Adaptive features for object classification

Heydar Maboudi Afkham, Stefan Carlsson, and Josephine Sullivan

Computer Vision and Active Perception Lab., KTH, Stockholm, Sweden {heydarma,stefanc,sullivan}@csc.kth.se http://csc.kth.se/cvap

Abstract. This work presents a method for building adaptive local/semiglobal features using a set of already extracted features. While for most methods local features are extracted independently of the task in hand, these features tend to change their representations in favor of different hypotheses to find the best possible representation. The features introduced in this paper take advantage of the part-based models at the feature level by combining the near by local features. This combination can either be local, which results in a more generic set of features suitable for bag-of-visual-words (BOVW) models or be semi-global, which results in a set of more object dependent features which are referred as parts. These representations capture the local variations around the local feature. At classification time, the best possible representation of these features is found and used in the calculations. This selection is done based on a latency defined at the feature level. The goal of this paper is to test how the adaptive features can improve the feature level likelihoods. The focus of the experiments of this paper is showing 1) how adaptive feature perform in BOVW scenarios and 2) how replacing single features with equivalent adaptive features improves the likelihoods obtained from them. The experiments of this paper are done on several classes of MSRCv1 and v2 datasets and it is shown that the method outperforms the baselines in all cases and the results are comparable to the state-of-the-art methods using BOVW models.

Keywords: Feature inference, Latent models, Clustering

1 Introduction

Local features are considered to be the building blocks of many computer vision and machine learning methods and their quality highly effects the method's outcome. There are many popular methods for extracting local features from images. Among them one can name sift [10], hog[2] and haar[16] features, which are widely used for object detection and recognition [4, 7] and texture features such as maximum response filter-banks [14] and MRF[13] which are used for texture recognition and classification. For most methods feature extraction is done independently from the method's task. For example in a normal inference problem a model tries to decide between two different classes. Usually for both

2 Lecture Notes in Computer Science: Authors' Instructions

hypothesises the same feature vector is fed to the model. In other words once the features are extracted, they remain constant through the whole process.

Computer vision methods deal with local features in different ways. Some such as boosting based object detectors [7, 16] and markov random fields [6], depend on how discriminative single features are and some, such as bag-of-visualwords [12], demand less discriminative features and depend on groups of features seen in a specific region on the image. For the methods that depend on the discriminative properties of local features the inference done at the feature level plays a critical role in the outcome of the method. Improving the quality of feature level inference can highly improve the quality of the object level inference done by most of such methods.

Many studies have shown that use of higher order statistics, ex. the joint relation between the features, can highly improve the quality of the features. Capturing joint relations is popular with the bag-of-words methods [9,8] since they deals with modeling joint relation between a finite number of data clusters. Unfortunately not many studies have focused on modeling joint relations in nondiscretized data to create features that capture joint relations. A recent study on this matter is done by Morioka *et al.* [11]. In their study they introduced a mechanism for pairing two sift feature vectors together and creating a Local Pairwise Codebook by clustering them. As shown in their work the clusters produced using these joint features are more informative than the clusters produced using single features. The idea behind their work is similar to the work in this paper while the methodology of this work does not limit the number of feature vectors used in creating more complex features.

The method in this paper uses the assumption that a set of features are extracted from the image and a relation is known between them that can be captured by a graph. For example the features can come from several patches in the image and their spatial relation can be presented as a graph. These features can be extracted using any feature extraction method. The basic idea behind this paper is to use local features and their relations to introduce a new set of dynamic and changeable intermediate semi-global features in terms of latent variables. These intermediate features will be referred as feature clouds. The latent variables enable the feature to change its representation in different scenarios and their value is determined by an optimization procedure to make it more discriminative for learning algorithms. This dynamic property of feature cloud provides a good ground for introducing more discriminative features than the ones previously extracted from the image. The performance of these features is analyzed in two different experiments on MSRC v1 and v2 [17] datasets. The first experiment deals with discriminative analysis of feature clouds and their inference at the feature level. These analysis show how more complex feature have an easier time locally identifying object regions in comparison more simple features. In the second experiment the features are employed in bag-of-visual-words model and it is shown that they have a better performance than the existing methods.



