

Video classification with Densely extracted HOG/HOF/MBH features: an evaluation of the accuracy/computational efficiency trade-off

J. Uijlings · I. C. Duta · E. Sangineto · Nicu Sebe

Received: 16 July 2014 / Revised: 11 September 2014 / Accepted: 11 September 2014
© Springer-Verlag London 2014

Abstract The current state-of-the-art in video classification is based on Bag-of-Words using local visual descriptors. Most commonly these are histogram of oriented gradients (HOG), histogram of optical flow (HOF) and motion boundary histograms (MBH) descriptors. While such approach is very powerful for classification, it is also computationally expensive. This paper addresses the problem of computational efficiency. Specifically: (1) We propose several speed-ups for densely sampled HOG, HOF and MBH descriptors and release Matlab code; (2) We investigate the trade-off between accuracy and computational efficiency of descriptors in terms of frame sampling rate and type of Optical Flow method; (3) We investigate the trade-off between accuracy and computational efficiency for computing the feature vocabulary, using and comparing most of the commonly adopted vector quantization techniques: k -means, hierarchical k -means, Random Forests, Fisher Vectors and VLAD.

Keywords Video classification · HOG · HOF · MBH · Computational efficiency

1 Introduction

The Bag-of-Words method [10,37] has been successfully adapted from the domain of still images to the domain of video using local, visual, space-time descriptors (e.g. [13,22,25,35,36,45]). Successful applications range from Human Action Recognition [24,25,32] to Event Detection [38] and Concept Classification [38,39]. However, analysing video is even more computationally expensive than analysing images. Hence, to deal with the enormous growing amount of digitalized video it is important to have not only accurate, but also computationally efficient methods.

In this paper, we take a powerful, commonly used Bag-of-Words pipeline for video classification and investigate how we can make it more computationally efficient while sacrificing as little accuracy as possible. The general pipeline is visualised in Fig. 1. In this pipeline we focus on densely sampled local visual descriptors only, since dense sampling has been found to be more accurate than keypoint-based sampling, both in images [20] and in video [46]. As type of local visual descriptors, we focus on the standard ones, which are based on local 3D volumes of histogram of oriented gradients (HOG) [11], histogram of optical flow (HOF) [12,25] and motion boundary histograms (MBH) [12]. For transforming the set of local descriptors extracted from a video into a fixed-length vector necessary for classification, we compare a variety of techniques: k -means, hierarchical k -means, Random Forests [4,16], Fisher Vectors [31] and Vector of Locally Aggregated Descriptors (VLAD) [19]. Starting from this pipeline, this evaluation paper makes the following contributions:

Fast dense HOG/HOF/MBH We exploit the nature of densely sampled descriptors to speed-up their computation. HOG,

J. Uijlings
University of Edinburgh, Edinburgh, UK
e-mail: jrr.ujlings@ed.ac.uk

I. C. Duta · E. Sangineto · N. Sebe (✉)
DISI, University of Trento, Trento, Italy
e-mail: sebe@disi.unitn.it

I. C. Duta
e-mail: duta@disi.unitn.it

E. Sangineto
e-mail: enver.sangineto@unitn.it




Fig. 1 General framework for video classification using a Bag-of-Words pipeline. The methods evaluated in this paper are instantiated in this diagram

HOF and MBH descriptors are created from subvolumes. These subvolumes can be shared by different descriptors similar to what was done in [42]. In this paper, we generalize their idea of reusing subregions to three dimensions. Matlab source code will be made available.¹

Evaluation of frame subsampling Videos consist of many frames, making them computationally expensive to analyse. However, subsequent frames also largely carry the same information. In this paper, we evaluate the trade-off between accuracy and computational efficiency when subsampling video frames.

Evaluation of optical flow Calculating optical flow is generally expensive and takes up much of the total HOF and MBH descriptor extraction time. But for optical flow there is also a trade-off between computational efficiency and accuracy. Moreover, optical flow methods are generally tested against optical flow benchmarks such as [2,7], but it is not immediately obvious that methods which perform well on these benchmarks would automatically also yield better HOF and MBH descriptors. Therefore, in this paper, we evaluate optical flow methods directly in our task of interest: video classification. Specifically, we compare the optical flow methods of Lukas-Kanade [28], Horn-Schunk [17], Farneback [15], Brox 04 [5], and Brox 11 [6].

Evaluation of descriptor encoding The classical way of transforming a set of local visual descriptors into a single fixed-length vector is using a k -means visual vocabulary and assigning local descriptors to the mean of the nearest cluster (e.g. [10]). However, both hierarchical k -means and Random Forests [30,42] are viable fast alternatives. Furthermore, the Fisher Vector [31] significantly outperforms classical k -means representation in many tasks, whereas VLAD [19] can be considered a simplified non-probabilistic version of the Fisher Vector [33] and it is computationally more efficient. In this paper, we evaluate the accuracy/efficiency trade-off of all five methods above in the context of video classification.

2 Related work

The most used local spatio-temporal descriptors are modelled after SIFT [27]: each local video volume is divided into blocks, for each block one aggregates responses (either oriented gradients or optical flow), and the final descriptor is a concatenation of the aggregated responses of several adjacent blocks. Both Dalal et al. [12] and Laptev et al. [25] proposed to aggregate 2D oriented gradient responses (HOG) and optical flow responses (HOF). In addition, Dalal et al. [12] also proposed to calculate changes of optical flow, or MBH. Both Scovanner et al. [36] and Kläser et al. [22] proposed to measure oriented gradients also in the temporal dimension, resulting in 3-dimensional gradient responses. Everts et al. [14] extended [22] to include colour channels. As Wang et al. [46] found little evidence that the 3D responses of [22] are better than HOG, in this evaluation paper we implemented and evaluated the descriptors which are most widely used: HOG, HOF and MBH.

Wang et al. [46] evaluated several interest point selection methods and several spatio-temporal descriptors. They found that dense sampling methods generally outperform interest points, especially on more difficult datasets. As this result was earlier found in image analysis [20,34], this paper focuses on dense sampling for videos. In [46] the evaluation was on accuracy only. In contrast, this paper focuses on the trade-off between computational efficiency and accuracy.

Recently, Wang et al. [45] proposed to use dense trajectories. In their method, the local video volume moves spatially through time; it tries to stay on the same part of the object. In addition, they use changes in optical flow rather than the optical flow itself. They show good improvements over normal HOG, HOF and MBH descriptors. Nevertheless, combining their dense trajectory descriptors with both normal HOG, HOF and MBH descriptors still gives significant improvements over dense trajectories alone [21,45]. In this paper we focus on HOG, HOF and MBH. Note that we evaluate the accuracy/efficiency trade-off for several optical flow methods which may be of interest also when using dense trajectories.

In [34], Sangineto proposes to use Integral Images [44] to efficiently compute densely extracted SURF features [3] in *still images*. The work of Uijlings et al. [42] proposes

¹ <http://homepages.inf.ed.ac.uk/juijling/index.php?page=software>.

several methods to speed-up the Bag-of-Words classification pipeline for *image* classification and provides a detailed evaluation on the trade-off between computational efficiency and classification accuracy. In this paper we perform such evaluation on *video* classification. Inspired by [42] we propose accelerated densely extracted HOG, HOF and MBH descriptors and provide efficient Matlab implementations. In addition, we evaluate various video-specific aspects such as frame sampling rate and the choice of optical flow method.

