Convex Reduction of High-Dimensional Kernels for Visual Classification

Efstratios Gavves[1], Cees G.M. Snoek[1] and Arnold W.M. Smeulders[1,2]

[1]ISLA, Informatics Institute, University of Amsterdam, Amsterdam, The Netherlands
[2]Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

Abstract

Limiting factors of fast and effective classifiers for large sets of images are their dependence on the number of images analyzed and *the dimensionality of the image representation. Considering the growing number of images as a given, we aim to reduce the image feature dimensionality in this paper. We propose reduced linear kernels that use only a portion of the dimensions to reconstruct a linear kernel. We formulate the search for these dimensions as a convex optimization problem, which can be solved efficiently. Different from existing kernel reduction methods, our reduced kernels are faster and maintain the accuracy benefits from non-linear embedding methods that mimic non-linear SVMs. We show these properties on both the Scenes and PASCAL VOC 2007 datasets. In addition, we demonstrate how our reduced kernels allow to compress Fisher vector for use with non-linear embeddings, leading to high accuracy. What is more, without using any labeled examples the selected and weighed kernel dimensions appear to correspond to visually meaningful patches in the images.*

1. Introduction

In classifying pictures in object and scene classes, the issue of compactness of data emerges on the agenda. In fact, the essential information content emerges due to the enormous progress that has been made, so much so that very large number of classes are now within reach.

From the TRECvid [20] and PASCAL experiments [4] it has become clear that large sets of features are necessary to capture the internal complexity of an arbitrary type of object in its distinction from the rest of the objects in the world. The feature set has to be capable of registering all internal and external variations. Therefore it is essential to start from a big set of features. For object classification this set has now grown up to 250K and 500K per images [14, 21] for the two ImageNet Large Scale Visual Recognition Challenge winners of this year.

However, large sets of features tend to obscure the understanding of the classification result, *i.e.* which ones of the features deliver most of the distinction. What is needed is a small selected set of features specific for each class. From a complexity point of view, both for storage and computation, selection of features is to be preferred over compression. Effective and semantically meaningful reduction of feature complexity is a necessary step before making the jump to routinely classifying anything on very large numbers of visual object classes.

In this paper we discuss fast and effective classifiers by selection of a small subset of visual features from an initially very large set of features. In general, classifiers are dependent on the number of images analyzed *and* on the dimensionality of the image representation. To cope with the increasing number of images, efficient linear SVM kernels are preferred. Yet, linear SVM kernels have repeatedly been shown [15, 24] to underperform in visual classification tasks. Two current-day extensions improve the performance of linear SVM kernels. The first one relies on image feature transformations mimicking non-linear kernel maps [24, 11, 15]. The second relies on richer image representations, such as Fisher vectors [13], known to be effective in combination with linear SVM kernels. While both solutions are effective in terms of classification accuracy, they both result in high-dimensional image representations and, hence, in inefficiency. We aim to reduce the complexity of the image representation *a posteriori* by feature selection when building the kernel, while maintaining the accuracy advantage of both non-linear kernel maps and Fisher vectors.

Dimensionality reduction for visual classification has been a topic of constant interest, since it allows for more compact and thus more efficient image representations. Recent approaches [12, 7] focus on supervised vocabulary reduction by optimizing statistical and information theory criteria. Both [12] and [7] use class labels from ground truth to select vocabulary words that maximize statistical interclass criteria [6]. However, large data sets are also accom-

panied by the presence of large number of classes, *e.g* ImageNet currently counts approximately 22K classes [3]. In the presence of so many, and often visually similar, classes optimizing such inter-class criteria is a futile task. In contrast, unsupervised data dimensionality reduction, such as PCA [1] or UKDR [25] overcomes this hurdle. We consider unsupervised kernel reduction.

We propose to perform kernel reduction before classification. Where standard kernels calculate distances between features dimension by dimension, before summing them, we propose to calculate the distance on *selected and weighed* feature dimensions. We pose selection and weighing of kernel dimensions as a convex optimization problem, searching for those dimensions that cooperatively approximate the original kernel value set. Such a treatment allows for class-independent feature dimensionality reduction. Experiments on the Scenes [8] and PASCAL VOC 2007 [4] datasets show that our reduction performs as accurate as PCA yet much faster. Moreover, in contrast to existing dimensionality reduction methods, our method can be applied independent of *and* additional to, both non-linear embedding and Fisher vectors. The method does not use any labeling of images, nevertheless we find that the selected dimensions tend to correspond to meaningful patches in the image.

2. Related Work

2.1. Enhancing linear kernels

Simple linear SVM kernels have repeatedly been shown [15, 24] to underperform in visual classification. However, linear SVM efficiency benefits are crucial when moving to large data and feature sets. Several approaches have been proposed that maintain the efficiency of linear SVM kernels, while making them more accurate.

Kernel theory dictates that for every non-linear kernel map $K(\mathbf{p},\mathbf{q})$, there exists a feature mapping $\phi(\cdot)$, such that $K(\mathbf{p},\mathbf{q}) = \phi(\mathbf{q})^T\phi(\mathbf{p})$. This is a simple dot product, thus equivalent to a linear kernel calculated using $\phi(\mathbf{p})$ instead of $\mathbf{p}$. Several methods have been proposed that approximate $\phi(\cdot)$.