Fig. 1. (Left) A node u (a patch from an image) is connected to its neighboring nodes (close by patches). Here tree node are selected (red patches) as a latent configuration. The quantitative value of this configuration is calculated as $(\phi_u, \phi_{v_2}, \phi_{v_5}, \phi_{v_7})$.(**Right**) For any given point z in the feature space the closest vector within \mathbf{F}_u is selected as f_z^* through an optimization process. This shows how z can influence the value of \mathbf{F}_u .

The outline of this paper is as follow. The related works to this paper are discussed in section 2. The feature clouds are discussed in detail in sections 3, 4 and 5. Finally section 6 discusses the behaviour of the feature clouds in some different scenarios.

2 Related Works

The goal of this paper is to present a method that takes advantage of part-based models at the feature level to come up with a set of intermediate features with discriminative properties. In this section a brief review of part based methods is provided and their differences with the presented method is pointed out. Later an example of studies that show how local features can effect the overall inference is discussed.

Part-based models have been widely used in object detection applications. A good example of such application can be found in the work of Felzenszwalb *et al.*[3]. These models consist of a fixed root feature and several part feature with their position as latent variables in relation to the root feature. The part features are learnt either using exact annotation [5] or as the result of an optimization problem [3]. Because of this latency the model can have many configurations an usually the best configuration is chosen among many due to the task in hand. In these models the part features are used to estimate a better confidence for the root feature.

Taking the part-based models to the feature level comes with several difficulties. To begin with there is larger variation at the feature level compared to the object level. Here each local feature can play the role of a root feature and completely different features can be equally good representatives for an object class. As an example consider the features obtained from the wheel and the door



Fig. 2. This figure shows how a car representative anchor point pulls out features from two different images. It can be seen that the features are avoiding cow texture since this texture is not a good candidate for supporting being a car hypothesis.

of a car, one wishes for a car model to return a high likelihood for both features despite their differences. Also there is no right or wrong way to look at local features. In this work the root features and their parts are calculated as the result of a clustering process. Since there is no best configuration for these features each root feature can have several good part configurations. These configurations will capture the local variation around the local feature and will later be used for training non-linear discriminative classifiers.

As mentioned in section 1 many methods benefit from discriminative behaviours of local features. A good example of such benefit can be seen in [6] where the authors Sanjiv Kumar *et al.* show how replacing generative models with discriminative models benefits the MRF solvers and improves their final results. Similar examples can be widely found in numerous computer vision studies. The key difference between these works and this method is the fact the features can change their value to result in a more discriminative behaviour.

3 Feature Clouds

4

To define the feature clouds, let G(V, E) be a graph with the extracted features as its nodes and the relation between them encoded as its edges. Also for each node, say u, let ϕ_u denote the feature vector associated with this node and \mathbf{N}_u denote the its neighbors. A cloud feature, with its root at node u and m latent parts, is the set of all possible vectors that are created by concatenating the feature vector of node u and the feature vectors of m nodes selected from \mathbf{N}_u . This set is formally defined as

$$\mathbf{F}_{u} = \{ (\phi_{u}, \phi_{v_{1}}, ..., \phi_{v_{m}}) | v_{i} \in \mathbf{N}_{u} \},$$
(1)

where $(\phi_u, \phi_{v_1}, ..., \phi_{v_m})$ is the concatenation of the feature vector of the nodes $u, v_1, ..., v_m$. In other words all possible configurations that can be made using u and its parts exist in the set \mathbf{F}_u . This can also be seen as the space of all variations around node u. These configurations are shown in figure 1 (Left). In this figure a

node u (a patch extracted from an image) is connected to its neighbours (its close by patches). In this configuration the three selected neighbors v_2, v_5, v_7 are shown using solid line edges. and the resulting feature vector for this configuration is $(\phi_u, \phi_{v_2}, \phi_{v_5}, \phi_{v_7})$. Here the set \mathbf{F}_u will contain all possible similar feature vectors made by selecting three nodes among the eight neighbors.

