

Mercer Kernels for Object Recognition with Local Features

Siwei Lyu

Department of Computer Science
Dartmouth College

Abstract

A new class of kernels for object recognition based on local image feature representations are introduced in this paper. These kernels satisfy the Mercer condition and incorporate multiple types of local features and semilocal constraints between them. Experimental results of SVM classifiers coupled with the proposed kernels are reported on recognition tasks with the COIL-100 database and compared with existing methods. The proposed kernels achieved competitive performance and were robust to changes in object configurations and image degradations.

1. Introduction

Kernel methods received attention originally as a “trick” to introduce non-linearity into support vector machines (SVM) [21]. Evaluating a kernel function between two data is equivalent to computing the scalar product of their images in a non-linearly mapped space (usually termed as feature space). It is realized later that kernel methods are more general. Similar to SVM, many linear algorithms (e.g., PCA and Fisher linear discriminant) depend on data through their scalar products. By substituting the scalar products with kernel evaluations, these algorithms can discover non-linear patterns in data. At the same time, they are still computationally efficient, as the kernel function is evaluated in the input space [20]. Instead of using general-purpose kernels (e.g., Gaussians), recent effort has been focused on designing kernels tailored to the requirements of a specific application. Such kernels better reflect the similarities between data and thus incorporate more domain knowledge.

One important application of kernel method in computer vision is appearance based object recognition. Object recognition remains one of the most challenging problems in computer vision. Changes in illumination, pose, viewing angle, occlusion, clutters and non-rigid deformations are just a few of the complicated problems a recognition system has to face. Many applications of kernel methods to object recognition are based on global image features (e.g., global grayvalue histograms) [14, 4, 15]. Though promising performance has been reported, these methods are plagued by the deficiencies of the global features, such as being sensi-

tive to image degradations (e.g., noise, occlusion and background clutters) and not robust under changes in object configurations (e.g., translation and scaling).

Recent years have seen impressive developments in using local features computed at interest points for matching and recognition [8, 17, 16, 10, 2]. Such approaches lead to robust and compact image representations that lend themselves to powerful pattern analysis algorithms. However, the local feature representations pose several challenges to kernel design. First, it requires the kernel to work efficiently on inputs of variable lengths, as images may have a different number of local features. Secondly, the kernel should measure similarity of two unordered sets of local features with no explicit correspondence available. Furthermore, several types of local features are usually collected and they need to be fused into the kernel. For better performance, semilocal spatial and geometrical constraints between interest points should also be incorporated. Finally, to guarantee unique global optimal solutions for the SVM algorithm, the kernel must also satisfy the Mercer condition. Unfortunately, existing methods (e.g., [1, 22, 23, 12, 7]) are not satisfactory in that they do not meet all of these requirements.

The major contribution of this paper is the definition of a new class of kernels for object recognition, based on local feature representations. It can be shown formally that this class of kernels satisfy the Mercer condition and reflect similarities between sets of local features. In addition, multiple local feature types and semilocal constraints are incorporated to reduce mismatches between local features, thus further improve the classification performance. Results are shown on testing the proposed kernels, coupled with SVM classification, on recognition tasks with the COIL-100 database.

2. Methods

In this section, after a brief review of Mercer kernels and local features, the proposed kernel is described and compared with previous approaches. Then kernels using multiple types of local features and semilocal constraints are introduced, followed by a summary of the overall algorithm.

2.1. Mercer kernel

For an input space \mathcal{X} , if there is a mapping $\phi : \mathcal{X} \rightarrow \mathcal{H}$ that maps any $x, z \in \mathcal{X}$ into a Hilbert space \mathcal{H} , then a kernel, $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathcal{R}$, is constructed as $K(x, z) = \langle \phi(x), \phi(z) \rangle_{\mathcal{H}}$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ is the scalar product operator in \mathcal{H} . Such kernel function K satisfy the Mercer condition: the kernel matrix formed by restricting K to any finite subset of \mathcal{X} is positive semi-definite, and hence K is usually termed as Mercer kernels. The Mercer condition is essential to kernel design, as it is the key requirement for a unique global optimal solution to the kernel-extended pattern analysis algorithms based on convex optimization (e.g., SVM) [20]. Besides satisfying the Mercer condition, many applications also require the designed kernel to reflect similarities between the data being studied. As kernels are elicited from scalar products, they are expected to have larger values for data that are more similar to each other.¹

2.2. Local feature representation

Local features are localized descriptors that provide distinct information about a specific location of an image. Many local features (e.g., [8, 17, 16, 10, 2]) are designed to be invariant under certain image transformations, such as rotation and scaling, so that they are relatively stable to changes in object configurations. Local features have proved to be very successful in appearance based object matching and recognition, as they are distinctive, robust to image degradation and transformation, and require no segmentation [11].