In their pioneering work, Maji and Berg [11] propose to approximate $\phi(\cdot)$ for the histogram intersection kernel with a modified version of Heaviside (step) functions. Generalizing on the Maji and Berg embedding, Vedaldi and Zisserman [24] describe the theoretical formulations for explicit feature maps of the χ^2, histogram intersection, Hellinger's and Jensen-Shannon kernel maps, calculated on their corresponding Fourier transform space. They furthermore propose a periodic approximation of these explicit feature maps, that allows for a concise, yet accurate enough representation for $\phi(\cdot)$. Rahimi and Recht [17] and Li et al. [10] also propose exploitation of the Fourier space of kernel maps, using random sampling instead of the exact Fourier formulation. Starting from the data point of view, Williams et al. [26] proposed the kernel PCA. The kernel PCA learns the non-linear embeddings from the data distribution. To improve the efficiency of kernel PCA when learning the non-linear embedding, Perronnin et al. propose the additive kernel PCA [15] that uses additivity to speed up the process. Both kernel PCA and additive kernel PCA are characterized of high computational complexity during embedding of new samples, although this complexity can be decreased in exchange of larger image feature dimensionality.

Rather than focusing on the kernel, other approaches emphasize enriching the image representation. In [13], Perronnin et al. propose to encode images using the Fisher vectors. Instead of mapping every local feature to a single element of a visual vocabulary, the difference between the local features and all the elements in the visual vocabulary is computed. In a similar fashion, Zhou et al. [27] propose the Super vector encoding, which also captures the difference between local features and vocabulary elements. Both methods have shown to perform exceptionally well in visual classification benchmarks like PASCAL VOC 2007 [2].

2.2. Reducing linear kernels

Data dimensionality reduction is a classical problem in the field of machine learning and statistics. The most popular method for unsupervised data reduction is PCA [1], which has been shown to work in a variety of contexts. Recently, Perronnin et al. [13] used PCA to reduce the dimensionality of SIFT features, so that their proposed Fisher vectors have a more manageable size. In [18], Raina et al. use PCA as a competitor to their LASSO optimization, which was applied on a transfer learning problem. Other methods for reducing the kernel dimensionality in an unsupervised manner include the work of Wang et al. [25]. Their desideratum is to mimic the autoencoders from the neural network literature, and like PCA, learn a transformation map which would reproduce the original kernel when applied inversely. They obtain low classification errors, especially for very low dimensions (less than 5). For larger dimensionality problems, which is the standard in visual classification, they achieve at best on par with PCA.

2.3. Contribution and novelty

State-of-the-art visual classification approaches such as the approximated feature maps [24] and Fisher vectors [13] enhance linear SVM kernels at the expense of image features with high dimensionality. Since both approaches build on the linear SVM kernel, reducing the linear kernel will result in substantial efficiency benefits. Methods for unsupervised kernel dimensionality reduction exist [1, 25], but it is unclear whether their reduced kernels benefit from methods

that enhance linear SVM models.

We propose to perform linear kernel reduction before classification. We pose selection and weighing of kernel dimensions as a convex optimization problem, searching for those dimensions that cooperatively approximate the original kernel value set. Such a treatment allows for class-independent feature dimensionality reduction. Our main contributions are that *i)* our reduction performs as accurate as PCA yet much faster, *ii)* our reduction is complementary to non-linear embeddings where others are not, and *iii)* combining our reduced kernel with Fisher vectors and approximate feature maps yields robust accuracy.

3. Convex Reduction of Kernels

Let $K(\mathbf{p},\mathbf{q})$ denote the abstract manifestation of a kernel, such as $K(\mathbf{p},\mathbf{q}) = \sum_i \min\{p_i, q_i\}$ for the histogram intersection kernel. With $N(\mathbf{p},\mathbf{q})$ we shall denote the convex reduced kernels we propose. With $\mathbf{x} = [x_1,...,x_D]^T, \mathbf{y} = [y_1,...,y_D]^T$ we denote the image feature composed of D dimensions. Given $\mathbf{x},\mathbf{y}$, let $K_{\mathbf{x},\mathbf{y}}$ denote the actual distance value calculated between the two images, that is $K_{\mathbf{x},\mathbf{y}} = K(\mathbf{p},\mathbf{q})|_{\mathbf{q}=\mathbf{y}}^{\mathbf{p}=\mathbf{x}}$. Finally, we shall denote the 1-d kernel distance values calculated per dimension i of $\mathbf{x},\mathbf{y}$ with $k^i_{\mathbf{x},\mathbf{y}}$, that is $k^i_{\mathbf{x},\mathbf{y}} = K(p_i,q_i)|_{q_i=y_i}^{p_i=x_i}$.

3.1. Theory

The most popular kernels applied for visual classification are the χ^2, histogram intersection and Hellinger's (or Bhattacharya) kernels. These kernels share two basic properties, additivity and homogeneity. A kernel is additive, when

$$K(\mathbf{p},\mathbf{q}) = \sum_i^D K(p_i,q_i), \tag{1}$$

This property is especially convenient, since performing 1-d, non-linear operations and adding them is always more efficient than performing a single multi-dimensional, non-linear operation. A kernel $K(\mathbf{p},\mathbf{q})$ is γ-homogeneous, if

$$K(c\,\mathbf{p}, c\,\mathbf{q}) = c^\gamma K(\mathbf{p},\mathbf{q}),\ \forall c \geq 0. \tag{2}$$

Homogeneity implies that scaling the kernel values by a constant c^γ has the same effect as scaling the corresponding image feature dimensions by the same constant c. The simplest kernel having both properties is the linear kernel, *i.e.* $K(\mathbf{p},\mathbf{q}) = \mathbf{q}^T\mathbf{p}$.

