

Affine Invariant Shape Description Using the Triangular Kernel

Hichem Sahbi

► To cite this version:

Hichem Sahbi. Affine Invariant Shape Description Using the Triangular Kernel. [Research Report] RR-5308, INRIA. 2004, pp.16. inria-00070692

HAL Id: inria-00070692

<https://hal.inria.fr/inria-00070692>

Submitted on 19 May 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.

Affine Invariant Shape Description Using the Triangular Kernel

Hichem Sahbi

N° 5308

September 2004

Thème COG



*R*apport
de recherche

Affine Invariant Shape Description Using the Triangular Kernel

Hichem Sahbi*

Thème COG — Systèmes cognitifs
Projets IMEDIA

Rapport de recherche n° 5308 — September 2004 — 13 pages

Abstract: We present in this report a novel approach for shape description based on kernel principal component analysis (KPCA). The strength of this method resides in the affine and particularly scale invariance of KPCA when using a particular kernel referred to as the triangular kernel. Beside this affine invariance, the method provides an effective way to capture the non-linearities in the shape geometry. Experiments conducted on the SQUID database show this affine invariance and the good performances for shape matching and retrieval.

Key-words: Statistical learning, kernel principal component analysis, scale-invariance, triangular kernel, shape description, image retrieval.

* IMEDIA Research Group-INRIA

Description des formes invariante par transformations affines a l'aide du noyau triangulaire

Résumé : Ce rapport décrit une nouvelle approche de description des formes basée sur l'analyse en composantes principales à noyaux (KPCA). Le point fort de cette méthode réside dans l'invariance par transformations affines, incluant l'échelle, de KPCA lors de l'utilisation d'un noyau particulier dit triangulaire. En plus de cette propriété d'invariance, cette méthode permet de décrire les non-linéarités dans la géométrie des formes. Les tests, effectués sur la base "SQUID", montrent l'invariance par échelle et les bonnes performances de généralisation pour des tâches de recherche et d'appariement des formes.

Mots-clés : Apprentissage statistique, analyse en composantes principales à noyaux, invariance par échelle, noyau triangulaire, description des formes, recherche d'images.

Contents

1	Introduction	4
2	Kernel PCA	5
3	Scale invariance	5
3.1	The triangular kernel	5
3.2	Invariance	6
4	Application in shape description	7
4.1	Indexing	7
4.2	Deformations	8
4.3	Dimensionality	9
4.4	Matching	10
5	Discussion and conclusion	11

1 Introduction

Solving the *semantic gap* in content based image retrieval basically requires relevant low level characteristics, also referred to as descriptors, including color, texture and shape [18]. The latter may be efficient to capture the discriminant information for many applications, for instance tracking silhouettes, character recognition, contour matching for medical imaging, 3D reconstruction, etc. Many signatures exist in the literature for the purpose of shape and curve description; among them the well studied edge orientation histogram, radon transform, etc. Other methods range from those based on: learning shape statistics [1, 6], geometry and a priori knowledges such as the curvature scale space (CSS) [11], skeleton and axial representations [13], context extraction and alignment [3, 17], algebraic description and invariant moments [8, 9]. A good shape description should be robust to affine transformations and also to the local non-linearities, the noise and the mirror effects.

Given a training set of $2D$ points sampled from a curve. As already well known, linear principal component analysis provides a set of orthogonal axes and the projection of training points on the span of these axes will be rotation, translation and also scale invariant up to a factor. Hence these principal axes may be efficient in order to recover the affine transformation and possibly to encode each training point using the coefficients of projection into these axes. Nevertheless, linear PCA maps each training point in the curve into a space of at most two dimensions, so this will be not sufficient to capture details and mainly the non-linearities in the shape geometry.

Kernel PCA, also known as the non-linear version of PCA, considers a positive definite kernel $k(x, x') = \langle \Phi(x), \Phi(x') \rangle$ where $\langle \cdot, \cdot \rangle$ stands for the inner product and Φ is a mapping from an input space into a higher (possibly infinite) dimensional space referred to as *the feature space*, where linear PCA can be performed. In contrast to the linear case, this non-linear version is affine invariant in the feature space but not in the input space when using many kernels for instance the well studied Gaussian. In this paper, we will show that our proposed kernel achieves this affine invariance both in the feature and the input spaces and we will discuss later, the fact that the VC dimension of a family of hypotheses (eigenspaces) trained using this kernel is infinite, so the dimension of span of the underlying principal directions increases with respect to the size of the training set. This increase of dimensionality makes it possible to capture the non-linearities in the shape geometry.