The Fisher Vector [31] has been shown to outperform standard vector quantization methods such as k -means in the context of Bag-of-Words. On the other hand, the recently proposed VLAD descriptors [19] can be seen as a non-probabilistic version of Fisher Vectors which are faster to compute [19,33]. In this paper, we evaluate the accuracy/efficiency trade-off using Fisher Vector and VLAD in the context of video classification.

3 Bag-of-Words for video

In this section we explain in detail the pipeline that we use. We mostly use off-the-shelf yet state-of-the-art components to construct our Bag-of-Words pipeline, which is necessary for a good evaluation paper. In addition, we explain how to create a fast implementation of densely sampled HOG and HOF descriptors, and also implicitly for MBH, being MBH based on HOG and Optical Flow. We make the HOG/HOF/MBH descriptor code publicly available.

3.1 Descriptor extraction

In this section we describe the details of our implementation for dense extraction of HOG, HOF and MBH descriptors. Specifically, in Sect. 3.1.1 we show how HOG and HOF can be efficiently extracted and aggregated from video blocks. Then, in Sect. 3.1.2 we deal with MBHs, which are largely based on HOG. Finally, since in this paper we compare our implementation with the widely used available code of Laptev [25], in Sect. 3.1.3, we show the parameters we have adopted in using Laptev's code. Both ours and the Laptev's system work on grey values only. Note that Laptev's implementation does not include MBH descriptors, thus the comparison performed in our experiments only concerns HOG and HOF.

3.1.1 Fast dense HOG/HOF descriptors

For both HOG and HOF descriptors, there are several steps. First one needs to calculate either gradient magnitude responses in horizontal and vertical directions (for HOG), or optical flow displacement vectors in horizontal and vertical directions (for HOF). Both result in a 2-dimensional vec-




Fig. 2 Blocks in a video volume can be reused for descriptor extraction. In our paper descriptors consist of 3×3 blocks in space and 2 blocks in time, shown in *blue* (color figure online)

tor field per frame. Then for each response the magnitude is quantized in o orientations, usually $o = 8$. Afterwards, one needs to aggregate these responses over blocks of pixels in both spatial and temporal directions. The next step is to concatenate responses of several adjacent pixel blocks. Finally, descriptors have to be normalized and sometimes PCA is performed to reduce their dimensionality, often leading to computational benefits or improved accuracy.

To calculate gradient magnitude responses we use HAAR-features. These are faster to compute than Gaussian Derivatives and have proven to work better for HOG [11]. Quantization in o orientations is done by dividing each response magnitude linearly over two adjacent orientation bins.

We use the classical Horn–Schunk [17] method for optical flow responses as a default. We use the version implemented by the Matlab Computer Vision System Toolbox. In addition, we evaluate four other optical flow methods: Lucas-Kanade [28], also using the Matlab Computer Vision System Toolbox, the method of Färneback [15], using OpenCV² with the mexopencv interface,³ Brox 04 [5], and Brox 11 [6] using the author's publicly available code.

Both HOG and HOF descriptors are created out of blocks. By choosing the sampling rate identically to the size of a single block, one can reuse these blocks. Figure 2 shows an example on how a video volume can be divided into blocks. Once responses per block are computed, descriptors can be formed by concatenating adjacent blocks. In this paper we use descriptors of 3×3 blocks in the spatial domain and 2 blocks in the temporal domain, as shown in blue in Fig. 2, but these parameters can be easily changed. Hence each block is reused 18 times (except for the blocks on the borders of the video volume).

² <http://opencv.org>.

³ <https://github.com/kyamagu/mexopencv>.

To aggregate responses over space we use the Matlab-friendly method proposed by [42]: Let R be an $N \times M$ matrix containing responses in a single orientation (be it gradient magnitude or optical flow magnitude). Let B_N and B_M be the number of elementary blocks from which HOG/HOF features are composed. Now it is possible to construct (sparse) matrices O and P of respectively $B_N \times N$ and $M \times B_M$ such that $ORP = A$, where A is a $B_N \times B_M$ matrix containing the aggregated responses for each block. O and P resemble diagonal matrices but are rectangular and the filled in elements follow the 'diagonal' of the rectangle instead of positions (i, i) . By proper instantiation of these matrices we perform interpolation between blocks, which provides the descriptors some translation invariance. For integration over time we add the responses of the frames belonging to a single block. For more details we refer the reader to the work of [42].

In this paper, we extract descriptors on a single scale where blocks consist of 8×8 pixels \times 6 frames, which at the same time is our dense sampling rate. Descriptors consist of $3 \times 3 \times 2$ blocks. Both for HOG and HOF the magnitude responses are divided into 8 orientations, resulting in 144-dimensional descriptors. PCA is performed to reduce the dimensionality by 50% resulting in 72-dimensional vectors. Afterwards, normalization is performed by the L1-norm followed by the square root, which effectively means that Euclidean distances between descriptors in fact reflect the often superior Hellinger distance [1].

3.1.2 Motion boundary histograms descriptor

Another commonly used descriptor for video classification tasks is MBH, proposed by Dalal et al. [12], who proved its robustness to camera and background motion. The intuitive idea of MBH is to represent the oriented gradients computed over the vertical and the horizontal optical flow components. The advantage of such representation is that constant camera movements tend to disappear and the description focuses on optical flow *differences* between frames (*motions boundaries*).

In more detail, the optical flow's horizontal and vertical components are separately represented using two scalar maps, which can be seen as grey-level "images" of the motion components. HOG are then computed for each of the two optical flow component images, using the same approach used for computing HOG in still images. Taken into account only flow differences, the information about changes in motion boundaries is kept and the constant motion information is removed, which leads to the cancellation of most of the effects of camera motion.

In our MBH implementation we follow the pipeline suggested in [12] and mentioned above. Once computed the horizontal and vertical optical flow components, HOG are

computed on each image component using the same efficient approach and the same parameters shown in Sect. 3.1.1. Also the block-based aggregation step is analogous to what described in Sect. 3.1.1.

The outcome of this process is a pair of horizontal (MBH_x) and vertical (MBH_y) descriptors [12], each one composed of 144 dimensions. We separately apply PCA to both MBH_x and MBH_y and we obtain two vectors of 72 dimensions each. The (PCA-reduced) MBH_x and MBH_y vectors can then be either separately used in the subsequent visual word assignment and classification stages (Fig. 1) or combined in order to get a unique descriptor. In [45] the authors state that late fusion of MBH_x and MBH_y gives a better performance than concatenating the two descriptors before the visual word assignment step. Hence in this paper we will report results for MBH_x and MBH_y separately, and a late fusion of the two which we simply denote as MBH. This late fusion combines the outcomes of the two (independent) classifications with equal weights. Finally, in Sect. 4.6, we will also show results concerning a late fusion strategy involving all the descriptors (MBH_x, MBH_y, HOG and HOF).