Given an additive kernel $K(\mathbf{p},\mathbf{q})$ and two image feature representations, $\mathbf{x},\mathbf{y}$, their kernel distance may be written as a sum of the kernel distances of the individual dimensions, that is

$$K_{\mathbf{x},\mathbf{y}} = k^1_{\mathbf{x},\mathbf{y}} + ... + k^D_{\mathbf{x},\mathbf{y}}. \tag{3}$$

Interestingly, eq. (3) implicitly assumes that when calculating the distance between two images, the 1-d distances between the individual dimensions have equal unit weight and thus equal importance. However, when two images are similar, it is mainly because some specific words contribute the most by having large values for both images. For example, when we have two images of boats, we expect that the distance between "water" visual words will contribute more to their similarity than the distance between "car wheel" visual words [22]. Therefore, we propose a different scaling factor c_i in front of each 1-d kernel distance and form a new kernel distance $N_{\mathbf{x},\mathbf{y}}$, that is

$$\begin{aligned} N_{\mathbf{x},\mathbf{y}} &= c_1 k^1_{\mathbf{x},\mathbf{y}} + ... + c_D k^D_{\mathbf{x},\mathbf{y}} \\ &= \mathbf{c}^T \mathbf{k}_{\mathbf{x},\mathbf{y}}, \end{aligned} \tag{4}$$

where $\mathbf{k}_{\mathbf{x},\mathbf{y}} = [k^1_{\mathbf{x},\mathbf{y}},...,k^D_{\mathbf{x},\mathbf{y}}]^T$. Eq. (4) implies that in order to find the similarity between two images, we may as well measure the distances $k^i_{\mathbf{x},\mathbf{y}}$ between each individual dimension and multiply them with the corresponding scaling factor c_i. By doing so, we place more importance to the distances between *certain* dimensions. Note that $N(\mathbf{p},\mathbf{q})$ is an additive kernel map as well.

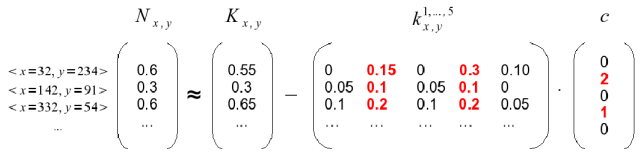

Figure 1. **Simple example of a convex reduced kernel using eq. (5).** The new kernel distances $N_{\mathbf{x},\mathbf{y}}$ approximate $K_{\mathbf{x},\mathbf{y}}$ for various pairs of images $<\mathbf{x},\mathbf{y}>$, by using only 2 out of 5 of the 1-d kernel distances.

3.2. Convex reduction

Starting from eq. (4), we want to compact $N(\mathbf{p},\mathbf{q})$. Therefore, we first separate the 1-d kernel dimensions contributing more to the final kernel distance values $N_{\mathbf{x},\mathbf{y}}$ from the less important dimensions. In order to distinguish the two, we denote with Δ the number of important dimensions, with c_i^+ the scaling factors of the important dimensions and with c_i^- the scaling factors of the unimportant dimensions. Compacting the kernel, or ignoring the unimportant dimensions, is equivalent to setting their scaling factors to zero, that is $c_i^- = 0$.

Intuitively, in order to obtain a successful compact approximation $N_{\mathbf{x},\mathbf{y}}$, we want two conditions to hold: *i)* for all possible image pairs $\mathbf{x},\mathbf{y}$ to have $N_{\mathbf{x},\mathbf{y}} \approx K_{\mathbf{x},\mathbf{y}}$ and *ii)* $\Delta \ll D$. The first condition is typically resolved in the literature by minimizing the squared difference between the target value $K_{\mathbf{x},\mathbf{y}}$ and the regressed value $N_{\mathbf{x},\mathbf{y}}$. The second condition is mathematically equivalent to minimizing the ℓ_1 norm of vector $\mathbf{c}$, that is driving as many elements in $\mathbf{c}$ to zero as possible. Taking into account eq. (4), we formulate the fol-

lowing optimization problem

$$\arg\min_{\mathbf{c}} \left\| K_{\mathbf{x},\mathbf{y}} - \mathbf{c}^T \mathbf{k}_{\mathbf{x},\mathbf{y}} \right\|^2 + \lambda \left\| c \right\|_{\ell_1}. \quad (5)$$

Eq. (5) is a regularized least squares problem, also known as LASSO convex optimization problem. A LASSO problem may be solved efficiently using non-negative quadratic programming, such as the feature-sign algorithm [9]. In order to phrase our optimization problem as a non-negative quadratic one, we assume the values of $\mathbf{c}$ to be positive. In order to efficiently solve eq. (5), we use 1-d kernel distances as training data to learn from the kernel data distribution at hand. We collect training data by measuring the kernel distances between random image pairs $\mathbf{x},\mathbf{y}$ to form the column vector $\mathbf{K}_{\mathbf{x},\mathbf{y}}$. At the same time, we store the individual 1-d kernel distances as computed from each one of the kernel dimensions $k^i_{\mathbf{x},\mathbf{y}}$, see Fig. 1. While all possible image pairs $\mathbf{x},\mathbf{y}$ can be used, we observe in our experiments that randomly sampling and using 5% of them is adequate enough, see Table 3. Parameter λ controls the sparsity of our solution, effectively affecting the coarseness of our approximation. The smaller the λ, the more dimensions are activated and therefore the better an approximation is obtained. An example is shown in Fig. 2.