Local features are usually collected at or in the neighboring region around interest points, which are specific positions in an image that carry distinctive features of the object being studied. Interest points are found by an interest point detector, popular choices are the Harris detector [6] and multi-resolution based detectors [19]. In this paper, we denote $p_i = (x_i, y_i)$ as the coordinate (in the image plane) of the i -th interest point detected in the image, and vector F_i as the local feature computed at or around p_i . An image I_a is represented by the set of local features corresponding to all interest points detected, denoted as $\mathcal{F}_a = \{F_1^{(a)}, \dots, F_{|\mathcal{F}_a|}^{(a)}\}$.

2.3. Related work

With local feature representation, an image is concisely represented by a set of local feature vectors. Accordingly, kernels that match images could be defined between two sets of local feature vectors. We start by enumerating some desirable properties of such kernels:

- The kernel should satisfy the Mercer condition;

- The computation of the kernel should be efficient in both time and space;
- The kernel should be able to handle inputs with variable lengths, as the number of interest points may vary across different images;
- The kernel should reflect similarities between two sets of local feature vectors.

It should be noted that the local feature representation does not provide correspondence between local features of two images, while only the correctly matched local features carry meaningful discriminant information. However, finding the optimal matching of local features is not always feasible in practice and many algorithms are based on heuristics. One important assumption common to most matching algorithms is that the correctly matched local features are more similar to each other than otherwise.

These properties void the use of off-the-shelf kernels, such as Gaussian and polynomial kernels, as the underlying data (sets of vectors) are not from a vector space. One can normalize the length of inputs by padding zeros. Whereas the inputs can be assumed in a vector space now, the computed quantity is of little interest to recognition. Notice, however, that it is relatively easy to build a Mercer kernel K_F on the local features, as they are usually vectors with identical dimensions. A natural idea is to construct composite kernels on the basis of such kernels, which work with sets of local features.

One simple example of such an approach is the summation kernel. On two local feature sets, $\mathcal{F}_a = \{F_1^{(a)}, \dots, F_{|\mathcal{F}_a|}^{(a)}\}$ and $\mathcal{F}_b = \{F_1^{(b)}, \dots, F_{|\mathcal{F}_b|}^{(b)}\}$, of two images, I_a and I_b , the summation kernel is defined as

$$K_S(\mathcal{F}_a, \mathcal{F}_b) = \frac{1}{|\mathcal{F}_a|} \frac{1}{|\mathcal{F}_b|} \sum_{i=1}^{|\mathcal{F}_a|} \sum_{j=1}^{|\mathcal{F}_b|} K_F(F_i^{(a)}, F_j^{(b)}). \quad (1)$$

With properties of the constructing Mercer kernels K_F , it is not hard to show that the summation kernel satisfies the Mercer condition [9]. However, its discriminative ability is compromised by the fact that all possible matchings between local features are combined with equal bias. The good matchings, highly out-numbered, could be easily swamped by the bad ones.

In [22], a kernel function based on matching local features was proposed

$$K_M(\mathcal{F}_a, \mathcal{F}_b) = \frac{1}{2} \frac{1}{|\mathcal{F}_a|} \sum_{i=1}^{|\mathcal{F}_a|} \max_{j=1, \dots, |\mathcal{F}_b|} K_F(F_i^{(a)}, F_j^{(b)}) + \frac{1}{2} \frac{1}{|\mathcal{F}_b|} \sum_{j=1}^{|\mathcal{F}_b|} \max_{i=1, \dots, |\mathcal{F}_a|} K_F(F_j^{(b)}, F_i^{(a)}). \quad (2)$$

Function K_M has the desired property of reflecting similarities of two sets of local feature vectors, as it only considers the similarities of the best matched local features. Unfortunately, despite the claim in [22], K_M is not a Mercer

¹Strictly speaking, normalized kernel evaluates the cosine similarity of two mapped data in the feature space.

kernel, for which a detailed proof is given in [9]. In [1], a similar non-Mercer kernel based on a sub-optimal matching between local features is used but measures are provided so that the probability of the kernel not being positive semi-definite is bounded. However, as pointed out earlier, the Mercer condition is essential to reliable recognition, Mercer kernels are still preferable in practice.

