Object Classification Using Heterogeneous Co-occurrence Features

Satoshi Ito and Susumu Kubota

Corporate Research & Development Center, Toshiba Corporation, Japan
satoshi13.ito@toshiba.co.jp

Abstract. Co-occurrence features are effective for object classification because observing co-occurrence of two events is far more informative than observing occurrence of each event separately. For example, a color co-occurrence histogram captures co-occurrence of pairs of colors at a given distance while a color histogram just expresses frequency of each color. As one of such co-occurrence features, CoHOG (co-occurrence histograms of oriented gradients) has been proposed and a method using CoHOG with a linear classifier has shown a comparable performance with state-of-the-art pedestrian detection methods. According to recent studies, it has been suggested that combining heterogeneous features such as texture, shape, and color is useful for object classification. Therefore, we introduce three heterogeneous features based on co-occurrence called color-CoHOG, CoHED, and CoHD, respectively. Each heterogeneous features are evaluated on the INRIA person dataset and the Oxford 17/102 category flower datasets. The experimental results show that color-CoHOG is effective for the INRIA person dataset and CoHED is effective for the Oxford flower datasets. By combining above heterogeneous features, the proposed method achieves comparable classification performance to state-of-the-art methods on the above datasets. The results suggest that the proposed method using heterogeneous features can be used as an off-the-shelf method for various object classification tasks.

1 Introduction

Object classification is one of the essential tasks in computer vision and histogram based features such as SIFT (scale invariant feature transform) [9], HOG (histograms of oriented gradients) [1], and a color histogram [17] are widely used features for object classification. A merit of histogram based features is robustness to the slight shift of an object position. However, these histogram based features have the limited discriminative power because they don't take any spatial information into account. One of the solutions to this problem is to extract features from multiple small regions in an image. However, if the regions are too small, features extracted from them become sensitive to the slight object translation. Another solution is to use co-occurrences of pairs of features extracted from different positions in an input image. For example, a color co-occurrence histogram (CCH) [6], also called color correlogram, captures co-occurrence of

K. Daniilidis, P. Maragos, N. Paragios (Eds.): ECCV 2010, Part II, LNCS 6312, pp. 209–222, 2010.

pairs of colors while a color histogram just expresses frequency of each color. In a similar way, edge co-occurrence matrices (ECMs) [13], originally applied to texture classification problem, express the spatial relationship of pairs of edge orientations. Recently, CoHOG (co-occurrence histograms of oriented gradients) [19], an extension of HOG to represent the spatial relationship between gradient orientations, has been proposed and its effectiveness for pedestrian detection and cat face detection has been shown in [19,8]. Methods using co-occurrences of more than two features have also been proposed in [20,15].

According to recent studies [4,16,10,11], it has been suggested that combining heterogeneous features such as texture, shape, and color is useful for object classification. Since heterogeneous features represent various aspects of objects and work complementarily, they achieve higher classification performance than homogeneous features and are applicable to a variety of object classification tasks. Therefore, we introduce three heterogeneous features based on co-occurrence called color-CoHOG, CoHED, and CoHD, respectively.

The remainder of the paper is organized as follows. CoHOG is briefly explained in Sect. 2. Then three heterogeneous features, color-CoHOG, CoHED, and CoHD are proposed in Sect. 3. Experiments are presented in Sects. 4 and 5. Finally, conclusions are given in Sect. 6.

2 Co-occurrence Histograms of Oriented Gradients

CoHOG (co-occurrence histograms of oriented gradients) [19], an extension of HOG [1], consists of multiple co-occurrence histograms of gradient orientations. Though the dimensionality of CoHOG is high, a linear classifier gives high classification performance. Therefore, computational cost of classification is lower than other complex classification methods such as kernel SVM. An algorithm of CoHOG calculation is shown in Algorithm 1. The number of elements of the co-occurrence histograms H is $m \times n \times d^2$ where d is the number of gradient orientation bins. For example, given 10 offsets, 10 small regions, and 10 bins for gradient orientation, the number of elements of H is $10,000$. In detail, please refer to [19].

3 Proposed Features

In this section, we propose three heterogeneous features based on co-occurrence called color-CoHOG, CoHED, and CoHD, respectively. Color-CoHOG, which is an extension of CoHOG to make use of color information, is co-occurrence of a color matching result and a pair of edge directions. CoHED is co-occurrence between edge orientation and color difference. CoHD is co-occurrence of a pair of color differences. Hence color-CoHOG and CoHED are co-occurrences of heterogeneous features and CoHD is co-occurrence of homogeneous features. Details are described in the following sections. We also explain a color histogram as a complementary feature of the above three features.