After solving the optimization of eq. (5), we obtain the scale vector $\mathbf{c}$ that approximates the original kernel distances $K_{\mathbf{x},\mathbf{y}}$ with a small number of selected and weighed 1-d kernel dimensions $k^i_{\mathbf{x},\mathbf{y}}$. In order to be able to apply the optimization solution directly on the image feature representations, we take advantage of γ-homogeneity, see eq. (2). More specifically

$$\begin{aligned} N(\mathbf{x},\mathbf{y}) &= \sum_i^{\Delta} c_i^{\gamma+} K(x_i, y_i) \\ &= \sum_i^{\Delta} K(c_i^+ x_i, c_i^+ y_i). \end{aligned} \quad (6)$$

Based on eq. (6) we only need to consider the dimensions, which correspond to the non-zero scale factors c_i^+. We then multiply the respective dimensions of the image features $\mathbf{x},\mathbf{y}$ with $\sqrt[\gamma]{c_i^+}$ and ignore the rest of the dimensions. When considering the 2-homogeneous linear kernel for example we multiply each image's feature dimension i with the square root of the corresponding scaling factor, *i.e.* $\sqrt{c_i}$.

The kernel $N(\mathbf{p},\mathbf{q})$ is valid if it meets two conditions, that is: *i)* being symmetric and *ii)* being semi-positive definite. Based on eq. (4) and given that $K(\mathbf{p},\mathbf{q})$ is a valid kernel, we derive that the first condition is met, since the image features in the 1- d kernel distances are interchangeable, *i.e.* $k^i_{\mathbf{x},\mathbf{y}} = k^i_{\mathbf{y},\mathbf{x}}$. Moreover, since $c_i \geq 0$ for all elements in vector $\mathbf{c}$ and $K(\mathbf{p},\mathbf{q})$ is positive semi-definite, then $N(\mathbf{p},\mathbf{q})$ is also positive semi-definite. We therefore conclude that the new kernel is valid.

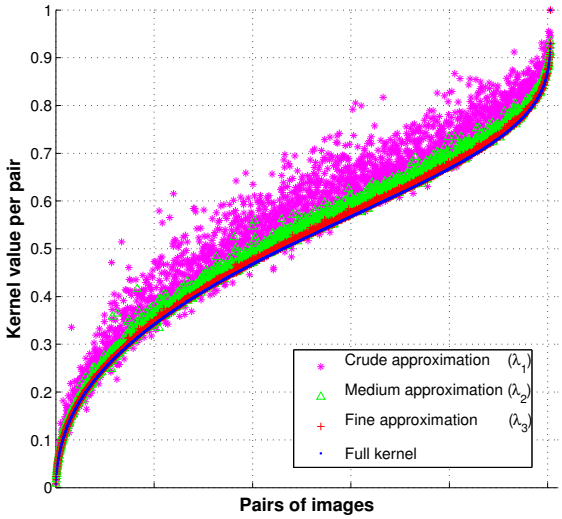

Figure 2. **Kernel value distribution for various image pairs** (best viewed in color). In blue we have the original kernel distance calculated. With the magenta, green and red we visualize convex reduced versions, that is $\lambda_1 > \lambda_2 > \lambda_3$. The λ_1 approximations uses very few (*i.e.* 66) dimensions, therefore accuracy of the kernel reconstruction is hurt.

An index of measuring the quality of our convex reduced kernel is the root mean square error $\varepsilon(\mathbf{c})$ between $N_{\mathbf{x},\mathbf{y}}$ and $K_{\mathbf{x},\mathbf{y}}$. If the error is small, then our reduced kernel properly approximates the original kernel distances. In practice, we require that the root mean square error given vector $\mathbf{c}$ over the average kernel distance value is below a threshold t, that is

$$\frac{\varepsilon(\mathbf{c})}{\hat{K}_{\mathbf{x},\mathbf{y}}} < t \quad (7)$$

We may use eq. (7) to define the optimal Δ for keeping a balance between performance loss and reduction ratio.

All the formulations above did not specify the kernel type. Although every kernel may be reduced, we observed that with the current linear model of eq. (4), we obtain the most accurate reconstruction when reducing the linear kernel. Moreover, by reducing the original linear kernel, we experimentally observe to have the additional benefit of being complementary to non-linear embeddings, such as [11, 15, 24]. In the remaining text we consider the linear kernel, that is $K(\mathbf{p},\mathbf{q}) = \mathbf{q}^T\mathbf{p}$.

3.3. Complexity

We identify four stages that need to be evaluated with respect to complexity. We separate these four stages into two groups, the *offline* and the *online* computations. The offline computations are fixed and independent of the dataset size; whether having $1,000$ or $1,000,000$ images, the offline optimization will spend the same time. Yet, the online computation timings will be severly affected when having

Table 1. **Experiment 1: Linear kernel reduction.** Our proposed convex reduced kernels are as accurate as PCA when approximating a full linear kernel, yet faster for online classification. We report rounded codebook sizes.