In the remainder of this paper, we use the following notations: $\mathcal{X} \subset \mathbb{R}^2$ denotes an *input space* and X is a random variable standing for the $2D$ training examples in \mathcal{X} . We denote by $\mathcal{S} = \{x_i, i = 1, \dots, N\}$ a training set generated i.i.d (independently and identically distributed) according to a particular and may be unknown probability distribution $P(X)$. Other notations will be introduced as we go along through different sections of this paper which is organized as follows: First, we review KPCA in §2, and we will show in §3 our main result of affine invariance using the triangular kernel. We present in §4 our main

application, i.e., shape description, then we conclude in §5 and we provide some directions for future work.

2 Kernel PCA

PCA is an unsupervised statistical analysis which provides a set of orthogonal axes where the projection of a training set using few of these axes, hopefully makes it possible to capture most of the statistical variance of the data. In practice, PCA have been successfully used in image processing, feature extraction, reconstruction, classification, etc [16, 2, 12].

Assuming centered training examples, i.e., $\sum_{i=1}^N x_i = 0$. In the linear case, PCA finds the principal axes by diagonalizing the covariance matrix $M = \frac{1}{N} \sum_{j=1}^N \Phi(x_j)\Phi(x_j)^t$ where x^t stands for the transpose of x . The principal orthogonal axes $\{V_k, k = 1, \dots, \min(n, N)\}$ can be found by solving the eigenproblem $M V_k = \lambda_k V_k$ where V_k and λ_k are respectively the k^{th} eigenvector and its underlying eigenvalue. It can be shown (see for instance [16]) that the solution of the above eigenproblem lies in the span of the training data, i.e.:

$$\forall k = 1, \dots, N, \quad \exists \alpha_{k1}, \dots, \alpha_{kN} \in \mathbb{R} \quad \text{s.t.} \quad V_k = \sum_{j=1}^N \alpha_{kj} \Phi(x_j) \quad (1)$$

where $\alpha_k = (\alpha_{k1}, \dots, \alpha_{kN})$ are found by solving the following eigenproblem [16]:

$$K \alpha_k = \lambda_k \alpha_k \quad (2)$$

here K is the Gram matrix¹ of the centered training set in the feature space.

3 Scale invariance

In this section, we describe the general form of the triangular kernel [14, 7] and the scale invariance of kernel principal component analysis.

3.1 The triangular kernel

The global form of the unrectified triangular kernel is:

$$k_T(x, x') = -\|x - x'\|^p, \quad p \in \mathbb{R} \quad (3)$$

This defines a conditionally positive definite kernel for $0 < p < 2$ [4]. This means that for any x_1, \dots, x_N and any $c_1, \dots, c_N \in \mathbb{R}$ such that $\sum_i c_i = 0$, we have $\sum_{i,j} c_i c_j k_T(x_i, x_j) \geq 0$.

¹Given a training set \mathcal{S} of size N , the Gram matrix K on \mathcal{S} is defined as $K_{ij} = k(x_i, x_j)$ where $i, j = 1, \dots, N$ are respectively the row and the column indices of K .

Due to the centering of the data in the feature space, we can show that this kernel can be used for KPCA [15].

Obviously, this kernel is invariant under translation and rotation but not under scaling. Nevertheless, it still has an interesting weak property of invariance that we could describe as an invariance "in shape". Given a scaling factor $\gamma > 0$, this weak invariance can formally be expressed as:

$$\begin{aligned} k_T(\gamma x, \gamma x') &= -\gamma^p \|x - x'\|^p \\ &= \gamma^p k_T(x, x') \end{aligned}$$

Thus, when the points are scaled by a certain factor γ , the value of the kernel scales by γ^p . We can also rewrite $k_T(\gamma x, \gamma x') = \langle \sqrt{\gamma^p} \Phi(x), \sqrt{\gamma^p} \Phi(x') \rangle$ which means that any scaling of a population in \mathcal{X} will also be interpreted as a scaling by a factor $\gamma^{p/2}$ in the feature space.