In practice only one of the vectors within \mathbf{F}_u is selected and used as its quantitative value. Since the size of this set can grow large this value is selected in an optimization process. This value is determined in relation to a fixed target point in the feature space. For any arbitrary point z in the feature space, the value of \mathbf{F}_u is fixed as the best fitting vector in \mathbf{F}_u to this point. This can be written as

$$f^{\star}_{(z,u)} = \operatorname*{arg\,min}_{f \in \mathbf{F}_u} \{ d(f,z) \}.$$

$$\tag{2}$$

Here d(.) is the euclidean distance. Here the point z is used as an anchor point in the feature space for fixing latent variables of the feature cloud. Figure 1 (Right) illustrates how f_z^* is selected. In this work z plays an important role in the classification process. Since for every given z the value of \mathbf{F}_u changes, z can be seen as a tool for pulling out different properties of \mathbf{F}_u . In the classification stage each \mathbf{F}_u will be fit to different z values, learnt during the training process, to verify whether \mathbf{F}_u contains configurations that belong to the object or not. The result of this selection can be seen in figure 2. In this figure a car related region is being extracted from a car image and a cow image. It can be seen that the positioning of the features completely differ in the two images and more importantly the features seem to be avoiding the texture of the cow. This is done because the optimization process tends to find the best available matches in the image to support a certain hypothesis (being on a car) and regions on the cow don't seem to be good candidates.

By having a prior knowledge of how the nodes are distributed in \mathbf{N}_u a partitioning can be imposed on this set. This partitioning will later be referred as the *architecture* of the cloud feature. This partitioning is designed in a way that each latent part comes from one partition. This partitioning can slightly reduce the complexity of the optimization.

Efficient solving of the optimization problem 2 can have a large effect on the performance and the running time of the methods using feaure. Assuming that the size of the extracted features is fixed and distance is measured using euclidean distance, the complexity of implemented method is calculated as $O(|\mathbf{N}_u|m)$, where is m is the number of latent parts. By introducing an architecture and partitioning \mathbf{N}_u , this complexity will be reduced to $O(|\mathbf{N}_u|)$. In other words this problem can be solved by visiting the neighbors of each node at most m times. It is possible to design more efficient algorithms for solving this optimization problem and this will be in the focus of the future works of this paper.

4 Latent Classifiers

In this problem the task of a classifier is to take a feature cloud and classify it into either being from the object or not. For a set of labeled features,



Fig. 3. This figure shows a summery of the classification process. Initially a node u is selected and its corresponding cloud is calculated. With the known anchor points z_1, \ldots, z_n , the value for the basis functions ϕ_i are calculated and using them a classification for the node u is achieved.

 $\{(\mathbf{F}_{u_1}, y_1), \dots, (\mathbf{F}_{u_N}, y_N)\}$ gathered from the training set, the goal is to design a function c that uses one or several configurations within cloud features to minimize the cost function

$$\sum_{i=1}^{N} |c(\mathbf{F}_{u_i}) - y_i|.$$
 (3)

A key difference between this approach and other available approaches is the fact that the local variations around the local feature are also modeled with optimization of c. This means that not only the model uses the value of the root feature but it also uses the dominant features appearing around the root feature regardless of their spatial position.

The optimization problem 3 is approximated by defining the function c as a linear basis regression function with model parameters $z_1, z_2, ..., z_M, W$. The values $z_1, z_2, ..., z_M$ are M anchor points in the feature space for fixing the latent variables capturing different configurations of the features and $W = (w_0, ..., w_M)$ contains the regression weights, these parameters are learnt during the training stage. Using these parameters, the regression function c is defined as

$$c(\mathbf{F}_{u}; z_{1}, \dots, z_{M}, W) = \sum_{m=1}^{M} w_{m} \Phi_{z_{m}}(\mathbf{F}_{u}) + w_{0}.$$
 (4)

Here the basis function Φ_{z_m} measures how good \mathbf{F}_u can be fit to z_m . This basis function can be written as any basis function for example a Gaussian basis function is defined as

$$\Phi_{z_m}(\mathbf{F}_u; s) = \exp(-\frac{d(f_{(z_m, u)}^{\star}, z_m)^2}{2s^2}).$$
(5)

Equations 4 and 5 clearly show that the decision made for \mathbf{F}_u depend both on the different values in \mathbf{F}_u and how good it can be fit to z_m values. Due to the dynamic section, \mathbf{F}_u can be fit to different z_m values which makes the scoring



Fig. 4. (Left) The Arc1 only depends on the patch itself. (Middle) The Arc2 (Arc3) Depends on the central patch together with a selection of patches in the scale details (context) of the central patch. (Right) The Arc4 Depends on the central patch and a selection of patches from from details and the context of the central patch.

processes harder for negative samples. The summery of this classification process can be seen in figure 3. During the training stage solving equation 4 to obtain W is straight forward once the z_m values are known.