In [23], a Mercer kernel is proposed for sets of vectors based on the concept of principal angles between two linear subspaces. However, this kernel showed poor recognition performance as reported in [5]. In [7], the Bhattacharyya kernel is introduced where a set of vectors is represented as a multivariate Gaussian. Though provably satisfying the Mercer condition, evaluating this kernel is cubic in the number of local features. Furthermore, good matchings do not necessarily distinguish themselves in such a setting. In [12], a kernel based on Kullback-Leibler divergence is proposed. However, as the authors pointed out, it is not clear if such a kernel satisfies the Mercer condition.

2.4. A Mercer kernel between local feature sets

As discussed earlier, only the correctly matched local features with large similarity measures provide meaningful discriminant information for recognition. This indicates that such matched pairs should dominate in the kernel evaluation, if we expect the kernel to measure similarities between two sets of local feature vectors. However, directly summing the maximum similarities as in the case of K_M results in inadmissible kernels that violate the Mercer condition.

In this paper, a new class of kernels are proposed that measure similarity between local feature sets and that provably satisfy the Mercer condition. The proposed kernel function is defined as

$$K_{\mathcal{F}}(\mathcal{F}_a, \mathcal{F}_b) = \frac{1}{|\mathcal{F}_a|} \frac{1}{|\mathcal{F}_b|} \sum_{i=1}^{|\mathcal{F}_a|} \sum_{j=1}^{|\mathcal{F}_b|} [K_F(F_i^{(a)}, F_j^{(b)})]^p \quad (3)$$

where integer $p \geq 1$ is the kernel parameter. This kernel satisfies the Mercer condition, a proof of which is given in [9]. With $p = 1$, the proposed kernel includes the summation kernel as a special case. Similar to the summation kernel, all possible matchings between local features in the two sets are considered in $K_{\mathcal{F}}$, but with different bias.

It is through the kernel parameter p that the correct matched local features are given dominant bias in $K_{\mathcal{F}}$. First note that $K_{\mathcal{F}}$ has a similar form as K_M : both are sums of some similarity measures over each local feature vector. Only the max function in K_M is replaced with a summation of monomials. Consider a local feature $F_i^{(a)}$ of \mathcal{F}_a (though the following results are also true for members of \mathcal{F}_b), and its $|\mathcal{F}_b|$ kernel evaluations with each member of \mathcal{F}_b , $K_F(F_i^{(a)}, F_1^{(b)}), \dots, K_F(F_i^{(a)}, F_{|\mathcal{F}_b|}^{(b)})$. Without loss of generality, let us assume $K_F(F_i^{(a)}, F_1^{(b)}) \geq \dots \geq$

$K_F(F_i^{(a)}, F_{|\mathcal{F}_b|}^{(b)})$. The contribution of the best matched local feature in \mathcal{F}_b with $F_i^{(a)}$ in the sum of Equation (3) is:

$$\kappa = [K_F(F_i^{(a)}, F_1^{(b)})]^p / \sum_{j=1}^{|\mathcal{F}_b|} [K_F(F_i^{(a)}, F_j^{(b)})]^p \quad (4)$$

The larger the value of p is, the more dominant is the best matched pair. As p approaches infinity, all but the maximal values will have a negligible fraction in the sum, making kernel $K_{\mathcal{F}}$ behave like K_M . Furthermore, if we require that the similarity of the best matched pair in the sum has a fraction above a given threshold ρ , a lower bound of p

can be computed as $p \geq \log \frac{1-\rho}{(\rho^{|\mathcal{F}_b|-1})} / \log \frac{K_F(F_i^{(a)}, F_2^{(b)})}{K_F(F_i^{(a)}, F_1^{(b)})}$,

where $F_2^{(b)}$ is the second best matching local feature in \mathcal{F}_b for $F_i^{(a)}$ (see [9] for details). A proper p can be chosen as the maximum of such lower bounds over all training data.

2.5. Multiple local feature types

So far, in constructing kernels on local feature sets, only one type of local feature is considered. However, it is usually possible to compute multiple types of local features at an interest point. As each individual type of local feature may carry distinctive information about the underlying object, it is desirable to have them fused into the designed kernel. Hereafter, we will refer to each type of local feature as a base local feature.

