}$ specifies a linear function on the set of entities: $y(\omega) = \sum_{i=1}^n K_i(\mathfrak{G}_i, \omega) + b = \sum_{i=1}^n a_i x_i(\omega) + b$ where $a_i = x_i(\mathfrak{G}_i)$.

The less the i th coefficient, i.e. the norm of the i th imaginary entity $\|\mathfrak{G}_i\| = x_i(\mathfrak{G}_i)$, the less is the contribution of this feature $x_i(\omega)$ to the value of the function.

7 Experimental results

As the essence of feature selection is shown to be the same as that of kernel fusion, we tested the proposed approach, for obviousness sake, on a set $\{(\mathbf{x}_j, y_j); j=1, \dots, N\}$ of $N=300$ pairs consisting of randomly chosen feature vectors $\mathbf{x}_j \in \mathbb{R}^n$, $n=100$, and scalars obtained by the rule $y_j = a_1 x_{j,1} + a_2 x_{j,2} + \xi_j$ with $a_1 = a_2 = 1$ and ξ_j as normal white noise with zero mean value and some variance σ^2 . So, only $n'=2$ features of $n=100$ were rational in the simulated data. In the experiment with regression estimation this set was taken immediately, whereas for the experiment with pattern recognition we took the set $\{(\mathbf{x}_j, \mathbf{g}_j); j=1, \dots, N\}$ where $\mathbf{g}_j = -1$ if $y_j < 0$ and $\mathbf{g}_j = 1$ if $y_j \geq 0$.

In both experiments, we randomly chose $N_{tr}=20$ pairs for training. So, the size of the training set was ten times greater than the number of rational features, but five times less than the full dimensionality of the feature vector. The remaining $N_{test}=280$ pairs we used as the test set.

The comparative results of training with equal weights at features (6)-(7) and with adaptive weights (8) are presented in the following two tables:

Regression estimation		
Error rate: ratio of the root-mean-square error in the test set to the actual root variance of the observation noise σ		
Feature set	Training procedure	
	equal weights	adaptive weights
2 rational features	1.01	inapplicable
all 100 features	166.75	2.16

Pattern recognition		
Error rate: misclassification percentage in the test set		
Feature set	Training procedure	
	equal weights	adaptive weights
2 rational features	0.36%	inapplicable
all 100 features	26.8%	0.36%

As was expected, the classical training criterion with equal weights shows a drastic increase in the error rate in both cases when confusing features (i.e. confusing kernel functions) participate in training. At the same time, the error rate with weights adaptation is little sensitive to the presence of purely noisy features. In both experiments, the weights at redundant features turned practically into computer zeros after 10 iterations.

8 Conclusions

A numerical feature, when assigned to entities of a certain kind, embeds the set of these entities into a one-dimensional linear space. The essence of assigning a kernel function in a set of entities is also embedding it into a hypothetical linear space through the notion of coaxiality of elements of a Euclidean metric space.

The important difference is that the dimensionality of the space induced by a kernel function will be, as a rule, greater than one, if not infinite at all. The main point of the way of fusing several kernels is the idea to consider the Cartesian product of the respective linear spaces, just as the multidimensional feature space formed by a vector of features is the Cartesian product of the respective one-dimensional ones.

Thus, treating the universal set of “all feasible” entities as a linear space practically wipes out the difference between a set of kernels and a set of features and, so, between the featureless and feature-based approach to data analysis. The featureless multi-kernel approach replaces the problem of choosing the features by that of choosing the kernels. According to which of these two problems is easier, the feature-based or the featureless approach should be preferred.

However, fusing too many kernels, just as training with too many features, will inevitably worsen the generalization performance of the decision rule inferred from a small training set unless some regularization measures are taken. The technique of kernel selection proposed here is only one of possible principles of kernel fusion and has its shortcomings. In particular, such a technique should involve elements of testing on a separate set immediately in the course of training, for instance, on the basis of the leave-one-out principle.

9 References

1. V. Vapnik. *Statistical Learning Theory*. John-Wiley & Sons, Inc. 1998.
2. M.A. Aizerman, E.M. Braverman, L.I. Rozonoer. Theoretical foundations of the potential function method in pattern recognition learning. *Automation and Remote Control*, 1964, Vol. 25, pp. 821-837.
3. Duin, R.P.W, De Ridder, D., Tax, D.M.J. Featureless classification. *Proceedings of the Workshop on Statistical Pattern Recognition*, Prague, June 1997.
4. Duin, R.P.W, De Ridder, D., Tax, D.M.J. Experiments with a featureless approach to pattern recognition. *Pattern Recognition Letters*, vol. 18, no. 11-13, 1997, pp. 1159-1166.
5. Duin, R.P.W, Pekalska, E., De Ridder, D. Relational discriminant analysis. *Pattern Recognition Letters*, Vol. 20, 1999, No. 11-13, pp. 1175-1181.
6. Bishop C.M., Tipping M.E. Variational relevance vector machines. In: C. Boutilier and M. Goldszmidt (Eds.), *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann, 2000, pp. 46–53.
7. V.V. Mottl. Metric spaces admitting linear operations and inner product. *Doklady Mathematics*, Vol. 67, No. 1, 2003, pp. 140–143.
8. Mottl V., Seredin O., Dvoenko S., Kulikowski C., Muchnik I. Featureless pattern recognition in an imaginary Hilbert space. *Proceedings of the 15th International Conference on Pattern Recognition*. Quebec City, Canada, August 11-15, 2002.
9. Kolmogorov A.N., Fomin S.V. *Introductory Real Analysis*. Prentice-Hall, Englewood Cliffs, 1970.