3.1.3 Existing HOG/HOF descriptors

We use the existing implementation of Laptev et al. [25]. We use the default parameters as suggested by the authors, which compared to our descriptors are as follows: They perform a dense sampling at multiple scales. At the finest scale, blocks are 12×12 pixels \times 6 frames, sampling rate is every 16 pixels by every 6 frames. They consider 8 spatial scales and 2 temporal scales for a total of 16 scales, where each scale increases the descriptor size by a factor of $\sqrt{2}$. In the end, they generate around 33% less descriptors than our single scale dense sampling method.

Unlike our descriptor extraction, the implementation of [25] uses 4 orientations for HOG and 5 orientations for HOF, resulting in respectively 72- and 90-dimensional descriptors.

3.2 Visual word assignment

We use five different ways of creating a single feature representation of a set of descriptors extracted from a single video: k -means, hierarchical k -means, Random Forests [4, 16], VLAD [19] and Fisher Vectors [31].

For hierarchical k -means we use the implementation made available by VLFeat [43]. For the regular k -means assignment, we make use of the fact that the descriptors are L2-normalized: Euclidean distances are proportional to dot products (cosine of angles) between the vectors. Hence finding the minimal euclidean distance is equivalent to finding the maximal dot product, yet more efficient to compute [42]. For both hierarchical k -means and regular k -means, we use

4,096 visual words. For hierarchical k -means, we learn a hierarchical tree of depth 2 with 64 branches per node of the tree (preliminary experiments showed a large decrease in accuracy when using a higher depth with fewer branches, but only marginal improvements in computational efficiency, data not shown). We normalize the resulting frequency histograms using the square root, which discounts frequently occurring visual words, followed by L1-normalization.

Random Forests are binary decision trees which are learned in a supervised way by randomly picking several descriptor dimensions at each node with several random thresholds and choose the one with the highest Entropy Gain. We follow the recommendations of [42], using 4 binary decision trees of depth 10, resulting in 4,096 visual words. The resulting vector is normalized by taking the square root followed by L1.

The Fisher Vector [18] as used in [31] encodes a set of descriptors D with respect to a Gaussian mixture model (GMM) which is trained to be a generative model of these descriptors. Specifically, the set of descriptors is represented as the gradient with respect to the parameters of the GMM. This can be intuitively explained in terms of the EM algorithm for GMMs: Let G_λ be the learned GMM with parameters λ . Now use the E-step to assign the set of descriptors D to G_λ . Then the M-step yields a vector F with adjustments on how λ should be updated to fit the data (i.e. how the GMM clusters should be adjusted). This vector F is exactly the Fisher Vector representation. We follow [31] and normalize the vector using a square root of the absolute values and afterwards keep the original sign ($(\text{sign}(f_i))\sqrt{|f_i|}$), followed by L2. In this paper we use two common cluster sizes for the GMM: 64 and 256 clusters [31]. Without a spatial pyramid [26], for our 72-dimensional HOG/HOF/MBH $_x$ /MBH $_y$ features this will yield vectors of 9,216 and 36,864 dimensions respectively. While not comparable with the dimensionality of other methods, Fisher Vectors (and VLAD) allow for linear Support Vector Machines rather than Histogram Intersection or χ^2 -kernels. Hence efficiency wise, the simpler classifiers will compensate for the larger dimension of the feature vectors.

The recently proposed VLAD [19] representation can be seen as a simplification of the Fisher Vector [19,33] in which: (1) a spherical GMM is used, (2) the soft assignment is replaced with a hard assignment and (3) only the gradient of G_λ with respect to the mean is considered (first order statistics). This leads to a lower dimensional representation, half of the dimensions of a Fisher Vector, in which second order statistics are also used. Following [19] we use for VLAD the same normalization scheme used for Fisher Vectors: We square root the VLAD vectors while keeping their sign, followed by L2-normalization. For good comparison to the Fisher Vectors, we use a dictionary of 128 and 512 clusters respectively, leading to features of dimensionality identical to the Fisher Vectors: 9,216 and 36,864 dimensions.

We use the Spatial Pyramid [26] in all our experiments. Specifically, we divide each video volume into the whole video and into three horizontal parts which intuitively roughly corresponds to a ground, object, and sky division (in outdoor scenes).

3.3 Classification

For classification we use Support Vector Machines which are powerful and widely used in a Bag-of-Words context (e.g. [10,26,42,43]). For k -means, hierarchical k -means, and Random Forests, we use SVMs with the Histogram Intersection kernel, using the fast classification method as proposed by [29]. For the Fisher Vector and VLAD, we use linear SVMs. For both types of SVMs, we make use of the publicly available LIBSVM library [8] and the fast Histogram Intersection classification of [29].

4 Experiments

Our baseline consists of densely sampled HOG, HOF and MBH(x/y) descriptors, all consisting of blocks of 8×8 pixels \times 6 frames. For HOF and MBH(x/y), optical flow is calculated using Horn–Schunk. Gradient and flow magnitude responses are quantized in 8 bins. The final descriptors consist of $3 \times 3 \times 2$ blocks. PCA always reduces dimensionality of descriptors by 50%. We use a spatial pyramid division of $1 \times 1 \times 1$ and $1 \times 3 \times 1$ [26] (we have no temporal division). Normalization after word assignment is done by either taking the square root while keeping the sign followed by L2 for the Fisher Kernel, or by the square root plus L1 for all other methods. We use SVMs for classification, with either a linear kernel for the Fisher Vectors or histogram intersection kernel for all other visual word assignment methods.

Starting from our baseline we perform four experiments: (1) We compare five different visual word assignment methods: k -means, hierarchical k -means, Random Forests, VLAD and the Fisher Kernel; (2) We compare our densely extracted descriptors with the descriptors provided by Laptev et al. [25]; (3) We evaluate the efficiency/accuracy trade-off by subsampling video frames for the descriptor extraction process; (4) For HOF and MBH(x/y) descriptors, we compare five different optical flow implementations: Horn–Schunk, Lukas-Kanade, Farneback [15], Brox 04 [5] and Brox 11 [6].

All timing experiments are performed on a single core of an Intel(R) Xeon(R) CPU E5620 2.40GHz. We use mainly Matlab, but most toolboxes used by us have mex-interfaces to C++ implementations for critical functions. All implementations are heavily optimized for speed. Since the computation involves many common operations that use standardized and optimized libraries (e.g. convolutions, matrix multipli-

cations) on large quantities of data, virtually the entire time is spent on core calculations while the overhead is negligible; using only C++ will not result in noticeable differences in the overall timing results presented in this paper.

Based on our experiments we provide two recommendations, one for real-time video classification and one for accurate video classification. Finally we give a comparison with the state-of-the-art.

4.1 Dataset

We perform all experiments on the UCF50 Human Action Recognition dataset [32]. This dataset contains 6,600 realistic videos taken from Youtube and as such has large variations in camera motion, object appearance and pose, illumination conditions, scale, etc. The 50 human action categories are mutually exclusive and include actions such as biking, diving, drumming, and fencing. The frames of the videos are 320×240 pixels. The video clips are relatively short with a length that varies around 70–200 frames. The dataset is divided in 25 predefined groups. Following the standard procedure we perform a leave-one-group-out cross-validation and report the average classification accuracy over all 25 folds. Optimization of the SVM slack parameter is done for every class for every fold on the training set (containing 24 groups).