The local features come from different regions of an object and these regions are no visually similar. This fact results in a large variation on the data used for training. This variation can not be captured by only one set of M configurations. To solve this problem a mixture model of K sets each with M different configurations is considered and each set is associated with a different regression model. When classifying a cloud feature, it is initially assigned to the model that minimizes the over all fitting cost, defined as

$$\arg\min_{k\in\{1...K\}} \{\sum_{m=1}^{M} d(f^{\star}_{(z^{(k)}_{m},u)}, z^{(k)}_{m})\}.$$
(6)

Here the model is chosen based on how good the feature \mathbf{F}_u is fit to all the configurations within the model.

5 Learning the parameters

The goal of the learning algorithms is to determine the K configurations sets together with the regression function. It is possible to design an optimization method to estimate all configurations together with the regression function, but such optimization is more suitable for an object level classification since the data contains less variation at that level. In this work the optimization is done two separate steps. The first step uses a generative method for finding the K sets of configurations. This generative method is an adaptation k-means algorithm with the cost function

$$\arg\min_{S} \sum_{k=1}^{K} \sum_{\mathbf{F}_{u} \in S_{k}} \left(\sum_{m=1}^{M} d(f_{z_{m}^{(k)}}^{\star}, z_{m}^{(k)}) \right).$$
(7)

This adaptation of k-means divides the features into K clusters and the elements of each cluster have a strong connection by sharing the M different configurations. This optimization can be solved using iterative methods used for solving 8



Fig. 5. The heat-maps shown in this figure are produced by mapping the scores of all the individual features on a plane. These show how the models locally learn the structures of the object.

the k-means problem. This method will be referred as L-KMEANS(K,M). To optimize this cost function initially the feature clouds are partitioned into K subsets based on how close they are to the M anchor points. After this partitioning the obtained values from the clouds of each partition are used for updating the anchor points. This process is continued until a convergence is achieved.

To determine the parameters initially L-KMEANS(K,M) procedure is ran over the positive features. This way the strong configurations appearing in the training set are formulated in terms of cluster centers resulted by the procedure. These cluster centers can be used for labeling all features. For each cluster the variations in the negative features is captured by running L-KMEANS(1,M) on the negative features assigned to that cluster. Finally, after identifying both positive and negative configurations, these configurations are used to fix the values of cloud features and the and regression function 4 is optimized to separate the positive features from the negative features.

6 Experiments and Results

The aim of the proposed methodology is to define a level of latency at the feature level to extract more discriminative local/semi-global features. This latency in selection of features can benefit many different pattern recognition tasks. To analyze the effect of this latency two different scenarios are considered. A scenario is to employ the feature clouds in a bag-of-visual-words (BOVW) model. The idea behind the bag of visual words model is the fact that objects are built using a series of local structures that are shared between different objects and a histogram of such local structures can lead to the identification of the objects. Meanwhile, another challenge is to identify unique structures (parts) on the objects and use them for the classification of the object. The fact that the parts are defined in a more global sense and are more related to the object class, makes them bad candidates for BOVW models. Instead the quality of these parts can

9

be evaluated based on how *good* they appear on the object. Both these scenarios are discussed in this section.

When defining the feature clouds in section 3, each cloud is defined based on a set of neighboring features. On an image the neighboring features can be located as features withing a spatial radius of δ around the root feature. In this context δ will be referred as the flexibility parameter. This parameter can be used to build clouds for different purposes. If the value of δ is low then the joint representation become more local therefor more suited for BOVW models. A disadvantage of building such local joint presentations is the fact that their performance of becomes close to the performance of single features upon which the joint features were built. When the value of δ grows large, the joint representations can be selected from a larger range of features on the object, which makes the selection more dependent on object class and the features become more class and viewpoint dependent. Having a semi-global feature provides a series of features that are independent and are rarely detected off the object.