Algorithm 1. CoHOG calculation

Input: I: a grayscale image, $\{(p_i, q_i)\}_{i=1}^{m}$: m offsets, $\{D_i\}_{i=1}^{n}$: n small regions in the image

compute a gradient orientation image G from I
initialize co-occurrence histograms H with zeros
for $i = 1$ to m **do**
  for $j = 1$ to n **do**
    for all $(x, y) \in D_j$ **do**
      if $(x + p_i, y + q_i)$ is inside of the image **then**
        $g_1 \leftarrow G(x, y)$
        $g_2 \leftarrow G(x + p_i, y + q_i)$
        $H(g_1, g_2, j, i) \leftarrow H(g_1, g_2, j, i) + 1$
      end if
    end for
  end for
end for
return H

3.1 Color-CoHOG

CoHOG calculation described in Algorithm 1 assumes that an input image is grayscale. Derivative masks such as Sobel filter are used to compute gradients. If a color image is given, the conversion from color to grayscale is necessary before CoHOG extraction. Therefore, we extend CoHOG to make use of color information and we apply two ideas. First, we calculate edge orientation in a color image instead of a grayscale one. Second, we use a result of color matching in order to take distinction of foreground and background into account. The details of the ideas are described below.

Deciding edge orientation in a color image is not a trivial problem and a lot of researches have been done [7,14,12]. We found that a method based on the double angle representation [5] gives the consistent results with reasonable computational cost. In the double angle representation, the directions θ and $\theta + 180$ degrees are equivalent and the orthogonal directions θ and $\theta + 90$ degrees are the vectors that point in opposite directions so that averaging gradients in different color channels makes sense (shown in Fig. 1). As a result, we obtain gradient orientations between 0 and 180 degrees since we make no distinction between θ and $\theta + 180$ degrees. In the experiments described in Sects. 4 and 5, Roberts filter is used to calculate initial gradients and then they are averaged in the double angle representation over the RGB channels and the spatial regions of 2×2 pixel size. Averaged gradient orientation is evenly divided into 4 bins.

Foreground-background discrimination is helpful to describe a shape (e.g., [12]). Taking this into account, we use a result of color matching between a pair of pixels at a given offset. This is based on the assumption that colors of a pair of pixels belonging to the same object are likely to be similar while colors of a pair of pixels located at different objects are likely to be dissimilar. In particular, we calculate two co-occurrence histograms per offset and small region, one is the

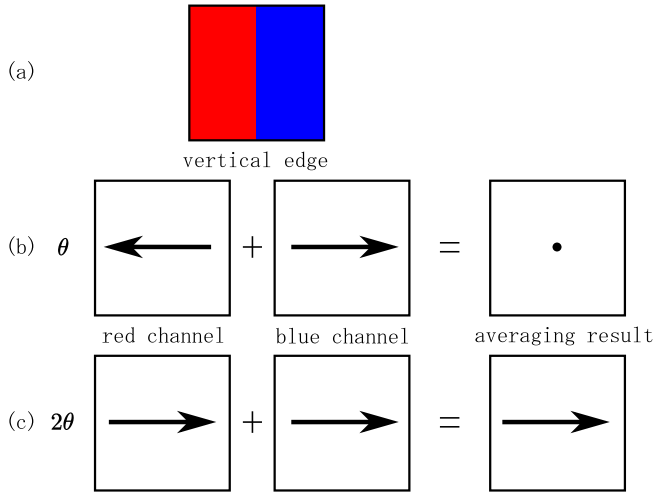

Fig. 1. (a) A vertical edge. (b) Averaging gradients, denoted by arrows, over red and blue channels in the single angle representation gives undesirable result. (c) Averaging gradients over red and blue channels in the double angle representation gives desirable result.

co-occurrence histogram of a pair of pixels at a given offset that have the *same* color and the other is the one of a pair of pixels that have *different* colors. For the computational efficiency, we quantize colors in Cb-Cr space into 17 clusters shown in Fig. 2a and compare the cluster labels to decide if a pair of pixels has the same color.

Our proposed feature named color-CoHOG is summarized in Algorithm 2. Whereas the original CoHOG captures texture information only, color-CoHOG can capture both texture and shape information since foreground-background discrimination is taken into account. The dimension of color-CoHOG is $m \times n \times 2 \times d^2$ where d is the number of quantized edge directions. In the experiments, since we use 16 offsets shown in Fig. 2b, color-CoHOG has $16 \times 1 \times 2 \times 4^2 = 512$ elements per small region.