4.2 Visual word assignment

In this experiment we compare the following visual word assignment methods: k -means, hierarchical k -means, Random Forests, VLAD and Fisher Vector. k -means, hierarchical k -means and Random Forests are similar in the sense that the final vector represents visual word counts. To compare these methods we ensure that all have 4,096 visual words. For k -means this means performing clustering with $k = 4,096$. For hierarchical k -means we use a hierarchy of depth 2 with 64 branches at each node. The Random Forest consists of 4 trees of depth 10. We choose to base our Fisher Vectors on standard sizes for the number of clusters: 64 and 256 clusters [9,31]. While Fisher Vectors are of higher dimensionality, the vectors work with linear classifiers. This means that Fisher Vectors are best compared with the other visual word assignment methods in terms of the accuracy/efficiency trade-off. Similarly, we adopted two standard cluster sizes for VLAD: 128 and 512 dimensions respectively [19] and we used linear classifiers as well.

The accuracy and computational efficiency for the various word assignment methods for our HOG, HOF and MBH(x/y) features are presented in Fig. 3 and Table 1. The first thing to notice is that the Fisher Vector with 256 clusters has the best accuracy of 0.765 for HOG, 0.795 for HOF, 0.796 for MBH x and 0.804 for MBH, while taking 3.39 s/video (per descriptor type). k -means has also good accuracy at 0.728 for HOG,

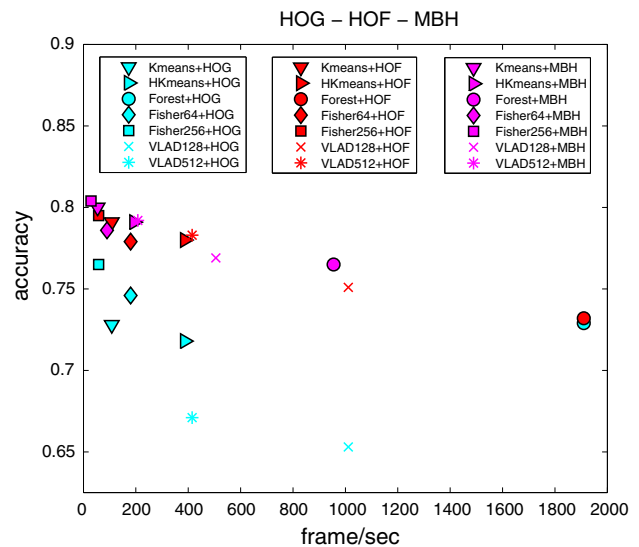


Fig. 3 Accuracy/efficiency trade-off for various word assignment methods and features. For a better readability of the figure, we omitted the results concerning MBH x and MBH y (see Table 1)

0.791 for HOF, 0.782 for MBH x and 0.8 for MBH. However, the computational time is at 1.81 s/video. This means that the Fisher Vector (with 256 clusters) for video classification is superior in accuracy but slightly slower compared to k -means. For computational efficiency, the Random Forest is by far the fastest and takes 0.1 s/video. The hierarchical k -means (hk -means) is four times slower at 0.47 s/video, and performs slightly worse on HOG (0.718 hk -means vs. 0.729 RF) but significantly better on HOF (0.780 hk -means vs. 0.732 RF) and on MBH x , MBH y and MBH (respectively, 0.774 vs. 0.738, 0.763 vs. 0.739 and 0.791 vs. 0.765).

In terms of classification time per video, we measure 0.017 s/video when using the fast Histogram Intersection-based classification for SVMs [29] for k -means, hk -means, and Random Forests. We measure 0.001 s/video for the linear classifier used on the Fisher Vector representation with 256 clusters. This means that the classification time is negligible compared to the word assignment time and is of little concern for video classification.

For the remainder of this paper, we choose to perform our evaluation on two word assignment methods: the Fisher Vector, which yields the most accurate results, and hk -means, which is the second fastest after Random Forests, while its accuracy for HOF and MBH(x/y) is much higher than using Random Forests.

4.3 Comparison with Laptev et al.

In this experiment we compare the publicly available code from [25] with our implementation. We compare only to the dense sampling option as [46] has already proven that dense

Table 1 Trade-off accuracy/efficiency for the following visual word assignment methods: *k*-means, hierarchical *k*-means (*hk*-means), Random Forest (RF), Fisher Kernel with 64 and 256 clusters (FK 64 and FK 256)

	<i>k</i> -means	<i>hk</i> -means	RF	FV 64	FV 256	VLAD 128	VLAD 512
HOG Acc	0.728	0.718	0.729	0.746	0.765	0.653	0.671
HOF Acc	0.791	0.780	0.732	0.779	0.795	0.751	0.783
MBHx Acc	0.782	0.774	0.738	0.767	0.796	0.749	0.774
MBHy Acc	0.772	0.763	0.739	0.759	0.787	0.737	0.765
MBH Acc	0.800	0.791	0.765	0.786	0.804	0.769	0.792
s/video	1.81	0.51	0.10	1.10	3.39	0.19	0.47
Frame/s	108	387	1,910	180	58	1,011	415

Assignment time for HOG and HOF is the same

Fig. 4 Accuracy comparison between [25] and our HOG/HOF descriptors

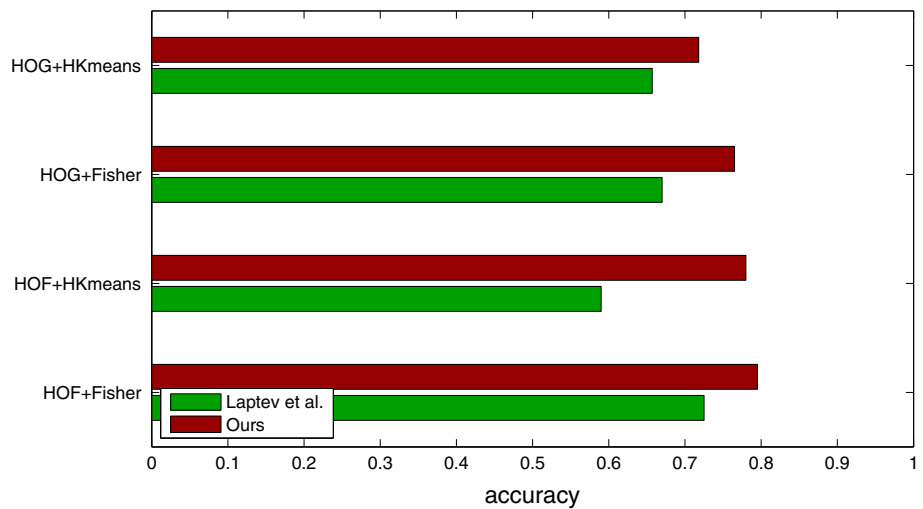


Fig. 5 Computational efficiency comparison between [25] and our HOG/HOF descriptors




Table 2 Comparing the dense HOG/HOF implementation of [25] and ours

	<i>hk</i> -means		FV 256		Efficiency	
	HOG	HOF	HOG	HOF	s/vid	Frame/s
[25]	0.657	0.590	0.670	0.725	141	1.4
Ours	0.718	0.780	0.765	0.795	15	12.8

The descriptor extraction time is measured for extracting both HOG and HOF features, as the binary provided by [25] does always both. Descriptor extraction time is independent of the visual word assignment method (RF or FV 256)

sampling outperforms the use of space-time interest points. Moreover, only HOG and HOF features are used for comparison because the code in [25] does not include any implementation for MBH features. Results are presented in Figs. 4 and 5 and in Table 2.