3.2 Invariance

In the following, we consider a situation where we scale the data by a factor $\gamma > 0$. Let's denote $\mathcal{S}^\gamma = \{\gamma x_1, \dots, \gamma x_n\}$ a training set for that population. For KPCA, our main interesting result to show is that the projection of any training example will be scale invariant, i.e.:

$$\forall x \in \mathbb{R}^n \quad \forall k = 1, \dots, N \quad \langle V_k^{(\gamma)}, \gamma x \rangle = \langle V_k^{(1)}, x \rangle \quad (4)$$

here $\{V_1^{(1)}, \dots, V_N^{(1)}\}$ and $\{V_1^{(\gamma)}, \dots, V_N^{(\gamma)}\}$ denote the eigenvectors of respectively the original and the scaled training sets. Notice that from (1) the eigenvectors are, of course, not affine invariant but the projections on these axes are affine invariant as shown below.

Proof: the proof is straightforward, and comes from the fact that the Gram matrix $K^{(\gamma)}$ of the scaled set can be written as:

$$K^{(\gamma)} = \gamma^p K^{(1)} \quad (5)$$

Using (2) it follows that:

$$\begin{aligned} K^{(\gamma)} \alpha_k^{(\gamma)} &= \gamma^p K^{(1)} \alpha_k^{(1)} \\ \Rightarrow \lambda_k^{(\gamma)} \alpha_k^{(\gamma)} &= \gamma^p \lambda_k^{(1)} \alpha_k^{(1)} \end{aligned} \quad (6)$$

here $\lambda_k^{(\gamma)}$ denotes the k^{th} eigenvalue at the scale γ . The above equality implies : $\forall k = 1, \dots, N$, $\lambda_k^{(\gamma)} = \gamma^p \lambda_k^{(1)}$ and $\alpha_k^{(\gamma)} = \alpha_k^{(1)}$. Using the latter equations, (4) is not valid but if we consider a new expansion of $V_k^{(\gamma)}$ as $\frac{1}{\lambda_k^{(\gamma)}} \sum_{j=1}^N \alpha_{kj}^{(\gamma)} \Phi(\gamma x_j)$, we can show that the projection in the span of these eigenvectors is scale invariant:

$$\begin{aligned}
\forall x \quad \langle V_k^{(\gamma)}, \gamma x \rangle &= \frac{1}{\lambda_1^{(\gamma)}} \sum_i \alpha_{ki}^{(\gamma)} k_T(\gamma x_i, \gamma x) = \frac{1}{\lambda_1^{(\gamma)}} \gamma^p \sum_i \alpha_{ki}^{(1)} k_T(x_i, x) \\
&= \frac{1}{\gamma^p \lambda_1^{(1)}} \gamma^p \sum_i \alpha_{ki}^{(1)} k_T(x_i, x) = \langle V_k^{(1)}, x \rangle \quad \square
\end{aligned} \tag{7}$$

And this shows that the ellipsoid modeling the eigenspace in the feature space can be normalized to be affine invariant. Notice that any Gram matrix built using this kernel is invertible for $0 < p < 2$ [10], so the VC-dimension [19] related to this kernel is infinite. Beside the affine invariance of this kernel (which is also achieved by the standard linear one), its discrimination power is high as the invertibility of the Gram matrix ensures that any training set can be approximated with an eigenspace with zero empirical error, i.e., all the data will live in that space.

4 Application in shape description

We ran our experiments on the SQUID² database consisting of 1100 curves. Let's define a curve as $\mathcal{L} = \{x(s) = (u(s), v(s)) \in \mathbb{R}^2, s \in [0, 1]\}$ and \mathcal{S} as a uniform sampling of \mathcal{L} . The former is considered enough representative to capture the whole shape of \mathcal{L} , so the issues related to adaptive sampling, under-sampling errors will not be considered through this paper. In practice, each curve in the SQUID database containing between 400-1600 points is randomly sampled in order to extract 128 (2D) training points which were used to synthesize 4 other curves with random orientations (in $[0^\circ, 360^\circ]$), scale factors (in $[0, 2]$) and locations (in ± 20 pixels). At the end, a total of 5500 curves were used in our experiments.

4.1 Indexing

It is easy to see that the triangular kernel is rotation and translation invariant, so under these transformations, the eigenvalues of KPCA on \mathcal{S} remain unchangeable. Only scaling the data in \mathcal{S} with γ , scales the eigenvalues by γ^p . Hence, the eigenvalues $\{\lambda_i^{(\gamma)}, i = 1, \dots, N\}$ can be normalized with respect to the largest value $\lambda_1^{(\gamma)}$ in order to cancel the factor γ^p . Therefore, $\{\lambda_i^{(\gamma)} / \lambda_1^{(\gamma)}, i = 1, \dots, m \ll N\}$ will be scale invariant and can be used as an affine description of a given curve.