Assume L different base local features are employed, and denote $f_l^{(a)} \in \mathcal{R}^{d_l}$ as the d_l -dimensional vector of the l -th base local feature computed at an interest point p_a , for $l = 1, \dots, L$. Also assume that the similarity of the l -th base local feature is properly measured by a Mercer kernel, $K_f^{(l)}$. Such kernels are termed as minor kernels in [5]. In [22], several minor kernels for some state-of-the-art local feature representations are listed. The local feature of p_a is a vector of dimension $\sum_{l=1}^L d_l$, formed by stacking all $f_l^{(a)}$ s as $F_a = (f_1^{(a)T}, \dots, f_L^{(a)T})^T$. A kernel between two such local features, F_a and F_b , is define as

$$K_F(F_a, F_b) = \prod_{l=1}^L K_f^{(l)}(f_l^{(a)}, f_l^{(b)}). \quad (5)$$

The function K_F satisfies the Mercer condition [9]. It can then be substituted into the definition of $K_{\mathcal{F}}$, Equation (3) to incorporate multiple types of local features.

2.6. Semilocal constraints

One problem of representing an image as an unordered set of local feature vectors is the assumption of independence of interest points, with which different objects with similar local feature vectors laid out differently in the image plane are indistinguishable by their local feature sets. On the other hand, as supported by the experimental results in [17], there are strong spatial correlations between the interest points

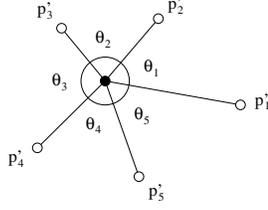


Figure 1: An example of semilocal group formed by an interest point (central filled dot) and its five nearest neighbors, p'_1, \dots, p'_5 . Hypothetical lines are added to show the neighboring angles.

and their corresponding local features in an image. Such correlations are termed semilocal constraints in [17]. For better recognition performance, it may be desirable to enforce such semilocal constraints in kernel design.

Following the method in [17], we use the local shape configuration to enforce semilocal constraints.² Specifically, an image is represented as a set of semilocal groups, which bundle together image information around spatially close interest points. One semilocal group is formed around each interest point (its central interest point) detected in an image. Each semilocal group is defined as a two component tuple, denoted as $g = \{\mathcal{F}, \Theta\}$. The first component, $\mathcal{F} = \{F_0, F_1, \dots, F_k\}$, is a set of local features collected at the central interest point as well as its k -nearest neighbors, p'_1, \dots, p'_k . The second component, $\Theta = (\theta_1, \dots, \theta_k)$, is a vector containing neighboring angles in the constellation spanned by the central interest point and its k -nearest neighbors, Figure 1. These neighboring angles convey the local geometrical constraints within the semilocal group. As pointed out in [17], if we suppose that the transformations of objects can be locally approximated by a similarity transformation, then these angles have to be locally consistent.

An image I_a is now represented by a set of semilocal groups, $G_a = \{g_1^{(a)}, \dots, g_{|G_a|}^{(a)}\}$. Correspondingly, the kernel matching images are now defined on two sets of semilocal groups. Similar to the approach taken in constructing kernel $K_{\mathcal{F}}$, we define a kernel between two sets of semilocal groups as

$$K_G(G_a, G_b) = \frac{1}{|G_a|} \frac{1}{|G_b|} \sum_{i=1}^{|G_a|} \sum_{j=1}^{|G_b|} [K_g(g_i^{(a)}, g_j^{(b)})]^p, \quad (6)$$

where K_g is a Mercer kernel between two semilocal groups to be specified later, and integer p is the kernel parameter. Kernel K_G satisfies the Mercer condition, as shown in [9]. Correct correspondence is still an important issue, as in the case of local features, only correctly matched semilocal groups are meaningful for recognition. The kernel parameter p in K_G has a similar role as its counterpart in ker-

²We are reluctant to use positions of interest points directly in the kernel, as in [22]. Such a setting makes the kernel vulnerable to changes in the spatial configurations of the object (e.g., translation).

nel $K_{\mathcal{F}}$, which gives preference to good matchings between semilocal groups.