The results show that for all settings there is a significant difference in accuracy between the dense implementation of [25] and our method. For the Fisher Vector, HOG descriptors yield 0.670 accuracy for [25] and 0.765 accuracy for our implementation and HOF descriptors yield 0.725 accuracy for [25] and 0.795 accuracy for our implementation. These are accuracy increases of 9 and 7% respectively. Similar differences are obtained using *hk*-means. Part of the difference can be explained by the fact that we sample differently: because we reuse blocks of the descriptors, our sampling rate is defined by the size of a single block. This means we sample descriptors every 8 pixels and every 6 frames at a single scale, whereas [25] samples every 16 pixels and every 6 frames at 10 increasingly coarse scales. For our method this yields around 150 descriptors per frame or around 29,000 descriptors per video whereas [25] generates around 90 descrip-

tors per frame or around 17,500 descriptors per video, which means we generate 66% more descriptors. While this may seem unfair towards [25], in this paper we are interested in the trade-off between accuracy and computational efficiency, which makes the exact locations from where descriptors are sampled irrelevant.

In terms of computational efficiency our method is more than 9 times faster: their method takes 141 s/video while our method takes 15 s/video. Our method is faster because we reuse blocks in our dense descriptor extraction method. Note that because the method of [25] samples fewer descriptors, visual word assignment time is faster. But using [25] the overall computation time will be completely dominated by descriptor extraction.

To conclude, our implementation is significantly faster and significantly more accurate than the version of [25].

4.4 Subsampling video frames

In video, subsequent video frames largely contain the same information. As the time for descriptor extraction is the largest bottleneck in video classification, we investigate how the accuracy behaves if we subsample video frames and hence speed-up the descriptor extraction process.

For a fair comparison, we want the descriptors always to describe the same video volume. In our baseline, each descriptor block consists of 8×8 pixels \times 6 frames. To subsample in such a way that every block describes the same video volume regardless of the sampling rate, we do the following: if we sample every 2 frames, we aggregate responses over 3 frames (i.e. of frame 2, 4 and 6). When sampling every 3 frames, we aggregate responses over 2 frames (i.e. frame 2 and 5), and when sampling every 6 frames in which we only consider a single frame per descriptor block (i.e. frame 3). Results are presented in Fig. 6 and Table 3.

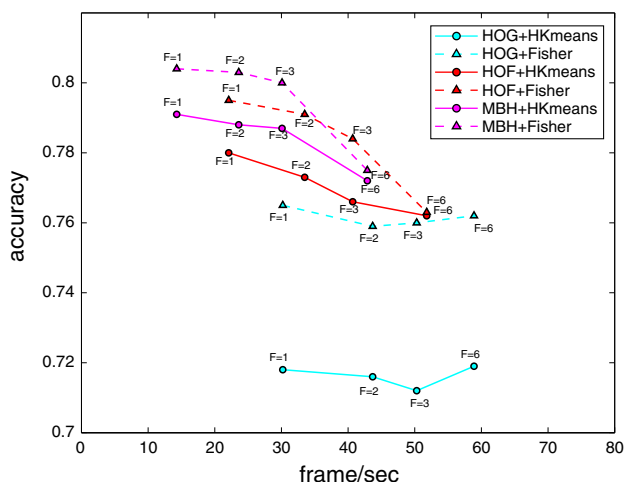


Fig. 6 Trade-off accuracy/efficiency when varying sampling rate. F stands for frames per block and it is directly related to sampling rate

For HOG descriptors, subsampling video frames have surprisingly little effect on the accuracy, both for hk -means and Fisher Vectors: using Fisher Vectors, a sampling rate of 1 yields an accuracy of 0.765 while a sampling rate of 6 yields 0.762 accuracy. The result of hk -means is basically constant, with slight oscillations. In terms of computational efficiency, a significant speed-up is achieved: sampling every 6 frames instead of every frame gives a speed-up from 6.5 s/video to 3.3 s/video.

For HOF descriptors, subsampling has a bigger impact: For the Fisher Vector, the accuracy is 0.795 using a sampling rate of 1, maintains a respectable 0.791 accuracy at a subsampling rate of 2 frames, while dropping significantly to 0.763 for sampling every 6 frames. Accuracy with hk -means is less affected and drops from 0.78 at sample rate of 1 to 0.762

Table 3 Trade-off between frame sampling rate and accuracy

	(Frames/block samplerate)	$\begin{pmatrix} 6 \\ 1 \end{pmatrix}$	$\begin{pmatrix} 3 \\ 2 \end{pmatrix}$	$\begin{pmatrix} 2 \\ 3 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 6 \end{pmatrix}$
HOG	hk -means	0.718	0.716	0.712	0.719
	FV 256	0.765	0.759	0.760	0.762
	s/vid	6.5	4.5	3.9	3.3
	Frame/s [†]	30.2	43.7	50.3	58.9
HOF	(Frames/block samplerate)	$\begin{pmatrix} 6 \\ 1 \end{pmatrix}$	$\begin{pmatrix} 3 \\ 2 \end{pmatrix}$	$\begin{pmatrix} 2 \\ 3 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 6 \end{pmatrix}$
	hk -means	0.780	0.773	0.766	0.762
	FV 256	0.795	0.791	0.784	0.763
	s/vid	8.9	5.9	4.8	3.8
	Frame/s [†]	22.1	33.5	40.7	51.8
MBHx	(Frames/block samplerate)	$\begin{pmatrix} 6 \\ 1 \end{pmatrix}$	$\begin{pmatrix} 3 \\ 2 \end{pmatrix}$	$\begin{pmatrix} 2 \\ 3 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 6 \end{pmatrix}$
	hk -means	0.774	0.767	0.769	0.758
	FV 256	0.796	0.794	0.788	0.771
	s/vid	9.4	6.1	5.0	3.9
	Frame/s [†]	20.9	32.1	39.4	50.7
MBHy	(Frames/block samplerate)	$\begin{pmatrix} 6 \\ 1 \end{pmatrix}$	$\begin{pmatrix} 3 \\ 2 \end{pmatrix}$	$\begin{pmatrix} 2 \\ 3 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 6 \end{pmatrix}$
	hk -means	0.763	0.757	0.752	0.741
	FV 256	0.787	0.785	0.772	0.750
	s/vid	9.4	6.1	5.0	3.9
	Frame/s [†]	20.9	32.1	39.4	50.7
MBH	(Frames/block samplerate)	$\begin{pmatrix} 6 \\ 1 \end{pmatrix}$	$\begin{pmatrix} 3 \\ 2 \end{pmatrix}$	$\begin{pmatrix} 2 \\ 3 \end{pmatrix}$	$\begin{pmatrix} 1 \\ 6 \end{pmatrix}$
	hk -means	0.791	0.788	0.787	0.772
	FV 256	0.804	0.803	0.800	0.775
	s/vid	13.7	8.3	6.5	4.6
	Frame/s [†]	14.3	23.6	30.1	42.9

We keep video volumes from which descriptors are extracted the same for all sampling rates

[†] Frames/s is measured in terms of the total number of frames of the video, not in terms of how many frames are actually processed during descriptor extraction

at sample rate of 6. Again, a good speed-up is obtained by subsampling. While descriptor extraction takes 8.9 s when using every frame, a sampling rate of 2 yields a factor 1.5 speed-up while sampling every 6 frames yields a factor 2.34 speed-up.