For the purpose of shape retrieval, the nearest neighbor classifier based on the Mahalanobis distance is used in the "eigenvalue description" space. Figure (4) illustrates some results; for each submitted query, the system finds first the 4 most similar shapes which differ only by affine transformations, then the system finds the other similar curves with an increasing order of the Mahalanobis distance.

²www.ee.surrey.ac.uk/Research/VSSP/imagedb/squid.htm

4.2 Deformations

Define $\Phi(x_{i,def}) = \Phi(x_i) + c V_k$ to be a transformation of a curve sample x_i , which is non-linear in the input space \mathcal{X} . Moving several samples $\mathcal{S} = \{x_1, \dots, x_N\}$ according to this transformation may be interpreted as a deformation S_{def} of S in \mathcal{X} . The new set $S_{def} = \{x_{1,def}, \dots, x_{N,def}\}$, referred to as the *prefeature* set is found by:

$$x_{i,def} = \arg \min_x \|\Phi(x) - \Phi(x_i) - c V_k\|^2, \quad i = 1, \dots, N \quad (8)$$

Solving this minimization problem makes it possible to understand visually the interpretation of a movement along the principal axe V_k . In practice, this minimization problem is not convex and difficult when using many kernels including the triangular and the Gaussian[5]. In order to understand the semantic of the principal axes, we proceed as following:




Figure 1: *Interpretation of the first principal components. By definition these level-curves characterize the 2D points in the input space \mathcal{X} with the same correlation with respect to respectively, from the left to the right, the first, second, third and the fourth eigenvectors. The first and the second principal axes correspond respectively to the variation in width and height while the other axes correspond to variation in details such as the length of the tail, diagonal characteristics etc. Usually it is difficult to interpret such dimensions.*

First, let's consider two training sets of 2D points, related to two fish contours and let apply KPCA on these sets (cf. figure 1). For each case, a test set consisting of 800×600 (2D) points were projected on the span of the 4 principal components. The level curves in figure (1) show for each principal component, the data with the same correlation with respect to this component in the feature space . We can see from the left-hand side image that the intensity value of the level curves decreases in the direction of shape elongation, and this clearly shows that the first principal component characterizes the shape width. Using the same reasoning, we can see that the second eigenvector characterizes the curve height but it is more difficult to interpret the third and the fourth eigenvectors. The latter may learn some details, for instance, the shape of the tail or the fin, or other non-linearities in the fish contours.

In the second experiments, we estimate the principal eigenvectors on different curves shown in figure (2, top). For each dimension, figure (2, bottom) shows the differences between the eigenvalue spectrums related to the left-hand side curve C_1 and the others denoted respectively $\{C_2, C_3, C_4\}$. We can see, through all these diagrams, that the differences

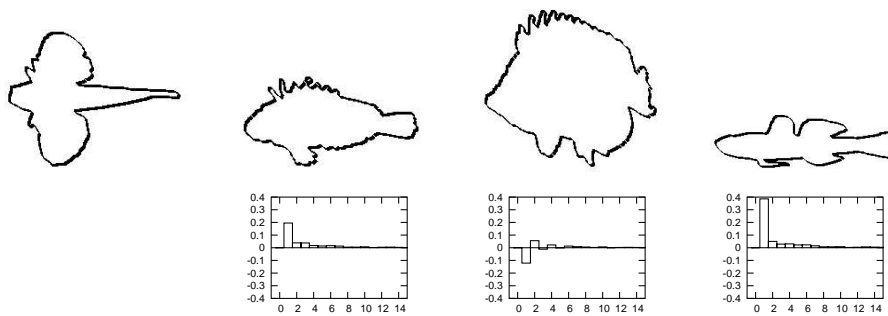


Figure 2: *Top: some fish contours randomly selected from the SQUID dataset. Bottom: eigenvalue differences between the first shape and the others.*

between the second eigenvalues characterize variations of the height between C_1 and C_i , $i = 2, 3, 4$. We can also interpret the third “eigenvalue differences” as variation in fluctuations, the fourth one can be some noise, etc. As all the first eigenvalues are normalized to 1, their differences of course vanish.