2.7. A circular-shift invariant kernel

In constructing K_G , Equation (6), kernel K_g between two semilocal groups is left unspecified. As a semilocal group consists of two parts, a natural way to design K_g is to use the product of two kernels individually defined on the two composing parts of g as:

$$K_g(g_a, g_b) = K_{\mathcal{F}}(\mathcal{F}_a, \mathcal{F}_b) K_{\Theta}(\Theta_a, \Theta_b), \quad (7)$$

where $K_{\mathcal{F}}$ is defined as in Equation (3). Kernel K_{Θ} is defined between two vectors of neighboring angles in the semilocal constellation. Special care is required to design such a kernel, as Θ is invariant under circular-shifts. For instance, consider again the example shown in Figure 1. A vector of neighboring angles as $(\theta_3, \theta_4, \theta_5, \theta_1, \theta_2)$ represents the same geometrical configuration as $(\theta_1, \theta_2, \theta_3, \theta_4, \theta_5)$. For this reason, kernel K_{Θ} , which measures the similarity between two vectors of neighboring angles, should not treat such two vectors as different (i.e., it should also be invariant under circular-shifts).

In the most general setting, for two n -dimensional vectors $x = (x_0, \dots, x_{n-1})^T \in \mathcal{R}^n$ and $y = (y_0, \dots, y_{n-1})^T \in \mathcal{R}^n$, formally we define function $c : \mathcal{R}^n \times \{0, \dots, n-1\} \rightarrow \mathcal{R}^n$ to be the circular-shift operator as $(c(y, l))_i = (y)_{(i+l) \bmod n}$, where $(y)_i$ is the i -th component of y and $0 \leq l, i \leq n-1$. Now consider function

$$K_G(x, y) = \sum_{l=0}^{n-1} [K(x, c(y, l))]^q, \quad (8)$$

where $K : \mathcal{R}^n \times \mathcal{R}^n \rightarrow \mathcal{R}$ is a Mercer kernel and satisfies that $K(x, y) = K(c(x, d), c(y, d))$ for $0 \leq d \leq n-1$. Many commonly used kernel functions such as Gaussian and polynomial are valid candidates for K . A proof to show that K_G satisfies the Mercer condition and invariant to circular shift is given in [9].

Notice that in constructing kernel K_G , we again employ a kernel parameter q to give dominant bias to good matchings. Finally, as both $K_{\mathcal{F}}$ and K_{Θ} satisfy the Mercer condition, it can be shown that their product K_g , Equation (7), is also a Mercer kernel [9].

2.8. Summary

The process of constructing a Mercer kernel for object recognition, built with multiple types of local features and semilocal constraints, is summarized in the following algorithm:

1. With minor kernels K_f defined on base local features, construct kernel $K_{\mathcal{F}}$ on local features with Equation (5);
2. Construct kernel $K_{\mathcal{F}}$ on local feature sets with Equation (3);

3. Obtain a vector of neighboring angles in a semilocal group, and construct kernel K_G with Equation (8);
4. Combine kernels K_F and K_G into kernel K_g with Equation (7);
5. Compute kernel K_G between two sets of semilocal groups with Equation (6);

3. Experiments

In this section, we present experimental results on recognition tasks using local features and SVM classification, bridged together by the proposed kernels. In principle, the proposed kernels can work with any pattern analysis algorithm that is able to be “kernelized”, i.e., depending on data through their scalar products. SVM was chosen for its performance and generalization ability.

3.1. Experimental setup

We performed our experiments on the COIL-100 database [13], a standard test benchmark for object recognition. The COIL-100 database contains 7200 color images of 100 different objects. All images are 128×128 pixels in size. They were obtained by placing the objects on a turntable and taking a picture every 5° of viewing angle of a 360° rotation.

In our experiments, the training set of all SVM classifiers consisted of 3600 images, 36 for each of the 100 objects that correspond to a 10° difference in the viewing angles. Shown in the top row of Figure 2 are five images from the training set. From the remaining images, five different testing sets were formed:

- **Set 1:** 3600 images with viewing angles other than those used in the training set.
- **Set 2:** 3600 images generated by randomly scaling, rotating and translating images in set 1.
- **Set 3:** 3600 images generated by adding Gaussian noise of average 12 dB to the images in set 1.
- **Set 4:** 3600 images generated by embedding the images in set 1 into randomly chosen backgrounds.³
- **Set 5:** 3600 images generated by artificially adding partial occlusions (stripes from randomly chosen images) to the images in set 1.

Set 1 and 2 test the generalization ability of the kernels and classifiers to changes in viewing angles and object positions. Set 3-5 are devised to test their resilience to common image degradations, namely additive noise, background clutters and partial occlusions.