The remaining rows of Table 3 present results obtained with different combinations of the vertical and the horizontal components of the MBH (Sect. 3.1.2). Note that when calculating both components of the MBH features, the optical flow has to be calculated only once, so computation time is faster than simply adding the times of MBH_x and MBH_y.

We observe a particular order of accuracy among these three combinations: using the only horizontal component (MBH_x) always results in a higher accuracy than using the only vertical component (MBH_y), independently of whether Fisher Vectors or *hk*-means is used as word assignment method. This sharp difference is probably due to the fact that in the test videos the horizontal motion is more frequent than the vertical one. Moreover, as expected, late fusion of the two components (MBH), always outperforms using MBH_x only. Concerning the drop of accuracy depending on the sample rate, for all the three descriptor combinations (MBH_x, MBH_y, MBH) and both word assignment methods (Fisher Vectors and *hk*-means), the accuracy loss as a function of the sample rate is similar to what happens with HOF and much higher than HOG. We believe that this is due to the fact that HOG are basically “static” features, representing the appearance of a given image window independently of possible motion information. As a consequence, they are less affected by optical flow errors [which is used to compute both HOF and MBH(*x/y*)] and better exploit the redundancy of consecutive video frames.

As for HOG and HOF and also for MBH(*x/y*) and MBH, we observe a significant computational efficiency gain using subsampling. For instance, sampling every 6 frames yields a factor of 2.4 speed-up for MBH(*x/y*) and a factor of 3 speed-up for MBH with respect to using all the frames.

To conclude, HOG descriptors can be sampled every 6 frames with negligible loss of accuracy yielding a speed-up of a factor 2. HOF and MBH descriptors can be sampled every 2 frames with negligible loss of accuracy yielding a speed-up of a factor 1.5 and 1.7 respectively. When speed is more important than accuracy, both HOF and MBH descriptors can also be sampled every 6 frames leading to 1–3% accuracy loss while gaining a significant speed-up of a factor 2.3–3.

4.5 Choice of optical flow

The results reported in the previous section show that both the HOF and the MBH(*x/y*) descriptors are much more expensive to extract than the HOG descriptors (Table 3). This is because calculating the optical flow is computationally expensive. In addition, not much research has been done

Table 4 Comparison of different optical flow methods used to compute HOF features

	Horn–Schunk	Lucas-Kanade	Färneback
<i>hk</i> -means	0.780	0.750	0.652
FV 256	0.795	0.747	0.641
s/video	8.8	7.2	19.0
Frame/s	22	27	10

Results obtained with no frame subsampling and at full original spatial resolution (320 × 240 pixels)

Table 5 Comparison of different optical flow methods used to compute HOF features

	Horn–Schunk	Lucas-Kanade	Färneback	Brox 04	Brox 11
<i>hk</i> -means	0.713	0.681	0.529	0.548	0.552
FV 256	0.718	0.697	0.542	0.638	0.652
s/video	2.9	2.8	0.76	7.2	12.4
Frame/s	68	69	257	27.4	16

Results obtained subsampling a frame every 6 and at reduced spatial resolution (80 × 60 pixels)

on how different optical flow methods affect HOF/MBH descriptors. Therefore, in this experiment we evaluate five available optical flow implementations to investigate both their computational efficiency and accuracy. In particular, we compare: (1) Farneback [15] from OpenCV using the mexopencv interface, (2) Lucas-Kanade [28] and (3) Horn–Schunk [17] from the Matlab Computer Vision Systems Toolbox, (4) Brox 04 [5] and (5) Brox 11 [6] using the available author’s code.

Results are presented in Tables 4 and 5. Specifically, while in Table 4 we used the same setting adopted in the other experiments of this paper, in Table 5 we downscaled the frame resolution of all the videos by a factor of 4 (i.e., using 80 × 60 pixel frames) and we subsampled every 6 frames (see Sect. 4.4). This scale and time subsampling was necessary to process our large video dataset with both Brox 04 and Brox 11, two state-of-the-art dense optical flow methods not able to process videos in real time. In fact, processing all the frames of our 6,600 videos at full spatial resolution with Brox 11 would require a few months.

With the original frame resolution (Table 4), and with both *hk*-means and Fisher Vectors, the three computationally feasible optical flow methods have the same ranking in terms of accuracy. For the Fisher Vector, Horn–Schunk performs best at an accuracy of 0.795, followed by Lucas-Kanade at an accuracy of 0.747, while the method of Farneback performs relatively poorly with an accuracy of 0.641. These results show that the optical flow method is crucial to the performance of the HOF descriptor: the choice of optical flow affects the results by up to 15%(!).

In terms of computational efficiency, Lucas-Kanade is the fastest at 27 frames/s, followed by Horn–Schunk at 22 frames/s, while Farneback is slower with 10 frames/s. However, while Lucas-Kanade is faster, its trade-off between efficiency and accuracy is not good: As seen in Table 3 Horn–Schunk with a frame sampling rate of 2 outperforms the Lucas-Kanade results in Table 4 in both speed (33 frames vs. 27 frames) and accuracy (0.77 vs. 0.75).

Table 5 reports results when we subsample frames and reduce the frame size by a factor 4, enabling comparison with the Brox methods. Note that for a fair comparison these times include the computation for reducing the frame sizes (although these times are negligible compared to the total description extraction time). It can be seen that both Brox methods are better than Farneback, but surprisingly not better than the Horn–Schunk and Lucas-Kanade method. One explanation is that this is due to the low resolution of the frames, which makes dense optical flow extraction not sufficiently accurate. Another possibility is that optical flow methods performing better on optical flow benchmarks are not necessarily optimal for use in classification; reducing mistakes in most parts of the flow may introduce artefacts elsewhere that negatively affect results in a classification framework.

In terms of computational efficiency, Brox 11 is the slowest, followed by Brox 04: even subsampled on reduced frames Brox 04 still processes only 27 frames/s. In contrast to results without downsampling, Farneback is here the fastest method. Apparently, there is some overhead in the Matlab optical flow implementations.