The experiments conducted in this paper are not designed to be compared with state of the art object detectors but to test the hypothesis proposed in the paper. The main idea behind the experiments is to evaluate these local features with different architectures and compare them with the baseline which only contains the root feature. The evaluation is straight forward, each feature is scored using the equations 4 and 6. Figure 5 shows how the heat-map of this score looks like for different images. For these figures the score was calculated for individual nodes and mapped on a plane. In the feature inference problem the results are presented in terms of different precision recall curves of these values. The experiments are conducted on several classes of MSRC v1 and v2[17] datasets.

6.1 Feature Clouds as semi-global features

Several parameters control the behaviour of the cloud features. As mentioned in section 3 the architecture of the cloud features imposes a strong prior on how the latent parts are placed together. The architecture controls the complexity of the feature by controlling the number of its latent parts. The flexibility of these features is controlled by the number of neighbors each node has in graph G. The larger the size of the neighbours the wider the search space is for the latent parts. The goal of the experiments to analyze how the complexity and flexibility of the cloud features effects their discriminative behavior. Unfortunately, as the flexibility and the complexity of the features increase the optimization processes in equations 7 and 4 become computationally expensive. Therefore the results are only provided on a few classes of this dataset.

The graph used in the experiments is build over the fixed size patches with 30 pixel side extracted from an image pyramid and described using a PHOG descriptor [1]. The choice of this graph is due to the future works of this paper when these features are used to build object level classifiers.

Let G = (V, E) denote the graph built over the patches extracted from the image pyramid with V containing all extracted patches and for two patches in

10 Lecture Notes in Computer Science: Authors' Instructions

V, say u and v, the edge uv belongs to E iff $|x_u - x_v| < t$ and $|s_u - s_v| < 1$. Here x_u is the spatial position, s_u is the scale level of node u and t is a given threshold which controls how patches are connected to each other.

In this work four different architectures are considered. These architectures are considered as prior information and are hard coded in the method. Although it is possible to learn the architectures, learning them requires more tools which are not in the scope of this paper. As mentioned in section 3 the architecture is imposed by partitioning the set \mathbf{N}_u . In this problem the neighbors of each node come from three different scale levels. There are, of course, many different ways that this set can be partitioned into m subsets. To reduce the number of possibilities only partitions with simple fixed scale and spatial relations to the central node are considered. Twelve subsets are formed by dividing the nodes (patches) in each scale into four quadrants. The scale levels and quadrants can be seen in the features shown in figure 4. A number of these subsets are selected to form different features architectures. Let this selection be denoted by $\overline{\mathbf{P}}_u$. Using a subset of the twelve partitions, the four architectures defined in this figure are,

- Arc 1: This architecture is created by having $\overline{\mathbf{P}}_u = \emptyset$. This architecture uses the descriptor of the central node as the descriptor. This feature will be used as a benchmark for analyzing dynamic architectures.
- Arc 2: Let $\overline{\mathbf{P}}_u$ partition the scale level below the scale level of u into four spatial quadrants. This architecture contains the data from node u and additional information about the details in this region.
- Arc 3: Let \mathbf{P}_u partition the scale level above the scale level of u into four spatial quadrants. This architecture contains the data from node u and additional information about the context in this region.
- Arc 4: Let $\overline{\mathbf{P}}_u$ partition the scale levels both above and below the scale level of u. This architecture contains the data from u together with information about details and the context of the region u has appeared in.

Here each node can be described using each of the four architectures and the goal is to verify the most suitable architecture for the object region. To train the classifiers any feature from the positive regions is considered positive and the rest of the extracted features are considered as negative features. This should be kept in mind that the problem being solved here is equivalent taking an arbitrary patch from an arbitrary scale and location and asking whether it belongs to the object or not. Due to the noise at the feature level this problem is a hard problem to solve by nature.