		Scenes				*VOC 2007*			
Method	Size	mAP	Preparation(s)	Learning(s)	Classification(s)	mAP	Preparation(s)	Learning(s)	Classification(s)
Full	4,000	0.767±0.001	0.18±0.01	10.39±0.05	0.13±0.00	0.388±0.002	0.36±0.01	20.89±0.14	0.32±0.01
PCA	600	0.763±0.001	1.44±0.01	1.14±0.01	0.04±0.00	0.381±0.001	3.10±0.06	3.59±0.01	0.08±0.00
PCA	1,200	0.766±0.001	2.74±0.08	2.71±0.15	0.05±0.00	0.387±0.002	6.17±0.15	6.70±0.02	0.12±0.00
PCA	2,200	0.766±0.001	4.88±0.01	4.14±0.01	0.08±0.00	0.388±0.002	10.33±0.35	11.14±0.04	0.18±0.00
Our kernel	600	0.747±0.001	0.11±0.00	1.14±0.01	0.03±0.00	0.364±0.002	0.20±0.00	3.58±0.01	0.08±0.00
Our kernel	1,200	0.757±0.002	0.19±0.01	2.88±0.12	0.06±0.00	0.376±0.002	0.42±0.01	6.72±0.01	0.12±0.00
Our kernel	2,200	0.761±0.000	0.41±0.01	4.42±0.16	0.08±0.00	0.379±0.002	0.80±0.01	11.13±0.02	0.18±0.00

$1,000,000$ instead of $1,000$ images. Since datasets become larger, we consider them separate for fairness.

1. Optimization time *(offline)*. In the first stage we optimize eq. (5). We first calculate 1-d kernel distances between pairs of images in the training set, with a $O(D^2 \cdot Q)$ for the linear kernel, where Q the number of pairs of images used to extract kernel distances. Then we solve eq. (5), with the complexity depending on the solver used. The optimization is performed in an offline manner once. Hence it does not affect learning and classification.

2. Preparation time *(online)*. The second stage refers to the complexity for obtaining the new image representations. We first reduce the original image features according to the new reduced kernel. During reduction, we require only scaling of the selected dimensions with an appropriate constant. Thus we have a $O(\Delta)$ complexity per image. Applying the non-linear embeddings [24, 15, 11] on top of our reduced kernel is optional. The complexity depends on the embedding. However, having fewer dimensions always results in faster embeddings. Therefore, embedding our convex reduced kernel is more efficient than embedding the full kernel.

3. Learning time *(online)*. The third stage involves the run-time complexity. After obtaining the reduced kernel, a linear SVM is applied for learning and classification. Linear SVM has learning complexity that depends on the algorithm used. However, state-of-the-art algorithms like [19] use sparse matrix algebra, hence complexity depends on the number of non-zero elements per feature. When using convex reduced kernel the dimensionality becomes smaller and the number of non-zero elements per feature drops, therefore learning becomes faster.

4. Classification time *(online)*. The forth stage refers to the classification. We multiply the obtained weight vector $\mathbf{w}$ with the respective image features (and perform one summation for the bias term). If we have M images, classification is a vector-matrix multiplication characterized by $O(D \cdot M)$ complexity. When applying our convex reduced kernel, this complexity drops to $O(\Delta \cdot M)$, which is a noticeable speedup, especially for large data and feature sets.

4. Experiments

4.1. Experimental setup

Datasets. For the experiments we use two popular datasets. The first one is the *Scenes* dataset introduced by Lazebnik et al. [8]. The *Scenes* dataset contains 4,485 medium size images of 15 indoor and outdoor scenes, such as "kitchen" or "forest". The second dataset is the *VOC 2007* dataset [4], which contains 9,963 images and is composed of 20 different objects. The *VOC 2007* dataset is a particularly challenging dataset, since one image may contain several of the objects. What is more, the instances of an object in a picture may exhibit large variation in appearance, size, context, etc. For both datasets we use the common training and test set divisions. Note that our convex reduced kernel is optimized on the training sets only.

Evaluation criteria. We study the algorithms with respect to their average precision accuracy, which is equivalent to the area under the precision-recall curve. We report the mean of average precision (mAP) over all scenes or objects. We also study the algorithms with respect to their computational efficiency. We focus on the *online* timings required by all the algorithms, that is the preparation timings, the learning timings and the classification timings, as discussed in Subsection 3.3. We repeat the experiments 10 times and report the mean and the standard deviation. The linear SVM solver we use, pegasos [19], returns approximate solutions, so we report the mean and standard deviation of mean average precision for the 10 runs. For completeness, we also report the *offline* optimization timings. All timings were computed on a standard Xeon machine at 2.93GHz.

Implementation. For all images we extract dense SIFT [23] features every 2 pixels on multiple scales, *i.e.* 4, 6, 8 and 10 pixels for bin size. We then construct a visual vocabulary of 4,000 visual words using approximate k-means [16]. Spatial pyramids [8] or other descriptors

Table 2. **Experiment 2: Non-linear kernel embedding.** In contrast to PCA, our kernel reduction is complementary to non-linear embeddings like approximated feature maps with χ^2 kernel [24]. Our kernel maintains the accuracy advantages of non-linear embeddings, but is more efficient for online classification, during all stages, when compared to non-linear embedding of the full 4,000 dimensional kernel.