To conclude, the choice of optical flow method drastically influences the power of the resulting HOF descriptor and it is not necessarily correlated with the performance on optical flow benchmarks. In addition, many optical flow methods aim for accuracy rather than computational efficiency (e.g. Sun et al. [41] provide a very good overview for accuracy but do not report computational efficiency). Indeed, except the Horn–Schunk, Lucas-Kanade, and Farneback methods we did not find any other freely available optical flow method fast enough for use in our classification pipeline. Our evaluation shows that the Horn–Schunk method has the best trade-off between accuracy and computational efficiency and that subsampling every two frames works better than switching to Lucas-Kanade optical flow. Horn–Schunk is therefore the current method of choice.

4.6 Recommendations for practitioners

Based on the results of the previous experiments, we can now give several recommendations when accuracy or computational efficiency is preferred. For calculating Optical Flow, Sect. 4.5 showed that the Matlab implementation of Horn–Schunk is always the method of choice. In terms of frame

sampling rate, for HOG descriptors we always recommend a sampling rate of every six frames. For HOF descriptor, if one wants accuracy we recommend a sampling rate of every two frames and if one wants computational efficiency we recommend a sampling rate of 6. The same holds for MBH(x/y) descriptors. For the word assignment method, the Fisher Vector is the method of choice for accuracy. For computational efficiency there are two candidates: hierarchical k -means and the Random Forest. Observe first that the descriptor extraction time is the most costly phase of the pipeline: Extracting HOF descriptors with a sampling rate of 6 frames takes 3.8 s/video to compute. And while the Random Forest is five times faster than hierarchical k -means, the difference is only 0.41 s/video, which is very small compared to the descriptor extraction phase. Furthermore, Table 1 showed a significant drop of accuracy from 0.780 for hierarchical k -means to 0.732 for Random Forests [and a similar drop of accuracy is observed with MBH(x/y)]. Therefore we recommend using hierarchical k -means for a fast video classification pipeline.

We found that late fusion of the classifier outputs gave slightly better results than early fusion of the descriptors (e.g. concatenating HOG and HOF). Hence in our recommendations we perform a late fusion with equal weights.

We tested different descriptor combinations, using equal-weights-based late fusion and with the goal of selecting: (1) the most accurate set of descriptors, possibly taking into account the complementarity of appearance/motion information of different features, and (2) the fastest solution with a sufficiently good accuracy degree. The final recommended pipelines are visualised in Figs. 7 and 8.

The most accurate pipeline (Fig. 7) combines all the descriptors we adopted in this paper: HOG, HOF, MBH x and MBH y . HOG are extracted using all the frames, while HOF and MBH(x/y) are extracted with a sampling rate of 2. The word assignment method used in this case is the Fisher Vector. Using this pipeline we can process 11 frames/s (for video frames of 320×240 pixels) at an accuracy of 0.818 on UCF50. Conversely, our recommended pipeline for computational efficiency (Fig. 8) is based on late fusion of only HOG and HOF, both extracted with a sampling rate of 6 and using hk -means. This second pipeline can process 28 frames/s at a respectable accuracy of 0.790.

4.7 Comparison with state-of-the-art

In this section we compare our descriptors to the state-of-the-art. Results of several recent works are given in Table 6. This comparison is done in terms of accuracy only, as most compared methods evaluate accuracy only. *This paper* in Table 6 indicates the late fusion of all the descriptors [HOG, HOF, MBH(x/y)]: see Sect. 4.6 and Fig. 7.

As can be seen, the method of [45] yields the best results. This method is a combination of Dense Trajectories and

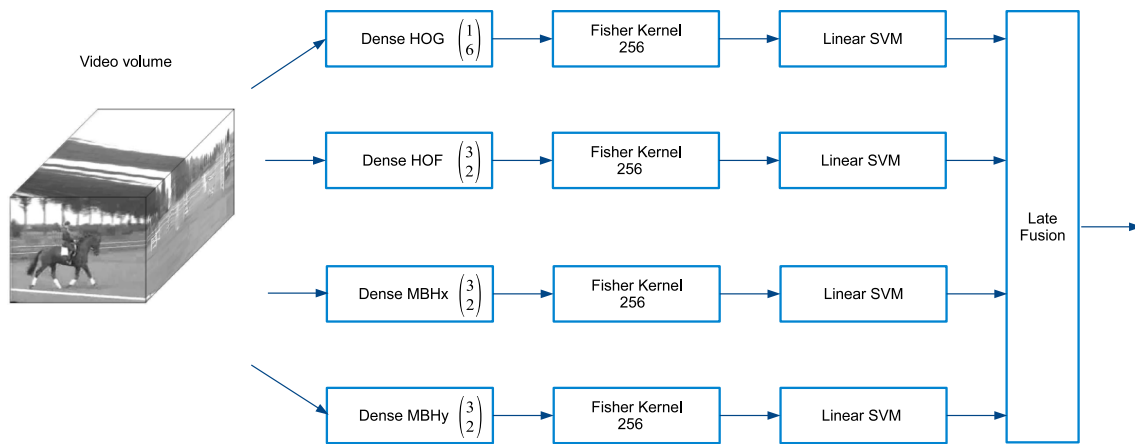


Fig. 7 Recommended pipeline for accurate video classification. This pipeline yields an accuracy of 0.818 on UCF50 while processing 9 frames/s




Fig. 8 Recommended pipeline for realtime video classification. This pipeline yields an accuracy of 0.790 on UCF50 while processing 28 frames/s

Table 6 Comparison with the state-of-the-art

Method	Accuracy (%)
Wang et al. [45]	0.856
This paper	0.818
Reddy et al. [32]	0.769
Solmaz et al. [40]	0.737
Everst et al. [14]	0.729
Klipper-Gross et al. [23]	0.727

STIP features [25]. As our results are better than [25], we expect that a combination of dense trajectories with our method would increase results further. In general, our method yields good performance compared to many recently proposed methods, which shows that we provide a strong implementation of densely sampled HOG, HOF and MBH(x/y) descriptors.

5 Conclusion

This paper presented an evaluation of the trade-off between computational efficiency and accuracy for video classification using a Bag-of-Words pipeline with HOG, HOF and MBH descriptors. Our first contribution is a strong and fast

Matlab implementation of densely sampled HOG, HOF and MBH descriptors, which we make publicly available.

In terms of visual word assignment, the most accurate method is the Fisher Kernel. Hierarchical k -means is more than six times faster while yielding an accuracy loss of less than 2% and is the method of choice for a fast video classification pipeline. HOG descriptors can be subsampled every six frames with a negligible loss in accuracy, while being two times faster. HOF and MBH descriptors can be subsampled every two frames with negligible loss in accuracy, being 1.5–1.7 times faster. When speed is essential, HOF and MBH descriptors may be subsampled every six frames.

For the HOF and MBH descriptors, we showed that the choice of optical flow algorithm has a large impact on the final performance. The difference between the best method, Horn–Schunk, and the second best method, Lucas–Kanade, is already 5%, while the difference with Färneback is a full 15%. Brox 04 and Brox 11 are computationally very demanding, and cannot be used in a real-time video classification scenario.