Classes {Car, Face, Plane, Cow, Bike} are chosen from this dataset. In this experiment all four architectures are used with varying t value to increase the flexibility of features. Here 256 models are used and each with 5 positive and 5 negative latent configurations. The results of this experiment can be seen in figure 6. These experiments reveal several properties about the cloud features. The first property is in fact that the improvement obtained on the feature like-lihood level depend on both the base feature and the architecture. At it can be seen in this figure the base feature (red curves) has an easier time capturing



Fig. 6. Results from the MSRCv1 dataset. Five classes {car, face, plane, cow, bike} were considered from this dataset. In this experiment the value of t controlling number of neighbors was varied from the value equal to half the patch size to three times larger than the patch size.

the properties of the *car* and *face* classes in comparison with the rest of the classes. Also between these two classes the architectures have an easier time capturing the relations in face region. Meanwhile it is clearly visible that for the *bike* class both the features and architectures are failing to capture the local

properties. This figure also shows that there is no best architecture for the all the object classes and the choice of the architecture is completely object dependent. This can be seen in the likelihoods obtained from the *plane* and the *cow* classes, where the base likelihoods are similar but the responses obtained from the different architectures are different.

6.2 Feature Clouds in BOVW Models

This experiment is conducted on 9 classes of MSRC v2 dataset following the experiment setting presented by Morioka *et al.* [11] for building local pairwise codebook (LPC). In their setting sift features were sampled at every 8 pixels from the images and LPC was build by clustering them. In their framework features with distance equal or less than 8 pixels are merged to build joint features. To adapt this scenario the graph G(V, E) from section 3 is constructed over such sift features. Here V contains all the sampled sift features and for every two features, say u and v, $uv \in E$ iff $|x_u - x_v| \leq \delta = 8$. In the concept of feature clouds each the anchor points is optimized for each for each class. Here a number of anchor points are calculated for the classes (together with the background class) using the cost function 7 and put together as N anchor points $\{c_1, ..., c_N\}$. In this experiment, equation 2 was used to build a histogram for each image. To define this formally let $\{\mathbf{F}_{u_1}, \ldots, \mathbf{F}_{u_M}\}$ be M feature clouds extracted from an image and H be a N bin histogram with H[n] representing its n^{th} bin. The value of H[n] is determined as

$$H[n] = \#\{\mathbf{F}_{u_i} : \forall m \neq n, d(f^{\star}_{(c_n, u_i)}, c_n) < d(f^{\star}_{(c_m, u_i)}, c_m)\}.$$
(8)

Here \mathbf{F}_{u_i} is assigned to most fitting anchor point. Similar to [11] the histograms where classified using non-linear SVM with histogram intersection kernel. In this experiment beside the appearance of latent parts their position was also modeled. This modeling was by changing the optimization 2 to

$$f_z^{\star} = \underset{f = (f_A, f_x) \in \mathbf{F}_u}{\arg\min} \{ \alpha d(f_A, z_A) + (1 - \alpha) d(f_x, z_x) \}.$$
(9)

Here f_A and z_A contain the appearance information of the configuration and the anchor point, while f_x and z_x contain the information about the relative position of the latent parts with respect to the root feature. The value alpha was considered as a constant value equal to 0.75.

For this experiment half of the images in each class were used as training images and the other half was used for testing. The SIFT features were calculated on a dense grid after smoothing the image using VLFeat's [15] Matlab API. In this experiment feature clouds with accuracy of $84.19 \pm 2.75\%$, having 820 words in the dictionary, have performed better than the best baseline accuracy (single features) with accuracy of $83.0 \pm 2.0\%$. Also since the leading dimensions of the words, calculated for the feature clouds, corresponding to the root features can be used for labeling root features, a secondary histogram can be calculated for these features. The performance of the concatenated histogram of cloud features and

root features classified the images in the dataset with the accuracy of $85.6 \pm 1.4\%$. Our results are higher than the results reported for LPC[11] $83.9 \pm 2.9\%$ and [18] $78.3 \pm 2.6\%$ for 2^{nd} order and $80.4 \pm 2.5\%$ for 10^{th} order spatial features.