		Scenes				*VOC 2007*			
Method	Size	mAP	Preparation(s)	Learning(s)	Classification(s)	mAP	Preparation(s)	Learning(s)	Classification(s)
Full	4,000	0.767±0.001	0.18±0.01	10.39±0.05	0.13±0.00	0.388±0.002	0.36±0.01	20.89±0.14	0.32±0.01
AFM(*Full*)	28,000	0.855±0.000	2.84±0.45	61.32±9.51	0.97±0.09	0.475±0.001	4.77±0.04	143.25±0.20	2.05±0.01
AFM(*PCA*-600)	4,500	0.746±0.001	1.88±0.10	7.55±0.11	0.13±0.0	0.351±0.001	4.11±0.16	23.21±3.08	0.36±0.00
AFM(*PCA*-1200)	8,700	0.753±0.001	3.72±0.03	14.65±0.15	0.24±0.01	0.346±0.001	7.55±0.45	42.37±0.07	0.65±0.03
AFM(*PCA*-2200)	15,300	0.753±0.001	4.07±0.72	23.38±0.08	0.39±0.01	0.345±0.001	13.46±0.25	70.54±0.11	1.11±0.00
AFM(*Our kernel*-600)	4,500	0.819±0.000	0.47±0.02	8.77±0.63	0.15±0.01	0.431±0.001	0.98±0.00	22.98±0.05	0.35±0.00
AFM(*Our kernel*-1200)	8,700	0.832±0.000	0.89±0.02	16.88±2.82	0.26±0.01	0.449±0.000	1.86±0.01	43.83±0.04	0.66±0.00
AFM(*Our kernel*-2200)	15,300	0.841±0.000	1.63±0.05	29.76±0.22	0.47±0.01	0.458±0.001	3.29±0.01	74.04±0.15	1.11±0.01

may be used to further improve accuracy, but we do not include them in our current experiments. In order to solve the convex-optimization problem of eq. (5), we use the feature-sign algorithm [9]. We learn the linear SVM using pegasos [19]. Prior to learning and classification, we perform ℓ_2 normalization. For the non-linear embeddings, we use the code provided [23] for the approximated feature maps [24]. We set its periods extension to 3. Preliminary experiments have shown similar results also for kernel-PCA [26]. Finally, we use the approximate version of Fisher vectors provided by [5], which does not include the variance term. We calculate the mixture parameters of the gaussian mixture model (GMM) from the approximate k-means clusters. Hence, our mixture model is equivalent to a GMM after one round.

4.2. Experiment 1: Linear kernel reduction

In the first experiment we compare our proposed convex reduced kernel with PCA [18]. We examine how well both approximate the full linear kernel, which uses a dimensionality of 4,000 words. The results are summarized in Table 1.

For both the *Scenes* and *VOC 2007* datasets, our kernel and PCA perform on par in terms of classification accuracy. Moreover, both approximate the full kernel well. We observe similar behavior for *Caltech-101* (data not shown). In terms of efficiency for offline optimization, our method needs approximately 3 minutes to harvest the 1-d kernel distances from 50K pairs of images, and it takes 4, 36 and 248 seconds for the 600, 1,200 and 2,200 dimensional kernels respectively. In the online stages of visual classification, the efficiency of our kernel becomes prevalent. Especially in the kernel preparation stage where PCA requires a matrix-matrix multiplication with a $O(\Delta \cdot D \cdot M)$ complexity, and ours is linear with $O(\Delta)$. Hence, we are 10x faster as PCA. Since the size of the reduced kernel is equal for learning and classification, our method and PCA are equally efficient for these stages. Depending on the reduction rate, both our method and PCA can be up to 10x faster than using the full linear kernel. In terms of total online classification time our method is most efficient.

4.3. Experiment 2: Non-linear kernel embedding

We test in the second experiment the complementarity of both our reduced kernel and PCA with the non-linear embedding using approximated feature maps [24]. The results are summarized in Table 2.

In contrast to PCA, our kernel reduction is complementary to approximated feature maps. Hence, our visual classification results are much more accurate for both the *Scenes* and *VOC 2007* datasets. We approach the accuracy of an embedding using a full kernel with approximately half of the number of dimensions. Regarding the efficiency, we observe substantial gains from using our reduction over the full kernel. Data preparation is up to 80% faster than using the full kernel for both *Scenes* and *VOC 2007* datasets. Learning becomes 50-85% faster for both the *Scenes* dataset and the *VOC 2007* dataset. Classification becomes 55-85% faster for the *Scenes* dataset and 45-85% faster for the *VOC 2007*dataset. Again, PCA is considerably slower in the data preparation time. The offline optimization is the same as in Section 4.2. Performing the reduction first on the linear kernel is theoretically inferior, yet has the apparent benefit that from that moment on we no longer need to apply subsequent optimizations on irrelevant kernel dimensions. For example, applying the approximate feature maps on the 600 most relevant kernel dimensions results in a 4,500-d kernel. When we compare it with the approximately similar sized regular 4,000-d kernel, we obtain a relative increase of 7% on *Scenes* and 11% on *VOC 2007* (see Table 2).