On each image, three types of local features along with their corresponding minor kernels were computed:

³Images used for backgrounds and partial occlusions in set 4 and set 5 are downloaded from <http://www.freefoto.com>.

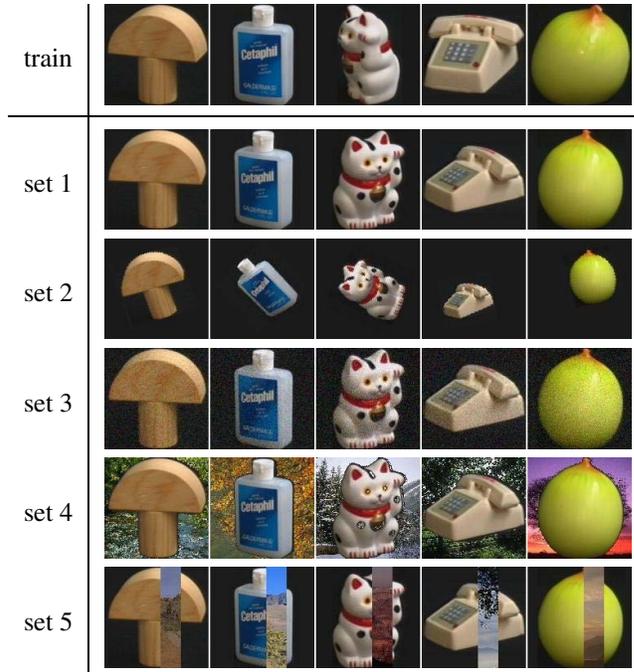


Figure 2: Examples of images used in our experiments. The first row are images from the training set. The remaining rows are examples from each of the five testing sets.

1. **Local jets** [17] are differential grayvalue invariants computed around an interest point. Each local jet is a vector of dimension 9 containing up to the third order derivatives. A Mercer kernel between two local jet features, x and z , is $K(x, z) = \exp\left(-\frac{(x-z)^T \Lambda^{-1} (x-z)}{2\sigma^2}\right)$, where Λ is the covariance matrix and $\sqrt{(x-z)^T \Lambda^{-1} (x-z)}$ is the Mahalanobis distance between x and z .
2. **Local histograms** [16] are local features consisting of histogram at different scales around interest points. Using 32 bins in computing the histogram and considering up to 3 scales, each feature is a 96 dimensional vector. A kernel based on the χ^2 -similarity between two feature vectors, x and z , $K(x, z) = \exp\left(-\frac{\chi^2(x, z)}{2\sigma^2}\right)$, is introduced in [22] and proved to satisfy the Mercer condition [4].
3. **Local phase-based features** [2] are comprised by local phases of a complex pyramid decomposition of the image. The features are 36-dimensional complex-valued vectors, and their similarity is measured by $\mathcal{C}(x, z) = \left|\frac{xz^*}{1+|x||z|}\right|$, from which a Mercer kernel is constructed as $K(x, z) = (\mathcal{C}(x, z) + 1)^q$.

For each of these local features, interest points were found by a Harris corner detector, showed to have high repeatability and robust performance [18]. Interest points too close

to the boundary were ignored to avoid image border effects. The parameters of the interest point detector were set so that, on average, approximately 100 interest points were found in an image. Semilocal groups, as described in section 2.6, were formed on each interest point using its five nearest neighbors. To have a basis of comparison, we also collected a global feature from each image. The global feature we used is the raw pixel representation [15], which was obtained by first converting a 128×128 color image into grayscale and resizing it to 32×32 pixels. A 1024-dimensional feature vector was formed by stacking the grayvalues of the resized image.

For the local feature representations, composite kernels as described in Section 2.8 were formed from the local features and their kernels. The kernel parameter, p , was set to 9 in all cases. For the global features, a Gaussian kernel was employed. The SVM classifiers were implemented with package LIBSVM [3], which was enhanced to work with kernels on local feature representations. As a standard preprocessing step in the literature, we used the normalized kernel evaluation in building the SVM classifier, as $K(x, y) \leftarrow \frac{K(x, y)}{\sqrt{K(x, x)}\sqrt{K(y, y)}}$. We employed a simple multi-class protocol for classification, namely a one-versus-the-rest scheme in training and a winner-takes-all strategy in testing. The regularization parameter of SVM was set to 10^3 in all classifiers.