4.3 Dimensionality

Let $\sum_{k=1}^m c_{ik} V_k$ be the expansion of a training sample $\Phi(x_i)$ in the span of m -principal eigenvectors. We define the *reconstruction error* as the deviation of the above expansion from the original one, i.e., $\Phi(x_i) = \sum_{k=1}^N c_{ik} V_k$ for $m < N$. This error is defined as:

$$\frac{1}{N} \sum_{i=1}^N \left(\sum_{k=m+1}^N c_{ik} V_k \right)^2 \quad (9)$$

Using (9) and $\langle V_k, V_l \rangle = 1_{\{k=l\}}$, this error, referred to as the distance from the feature space (DFFS)[12], can simply be rewritten as $\sum_{i=1}^N \sum_{k=m+1}^N c_{ik}^2$.

The issue of selecting m , i.e., the number of eigenvectors sufficient to capture the major statistical variance of the data has been tackled in many works for instance [12] and it is interpreted as keeping only few m -dimensions while minimizing the DFFS error. Once these m -principal components are selected, the $N - m$ remaining eigenvectors correspond to the insignificant and useless details, such as noise, not necessary for discrimination. In our experiments, we select m eigenvectors in order to capture at least 95% of the statistical variance, i.e.,:

$$\frac{\sum_{k=1}^m \lambda_k}{\sum_{k=1}^N \lambda_k} \times 100 \geq 95\% \quad (10)$$

When averaging this criteria through all the curves in the SQUID database and according to our experiments, we found that among the 128 possible eigenvectors, only the 23-principal axes capture more than 95% of the statistical variance of the data (cf. figure 3, left) and for this value, the DFFS error is relatively small (cf. figure 3, right).




Figure 3: *Left: percentage of the statistical variance with respect to the number of dimensions given by the average of $\sum_{k=1}^m \lambda_k / \sum_{k=1}^N \lambda_k$ over all the curves. We can see that 23 eigenvectors capture more than 95% of the variance. Right: the reconstruction error with respect to the number of dimensions.*

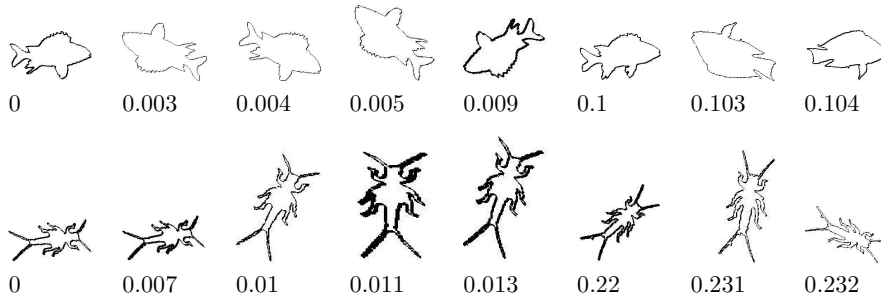


Figure 4: *Left images are query shapes while the others are some results sorted from left to right according to their dissimilarity.*

4.4 Matching

Many applications for instance 3D object reconstruction, require recovering the parameters of an affine transformation by finding a set of *good matches* between points in two curves. Given two curves \mathcal{S}_1 and \mathcal{S}_2 , we define $(x_i, x'_j) \in \mathcal{S}_1 \times \mathcal{S}_2$ as a good match, if:

$$x'_j = \arg \min_{x'_k \in \mathcal{S}_2} \|y_i - y'_k\|^2 \quad (11)$$

Here y_i, y'_k correspond respectively to the projection of x_i and x'_k in the span of the m-selected eigenvectors. We used this criteria in order to find these good matches through the SQUID dataset, figure (5) shows some results. Notice that we did not consider any smoothness or neighborhood criteria such as: every two neighboring points in a curve should have neighboring matches in the other one; this can reduce false matches.




Figure 5: *Some matching results. The matching scores using the Euclidean distance were also reported. For ease of visualization, only a subset of matches is displayed.*

5 Discussion and conclusion

One can state that this approach may work with other kernels such as the Gaussian and the linear. For instance with the linear kernel the method can be made scale invariant but as stated in the introduction the dimension of the underlying eigenspace will not exceed two. Therefore, the coefficients of projection in this eigenspace will not be sufficient to discriminate different and complex shapes especially those containing fine details. While this problem does not appear when using the Gaussian kernel, this one is not scale invariant and it necessitates finding different σ for different curves, so it may be meaningless to compare the underlying eigenvalues.