We show some classification examples for the *VOC 2007* dataset in Fig. 4. In the pictures the 10 dimensions with the largest c_i values are visualized. Although images are densely sampled, we find that features mapped on these dimensions tend to be located on salient locations, often on objects. It is interesting that

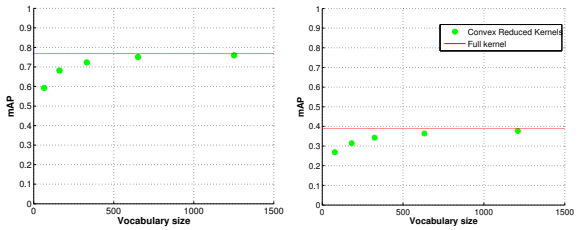

(a) *Scenes* (b) *VOC 2007*

Figure 3. **Experiment 3: Convex reduction parameters.** Mean average precision given λ.

no labeled data are used to find these features lying on salient locations. We conclude that our convex reduced kernel is complementary to non-linear embeddings, and should be used together with these embeddings, such as the approximated features maps, to obtain an accurate and efficient visual classification.

4.4. Experiment 3: Convex reduction parameters

In experiment 3 we first compare the performance of our reduced kernels with respect to the image feature dimensionality, which is controlled by parameter λ of eq. (5). We show results in Fig. 3. For the *Scenes* dataset we observe in Fig. 3(a) that a reduction to only 60 dimensions leads to a crude approximation. As a result the error ratio, see eq. (7), is high; around 14%. Increasing the number of dimensions results in smaller error ratios and better classification accuracy. The same holds for the *VOC 2007* dataset, see Fig. 3(b).

We also test the influence of the number of samples used for solving the optimization of eq. (5). We sample the 1-d kernel distances geometrically in the interval $[10K, 200K]$ for the *Scenes* dataset. The results are shown in Table 3. We observe that using the kernel distances between 50K to 100K data samples, which corresponds to 2-5% of the number of the $2.5M$ possible pairs of images in the training set, is adequate for a good approximation. Similar results were obtained for the *VOC 2007* dataset (data not shown).

Table 3. **Experiment 3: Convex reduction parameters.** Performance of the kernel approximation with respect to the sample size used for training in eq. (5). The results are averaged over 3 runs including standard deviation. Using $50K$ pairs of images, that is 2% of the number of possible pairs in the training set of the *Scenes* dataset is adequate to solve eq. (5). Similar results are obtained for the *VOC 2007* dataset (data not shown).

Method	**10K**	**20K**	**50K**	**100K**
Our kernel	0.740±0.002	0.744±0.000	0.748±0.000	0.750±0.001
AFM(*Our kernel*)	0.816±0.001	0.818±0.001	0.819±0.001	0.819±0.000

Table 4. **Experiment 4: Non-linear Fisher kernel.** Using our 600-d kernel as a basis, we can apply additional optimizations like Fisher vectors and approximate feature maps, whose computation using the Full 4,000-d kernel is prohibitive. Combining our kernel with Fisher vectors and approximate feature maps is fruitful (best results bold).

	Size		***Scenes***		***VOC 2007***	
Method	Full	Ours	Full	Ours	Full	Ours
Kernel	4K	600	0.767	0.747	0.388	0.364
AFM(Kernel)	28K	4.5K	0.855	0.819	0.475	0.431
Fisher(Kernel)	256K	40K	N/A	0.848	N/A	0.457
AFM(Fisher(Kernel))	1.8M	290K	N/A	**0.883**	N/A	**0.529**

4.5. Experiment 4: Non-linear Fisher kernel

In our fourth experiment, we take advantage of the convex reduced kernels to combine them with the popular Fisher vectors [13]. Extracting Fisher vectors using the full kernel is computationally too expensive (our full kernel is composed of 4,000 dimensions, contrast to the 256 in [13]). Using the Fisher representation for the reduced kernels results in image feature of large, yet manageable dimensionality. The reduced feature size allows us to embed the reduced Fisher vectors on the approximated feature maps [24], so that it benefits from non-linear SVMs. Using the full kernel of 4,000 dimensions, combined with Fisher vectors and approximated feature maps would require prohibitively expensive computations, resulting in vectors of 1.8M dimensions, too large to handle. We show results in Table 4. For completeness, we also include results from experiments 1 and 2.

We observe an explosion of dimensions when using Fisher vectors, resulting in image features of 40K dimensions. For the *Scenes* dataset the increase comes with an accuracy boost from 0.747 to 0.848. When we apply the approximated feature maps on the Fisher vectors and our kernel reduction we improve even further, scoring 0.883 in *mAP*. Since for the *Scenes* dataset also precision is used for evaluation, we report our current best result of 0.840, not yet including spatial pyramids and their newest variants. For the *VOC 2007* dataset the absolute improvement of embedding the reduced Fisher vectors on the approximated feature maps is even larger, going from 0.364 in mAP to 0.529, a serious improvement. We conclude that our convex reduced kernel is complementary to Fisher kernels. What is more, using our convex kernel to reduce Fisher vectors makes non-linear embeddings computationally feasible, and results in robust classification accuracy.