3.2. Results

Shown in Table 1 is the performance of different types of kernels with the local jets on all testing sets. For comparison, the performance of the global feature (raw pixel representation) with a Gaussian kernel is also included. Performance is evaluated in error rates, which is the percentage of all misclassification cases in all testing examples. Several points are worth noting about this set of results. First, the local jet features out-performed the global features on all testing sets. The difference is more significant with the presence of image transformations and degradations (set 2-5). Furthermore, notice that the proposed kernel, $K_{\mathcal{F}}$, achieved competitive performance to that of the matching kernel K_M . As K_M is less sensitive to mismatches in local features, it had lowest error rate in some cases (set 3,4). However, its drawback is that there is no guarantee of a unique global optimal solution to the SVM training.

Shown in Figure 3 is a plot of the contribution of the best matched local feature pairs in the evaluation of kernel $K_{\mathcal{F}}$, Equation (4), with regards to the kernel parameter p . For stability, we reported here the average over kernel evaluations of all training image pairs. Notice that after $p \geq 9$, this ratio is plateaued to be more than 99%, indicating that the best matched pairs of local features has dominated in the kernel evaluation. This fact is further supported by the

	set 1	set 2	set 3	set 4	set 5
K_R	11.3	53.9	38.4	57.3	49.6
K_S	10.9	29.4	23.9	39.7	36.4
K_M	1.8	27.2	16.9	20.5	28.8
$K_{\mathcal{F}}$	1.4	23.8	17.7	25.3	24.6

Table 1: Error rates (in percentage) of the Gaussian kernel with a global feature (raw pixels) and different kernels, K_R (Equation (1)), K_S (Equation (2)) and K_M (Equation (3)), with the local jet features.

	set 1	set 2	set 3	set 4	set 5
local jets	1.4	23.8	17.7	25.3	24.6
local histograms	7.6	34.1	22.8	28.6	21.2
local phases	10.2	39.4	28.5	29.1	27.9

Table 2: Error rates (in percentage) of different local features with kernel $K_{\mathcal{F}}$, Equation (3).

	set 1	set 2	set 3	set 4	set 5
K_F	6.1	23.5	14.4	19.9	21.2
K_G	1.2	17.6	8.3	12.7	10.9

Table 3: Error rates (in percentage) of kernels using multiple types of local features Equation (5) and semilocal constraints, Equation (6).

corresponding classification error rates of $K_{\mathcal{F}}$ on test set 1, Figure 4. With p chosen greater than 9, the performance does not improved significantly.

In the second series of experiments, we tested the proposed kernel combined with different types of local features. Shown in Table 2 are the results of this experiment. Note that the local jets work well under noise, but suffer from background clutter and occlusion. The local histograms, on the other hand, are more robust in the face of partial occlusions. The local phase-based features perform worst in all the experiments. We further combined all types of local features as in Equation (5), and reported its performance in the first row of Table 3. It seems that fusion of local features does not necessarily improve the performance (set 1). However, in cases of image degradations, this approach achieved better results, possibly because multiple types of local features provide complementary information that helps to reduce ambiguity in classification.

Finally, we constructed an SVM classifier using the kernel defined in Equation (6), to further incorporate the semilocal constraints. Such a kernel, equipped with the most comprehensive domain knowledge, was expected to work best. Shown in the second row of Table 3 is the performance of this kernel on all testing sets. Compared to other kernels, it indeed achieved the lowest error rate, which suggests the efficacy of semilocal constraints.

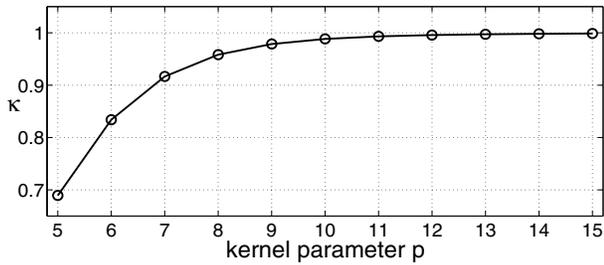


Figure 3: Contribution of the best matched local feature pairs, κ , Equation (4), in kernel $K_{\mathcal{F}}$ with local jet features as a function of the kernel parameter p .