Compared to the state-of-the-art, the Dense Trajectory method of [45] obtains better results. Nevertheless, the huge difference for the choice of optical flow methods suggests this would also influence dense trajectories. Furthermore, Dense Trajectories still benefit from a combination with normal HOG, HOF and MBH descriptors [21, 45]. Finally, com-

parisons with other recent methods on UCF50 show that we provide a strong implementation of dense HOG, HOF and MBH descriptors to the community.

Acknowledgments This work was supported by the European 7th Framework Program, under grant xLiMe (FP7-611346) and by the FIRB project S-PATTERNS.

References

- Arandjelović R, Zisserman A (2012) Three things everyone should know to improve object retrieval. In: CVPR
- Baker S, Scharstein D, Lewis JP, Roth S, Black MJ, Szeliski R (2011) A database and evaluation methodology for optical flow. *Int J Comput Vis* 92:1–31
- Bay H, Ess A, Tuytelaars T, Van L (2008) Speeded-Up Robust Features (SURF). *Comput Vis Image Underst* 110:346–359
- Breiman L (2001) Random forests. *Mach Learn* 45(1):5–32
- Brox T, Bruhn A, Papenbergh N, Weickert J (2004) High accuracy optical flow estimation based on a theory for warping. In: ECCV, pp 25–36
- Brox T, Malik J (2011) Large displacement optical flow: descriptor matching in variational motion estimation. *PAMI* 33(3):500–513
- Butler DJ, Wulff J, Stanley GB, Black MJ (2012) A naturalistic open source movie for optical flow evaluation. In: ECCV
- Chang C-C, Lin C-J (2011) LIBSVM: A library for support vector machines. *ACM Trans Intell Syst Technol*. <http://www.csie.ntu.edu.tw/~cjlin/libsvm>
- Chatfield K, Lempitsky V, Vedaldi A, Zisserman A (2011) The devil is in the details: an evaluation of recent feature encoding methods. In: BMVC
- Csurka G, Dance CR, Fan L, Willamowski J, Bray C (2004) Visual categorization with bags of keypoints. In: ECCV international workshop on statistical learning in computer vision, Prague
- Dalal N, Triggs B (2005) Histograms of oriented gradients for human detection. In: CVPR
- Dalal N, Triggs B, Schmid C (2006) Human detection using oriented histograms of flow and appearance. In: ECCV
- Dollár P, Rabaud V, Cottrell G, Belongie S (2005) Behavior recognition via sparse spatio-temporal features. In: VS-PETS
- Everts I, van Gemert J, Gevers T (2013) Evaluation of color STIPs for human action recognition. In: CVPR
- Farneback G (2003) Two-frame motion estimation based on polynomial expansion. In: Scandinavian conference on image analysis
- Geurts P, Ernst D, Wehenkel L (2006) Extremely randomized trees. *Mach Learn* 63(1):3–42
- Horn B, Schunck B (1981) Determining optical flow. *Artif Intell* 17:185–203
- Jaakkola T, Haussler D (1999) Exploiting generative models in discriminative classifiers. In: NIPS
- Jégou H, Douze M, Schmid C, Pérez P (2010) Aggregating local descriptors into a compact image representation. In: CVPR, pp 3304–3311
- Jurie F, Triggs B (2005) Creating efficient codebooks for visual recognition. In: ICCV
- Karaman S, Seidenari L, Bagdanov A, del Bimbo A (2013) L1-regularized logistic regression stacking and transductive CRF smoothing for action recognition in video. In: ICCV workshop on action recognition with a large number of classes
- Kläser A, Marszałek M, Schmid C (2008) A spatio-temporal descriptor based on 3d-gradients. In: BMVC
- Kliper-Gross O, Gurovich Y, Hassner T, Wolf L (2012) Motion interchange patterns for action recognition in unconstrained videos. In: ECCV
- Kuehne H, Jhuang H, Garrote E, Poggio T, Serre T (2011) HMDB: a large video database for human motion recognition. In: ICCV
- Laptev I, Marszałek M, Schmid C, Rozenfeld B (2008) Learning realistic human actions from movies. In: CVPR
- Lazebnik S, Schmid C, Ponce J (2006) Spatial pyramid matching for recognizing natural scene categories. In: CVPR. Beyond Bags of Features
- Lowe DG (2004) Distinctive image features from scale-invariant keypoints. *IJCV* 60:91–110
- Lucas B, Kanade T (1981) An iterative image registration technique with an application to stereo vision. In: International joint conference on artificial intelligence
- Maji S, Berg AC, Malik J (2008) Classification using intersection kernel support vector machines is efficient. In: CVPR
- Moosmann F, Nowak E, Jurie F (2008) Randomized clustering forests for image classification. *IEEE Trans Pattern Anal Mach Intell* 9:1632–1646
- Perronnin F, Sanchez J, Mensink T (2010) Improving the Fisher kernel for large-scale image classification. In: ECCV
- Reddy K, Shah M (2013) Recognizing 50 human action categories of web videos. *Mach Vis Appl* 24(5):971–981
- Sánchez J, Perronnin F, Mensink T, Verbeek JJ (2013) Image classification with the fisher vector: theory and practice. *Int J Comput Vis* 105(3):222–245
- Sangineto E (2013) Pose and expression independent facial landmark localization using dense-SURF and the Hausdorff distance. *IEEE Trans Pattern Anal Mach Intell* 35(3):624–638
- Schuld C, Laptev I, Caputo B (2004) Recognizing human actions: a local svm approach. In: ICIP
- Scovanner P, Ali S, Shah M (2007) A 3-dimensional sift descriptor and its application to action recognition. In: ACM MM
- Sivic J, Zisserman A (2003) Video Google: a text retrieval approach to object matching in videos. In: ICCV
- Smeaton AF, Over P, Kraaij W (2006) Evaluation campaigns and TRECVID. In: ACM SIGMM international workshop on multimedia information retrieval (MIR)
- Snoek CGM, Worring M, Gemert J, Geusebroek J, Smeulders A (2006) The challenge problem for automated detection of 101 semantic concepts in multimedia. In: ACM MM
- Solmaz B, Assari SM, Shah M (2013) Classifying web videos using a global video descriptor. *Mach Vis Appl* 24(7):1473–1485
- Sun D, Roth S, Black M (2014) A quantitative analysis of current practices in optical flow estimation and the principles behind them. *Int J Comput Vis* 106:115–137
- Uijlings JRR, Smeulders AWM, Scha RJH (2010) Real-time visual concept classification. *IEEE Trans Multimed* 12(7):665–681
- Vedaldi A, Fulkerson B (2010) VLFeat—an open and portable library of computer vision algorithms. In: ACM MM
- Viola P, Jones M (2001) Rapid object detection using a boosted cascade of simple features. *Proc CVPR* 1:511–518
- Wang H, Kläser A, Schmid C, Liu C (2013) Dense trajectories and motion boundary descriptors for action recognition. *Int J Comput Vis* 103:60–79
- Wang H, Ullah M, Kläser A, Laptev I, Schmid C (2009) Evaluation of local spatio-temporal features for action recognition. In: BMVC