5. Conclusion

We propose to perform linear kernel reduction before classification. We pose selection and weighting of kernel dimensions as a convex optimization problem by search-

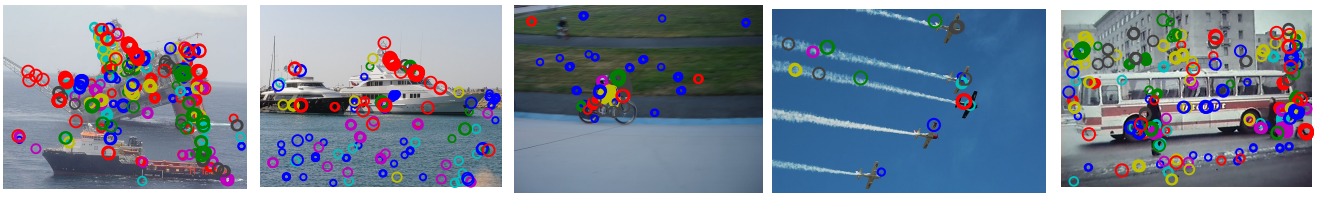

(a) (b) (c) (d) (e)

Figure 4. **Ten kernel dimensions with the largest weights c_i after reduction for *VOC 2007* images** (best viewed in color). Note that dimensions were found based on pairwise 1-d kernel distances between pairs of images only. No individual bag-of-words image features, nor labeled examples were used.

ing for those dimensions that cooperatively approximate the original kernel value set. Our proposed algorithm performs the reduction in an unsupervised manner, making it independent of the number of classes in the classification. We show on both the Scenes and VOC 2007 datasets that our reduced linear kernels are more efficient and as accurate as existing reduction methods. In contrast to existing reduction methods, our reduced kernels benefit from the state-of-the-art in visual classification, such as the approximated feature maps [24] and Fisher vectors [13]. We advocate to use our convex reduced kernel together with Fisher kernels and approximate feature maps to obtain the highest accuracy most efficiently.

6. Acknowledgements

IM-Pact BeeldCanon, STW SEARCHER and the Dutch national program COMMIT support this research.

References

[1] C. M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.

[2] K. Chatfield, V. Lempitsky, A. Vedaldi, and A. Zisserman. The devil is in the details: an evaluation of recent feature encoding methods. In *BMVC*, 2011.

[3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR*, 2009.

[4] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010.

[5] H. Jégou and B. Fulkerson. Yael. gforge.inria.fr/projects/yael.

[6] F. Jurie and B. Triggs. Creating efficient codebooks for visual recognition. In *ICCV*, 2005.

[7] S. Lazebnik and M. Raginsky. Supervised learning of quantizer codebooks by information loss minimization. *IEEE PAMI*, 2009.

[8] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, 2006.

[9] H. Lee, A. Battle, R. Raina, and A. Y. Ng. Efficient sparse coding algorithms. In *NIPS*, 2007.

[10] F. Li, C. Ionescu, and C. Sminchisescu. Random fourier approximations for skewed multiplicative histogram kernels. In *DAGM*, 2010.

[11] S. Maji and A. Berg. Max-margin additive classifiers for detection. In *ICCV*, 2009.

[12] P. K. Mallapragada, R. Jin, and A. K. Jain. Online Visual Vocabulary Pruning Using Pairwise Constraints. In *CVPR*, 2010.

[13] F. Perronnin and C. R. Dance. Fisher kernels on visual vocabularies for image categorization. In *CVPR*, 2007.

[14] F. Perronnin, J. Sánchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *ECCV*, 2010.

[15] F. Perronnin, J. Senchez, and Y. L. Xerox. Large-scale image categorization with explicit data embedding. In *CVPR*, 2010.

[16] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman. Object retrieval with large vocabularies and fast spatial matching. In *CVPR*, 2007.

[17] A. Rahimi and B. Recht. Random features for large-scale kernel machines. In *NIPS*, 2007.

[18] R. Raina, A. Battle, H. Lee, B. Packer, and A. Y. Ng. Self-taught learning: transfer learning from unlabeled data. In *ICML*, 2007.

[19] S. S. Shwartz, Y. Singer, and N. Srebro. Pegasos: Primal Estimated sub-GrAdient SOlver for SVM. In *ICML*, 2007.

[20] A. F. Smeaton, P. Over, and W. Kraaij. Evaluation campaigns and trecvid. In *MIR*, 2006.

[21] K. E. A. van de Sande, T. Gevers, and C. G. M. Snoek. Evaluating color descriptors for object and scene recognition. *IEEE PAMI*, 2010.

[22] J. C. van Gemert, C. G. M. Snoek, C. J. Veenman, A. W. M. Smeulders, and J. M. Geusebroek. Comparing compact codebooks for visual categorization. *CVIU*, 2010.

[23] A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision algorithms. http://www.vlfeat.org/.

[24] A. Vedaldi and A. Zisserman. Efficient additive kernels via explicit feature maps. *IEEE PAMI*, 2012.

[25] M. Wang, F. Sha, and M. Jordan. Unsupervised Kernel Dimension Reduction. In *NIPS*, 2010.

[26] C. Williams and M. Seeger. The Effect of the Input Density Distribution on Kernel-based Classifiers. In *ICML*, 2000.

[27] X. Zhou, K. Yu, T. Zhang, and T. S. Huang. Image classification using super-vector coding of local image descriptors. In *ECCV*, 2010.