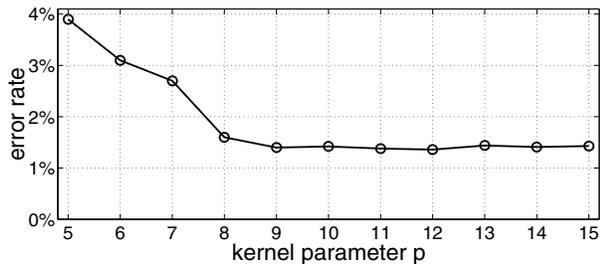


Figure 4: Classification performance of kernel $K_{\mathcal{F}}$ with local jets on set 1 as a function of the kernel parameter p .

Acknowledgment

This work was supported by Hany Farid under an Alfred P. Sloan Fellowship, an NSF CAREER Award (IIS99-83806), an NSF Infrastructure Grant (EIA-98-02068), and under Award No. 2000-DT-CX-K001 from the Office for Domestic Preparedness, U.S. Department of Homeland Security (points of view in this document are those of the authors and do not necessarily represent the official position of the U.S. Department of Homeland Security).

References

- [1] S. Boughorbel, J. Tarel, and F. Fleuret. Non-mercator kernel for SVM object recognition. In *British Machine Vision Conference (BMVC)*, 2004.
- [2] G. Carneiro and A. Jepson. Phase-based local features. In *European Conference on Computer Vision (ECCV)*, 2001.
- [3] C. Chang and C. Lin. *LIBSVM: a library for support vector machines*, 2001. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [4] O. Chapelle, P. Haffner, and V. Vapnik. SVMs for histogram based image classification. *IEEE Transactions on Neural Networks*, 10(5), 1999.
- [5] J. Eichhorn and O. Chapelle. Object categorization with SVM: Kernels for local features. In *Advances in Neural Information Processing Systems (NIPS)*, page to appear, 2004.
- [6] C. Harris and M. Stephens. A combined corner and edge detector. In *Alvey Vision Conference*, 1988.
- [7] R. Kondor and T. Jebra. A kernel between sets of vectors. In *International Conference on Machine Learning (ICML)*, 2003.
- [8] D. Lowe. Distinctive image features from scale-invariant keypoints. *International Journal on Computer Vision*, 60(2):91–110, 2004.
- [9] S. Lyu. Mercer kernels for object recognition with local features. Technical Report TR2004-520, Dartmouth College, 2004.
- [10] K. Mikolajczyk and C. Schmid. Indexing based on scale invariant interest points. *International Journal on Computer Vision*, pages 523–531, 2001.
- [11] K. Mikolajczyk and C. Schmid. A performance evaluation of local descriptors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003.
- [12] P. Monreano, P. Ho, and N. Vasconcelos. A Kullback-Leibler divergence based kernel for SVM classification in multimedia applications. In *Advances in Neural Information Processing Systems (NIPS)*, 2003.
- [13] S. A. Nene, S. K. Nayar, and H. Murase. Columbia object image library (coil-100). Technical Report CUCS-006-96, Columbia University, 1996.
- [14] F. Odone, A. Barla, and A. Verri. Building kernels from binary strings for image matching. *IEEE Transaction on Image Processing*, (to appear).
- [15] M. Pontil and A. Verri. Support vector machines for 3D object recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(6):637–646, 1998.
- [16] F. Schaffalitzky and A. Zissermann. Viewpoint invariant texture matching and wide baseline stereo. In *IEEE International Conference on Computer Vision (ICCV)*, 2001.
- [17] C. Schmid and R. Mohr. Local grayvalue invariants for image retrieval. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 19(5):530–535, 1997.
- [18] C. Schmid, R. Mohr, and C. Bauckhage. Evaluation of interest point detectors. *International Journal on Computer Vision*, 7(1):11–32, 2000.
- [19] N. Sebe, Q. Tian, E. Louprias, M. Lew, and T. Huang. Evaluation of salient point techniques. In *International Conference on Image and Video Retrieval*, 2004.
- [20] J. Shawe-Taylor and N. Cristianini. *Kernel Methods for Pattern Analysis*. Cambridge, 2004.
- [21] V. Vapnik. *Statistical Learning Theory*. Wiley, 1998.
- [22] C. Wallraven, B. Caputo, and A. Graf. Recognition with local features: the kernel recipe. In *IEEE International Conference on Computer Vision (ICCV)*, pages 257–264, 2003.
- [23] L. Wolf and A. Shashua. Kernel principle angles for classification machines with applications to image sequence interpretation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2003.