7 Conclusion and Future Works

The main objective of this work has been to investigate the improvement in discriminability obtained by substituting simple local features with more adaptive composite hierarchical structures that are computed at recognition time from a set of potential structures denoted as feature clouds. This is motivated by the fact that even at local feature level, intra class object variation is very large, implying that generic single feature classifiers that try to capture this variation will be very difficult to design. In our approach this difficulty is circumvented by the introduction of the cloud features that capture the intra class variation an feature level. The price paid is of course a more complex process for the extraction of local features that are computed in an optimization process in order to yield maximally efficient features. We believe however that this process can be made efficient by considering the dependencies and similarities between local feature variations that are induced by the global intra class object variation.

There are many ways to improve the performance and accuracy of the feature clouds and investigate their applications. As mentioned in the text, coming up with better optimization algorithms will decrease the usage cost of these features. Meanwhile designing algorithms for learning the architecture rather than hardcoding them will increase the accuracy of these features. As for the applications, these features can be used in different object detection and recognition platforms. A direct follow up of this work is using these features to build more robust object detectors for detecting object classes. Since the cloud features are results of clustering process rather than discriminative analysis, they can also be used in bag-of-words models and will result in more discriminative words and smoothed labeled regions.

Acknowledgements

This work was supported by The Swedish Foundation for Strategic Research in the project Wearable Visual Information Systems. 14 Lecture Notes in Computer Science: Authors' Instructions

References

- Anna Bosch, Andrew Zisserman, and Xavier Munoz. Representing shape with a spatial pyramid kernel. In *CIVR*, pages 401–408. Association for Computing Machinery, July 2007.
- Navneet Dalal and Bill Triggs. Histograms of oriented gradients for human detection. In CVPR, pages 886–893, 2005.
- 3. Pedro Felzenszwalb, Ross Girshick, David Mcallester, and Deva Ramanan. Object detection with discriminatively trained part based models. Technical report, 2009.
- Pedro F. Felzenszwalb, Ross B. Girshick, David McAllester, and Deva Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transac*tions on Pattern Analysis and Machine Intelligence, 32:1627–1645, 2010.
- Pedro F. Felzenszwalb and Daniel P. Huttenlocher. Pictorial structures for object recognition. *IJCV*, 61:55–79, January 2005.
- Sanjiv Kumar and Martial Hebert. Discriminative random fields. IJCV, 68:179– 201, 2006.
- Ivan Laptev. Improvements of object detection using boosted histograms. In BMVC, pages 949–958, 2006.
- 8. H. Ling and S. Soatto. Proximity distribution kernels for geometric context in category recognition. In *ICCV*, 2007.
- 9. David Liu, Gang Hua, Paul Viola, and Tsuhan Chen. Integrated feature selection and higher-order spatial feature extraction for object categorization. *CVPR*, 0:1–8, 2008.
- 10. David G. Lowe. Distinctive image features from scale-invariant keypoints, 2003.
- 11. Nobuyuki Morioka and Shin'ichi Satoh. Building compact local pairwise codebook with joint feature space clustering. In *ECCV*, page 14, Crete, Greece, September 2010.
- S. Savarese, J. Winn, and A. Criminisi. Discriminative object class models of appearance and shape by correlatons. *CVPR*, 2:2033–2040, 2006.
- Manik Varma and Andrew Zisserman. Texture classification: Are filter banks necessary. In CVPR, page 2003.
- Manik Varma and Andrew Zisserman. Classifying images of materials: Achieving viewpoint and illumination independence. In *ECCV*, pages 255–271, London, UK, 2002. Springer-Verlag.
- 15. A. Vedaldi and B. Fulkerson. VLFeat: An open and portable library of computer vision algorithms. http://www.vlfeat.org/, 2008.
- Paul Viola and Michael Jones. Rapid object detection using a boosted cascade of simple features. CVPR, 1:511, 2001.
- J. Winn, A. Criminisi, and T. Minka. Object categorization by learned universal visual dictionary. In *ICCV*, ICCV '05, pages 1800–1807, Washington, DC, USA, 2005. IEEE Computer Society.
- Yimeng Zhang and Tsuhan Chen. Efficient kernels for identifying unbounded-order spatial features. CVPR, IEEE Computer Society Conference on, 0:1762–1769, 2009.