As a future work, we aim to address the main limitation of this approach which resides in the sensitivity of the estimated eigenvectors and eigenvalues to partial occlusion. We envisage also extensive evaluation on different databases and also the application of this

approach for face and mainly expression analysis. Indeed, affine invariant shape description of mouth contour can be a preliminary step for classification.

References

- [1] A. Baumberg and D. Hogg. An adaptive eigenshape model. *in proc of the British Machine Vision Conference*, 1995.
- [2] P. Belhumeur, Joao P. Hespanha, and D. Kriegman. Eigenfaces vs fisherfaces: Recognition using class specific linear projection. *IEEE transactions on pattern analysis and machine intelligence*, 19(7):711–720, 1997.
- [3] Serge Belongie, Jitendra Malik, and Jan Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 24(4):509–522, 2002.
- [4] C. Berg. Harmonic analysis on semigroups: Theory of positive definite and related functions. *Springer Verlag*, 1984.
- [5] Chris Burges and B. Schölkopf. Improving the accuracy and speed of support vector machines. In Michael C. Mozer, Michael I. Jordan, and Thomas Petsche, editors, *Advances in Neural Information Processing Systems*, volume 9, pages 375–381. The MIT Press, 1997.
- [6] D. Cremers, T. Kohlberger, and C. Schnörr. Nonlinear shape statistics via kernel spaces. *German National Conference on Pattern Recognition (DAGM)*, 2001.
- [7] F. Fleuret and H. Sahbi. Scale-invariance of support vector machines based on the triangular kernel. *Third International Workshop on Statistical and Computational Theories of Vision (part of ICCV)*, 2003.
- [8] Nishida Hirobumi and Mori Shunji. An algebraic approach to automatic construction of structural models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(12), 1993.
- [9] Subrahmonia J., Cooper D.B., and Keren D. Practical reliable bayesian recognition of 2d and 3d objects using implicit polynomials and algebraic invariants. *PAMI*, 18(5):505–519, 1996.
- [10] C.A. Micchelli. Interpolation of scattered data: Distance matrices and conditionally positive definite functions. *Constr Approx*, 2(11), 1986.
- [11] F. Mokhtarian, S. Abbasi, and J. Kittler. Robust and efficient shape indexing through curvature scale space. *In proceedings of British Machine Vision Conference*, pages 53–62, 1996.

-
- [12] A. Pentland, B. Moghaddam, and T. Starner. View-based and modular eigenspace for face recognition. *In Proceedings of the International Conference on Computer Vision*, pages 84–91, 1994.
 - [13] A. Rosenfeld. Axial representation of shape. *CVGIP*, 33(2):156–173, 1986.
 - [14] H. Sahbi and F. Fleuret. Scale-invariance of support vector machines based on the triangular kernel. *INRIA Research Report, N 4601*, 2002.
 - [15] Bernhard Scholkopf. The kernel trick for distances. *in proceedings of NIPS*, pages 301–307, 2000.
 - [16] B. Schölkopf, A.J. Smola, and K.-R. Müller. Kernel principal component analysis. *Advances in Kernel Methods - Support Vector Learning (Eds.) B. Schölkopf, C.J.C. Burges and A.J. Smola, MIT Press, Cambridge*, pages 327–352, 1999.
 - [17] Thomas B. Sebastian, Philip Klien, and Benjamin B. Kimia. On aligning curves. *PAMI*, 25(1):116–125, 2003.
 - [18] Arnold W. M. Smeulders, Marcel Worring, Simone Santini, Amarnath Gupta, and Ramesh Jain. Content-based image retrieval at the end of the early years. *in PAMI*, 22(12):1349–1380, 2000.
 - [19] Vladimir N. Vapnik. Statistical learning theory. *A Wiley-Interscience Publication*, 1998.



Unité de recherche INRIA Rocquencourt
Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex (France)

Unité de recherche INRIA Futurs : Parc Club Orsay Université - ZAC des Vignes
4, rue Jacques Monod - 91893 ORSAY Cedex (France)

Unité de recherche INRIA Lorraine : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex (France)

Unité de recherche INRIA Rennes : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Unité de recherche INRIA Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Unité de recherche INRIA Sophia Antipolis : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399