3.2 CoHED

We propose a feature CoHED (Co-occurrence Histograms of pairs of Edge orientations and color Differences) that expresses the relationships between an edge orientation and the change of colors across the edge. Once an edge orientation at the point p_0 is determined, two points p_1 and p_2 are located at the two opposite sides of the edge point p_0 (shown in Fig. 3a). Edge orientations are calculated in the same manner as described in Sect. 3.1 and color differences between p_1 and p_2 are calculated in YCbCr color space. Then color differences are quantized to 8 directions in each color plane, that is, Y-Cb plane, Y-Cr plane, and Cb-Cr plane. Calculation of a co-occurrence histogram with color difference in Y-Cb plane is as follows;

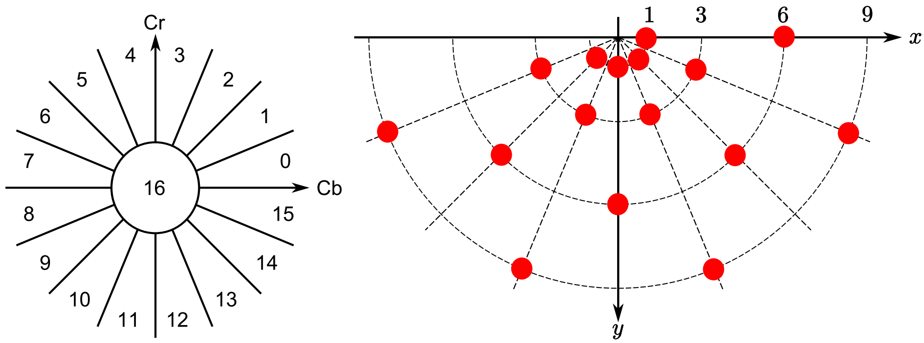

(a) (b)

Fig. 2. (a) The figure shows color labels in Cb-Cr space. The label 16 corresponds to neutral gray. (b) The figure shows 16 offsets (drawn as *filled circles*) used for color-CoHOG calculation.

Algorithm 2. Color-CoHOG calculation

Input: I: a color image, $\{(p_i, q_i)\}_{i=1}^{m}$: m offsets, $\{D_i\}_{i=1}^{n}$: n small regions in the image

compute an edge direction image G from I using the double angle representation
compute color labels C of pixels in the image
initialize co-occurrence histograms H with zeros
for $i = 1$ to m **do**
    for $j = 1$ to n **do**
        for all $(x, y) \in D_j$ **do**
            if $(x + p_i, y + q_i)$ is inside of the image **then**
                $g_1 \leftarrow G(x, y)$
                $g_2 \leftarrow G(x + p_i, y + q_i)$
                if $C(x, y)$ is equal to $C(x + p_i, y + q_i)$ **then**
                    $c \leftarrow 1$
                else
                    $c \leftarrow 0$
                end if
                $H(g_1, g_2, c, j, i) \leftarrow H(g_1, g_2, c, j, i) + 1$
            end if
        end for
    end for
end for
return H

$$H_{\mathrm{Y-Cb}}(g, c) \leftarrow H_{\mathrm{Y-Cb}}(g, c) + |dy| + |du| \tag{1}$$

where g is the edge orientation at p_0, c is the quantized color difference between p_1 and p_2 in Y-Cb plane, and dy and du are differences between p_1 and p_2 in Y channel and Cb channel, respectively. CoHED is computed by weighted voting ($|dy| + |du|$ in (1) corresponds to a voting weight) while other co-occurrence

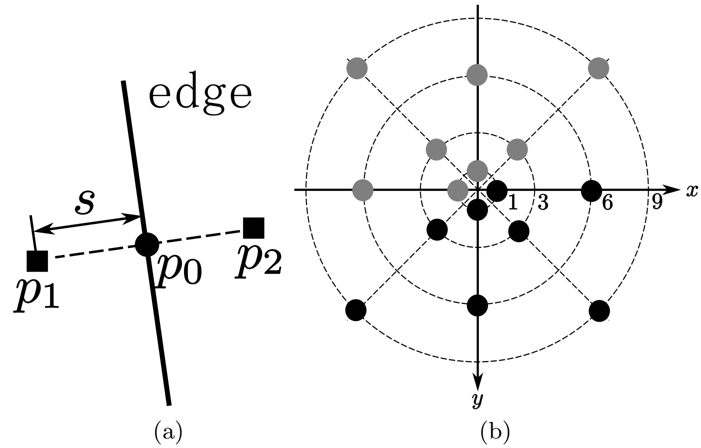

Fig. 3. (a) Positions of three points p_0, p_1 and p_2 used for CoHED. Once edge orientation at p_0 is decided, p_1 and p_2 are located at the two opposite sides of the edge. (b) Eight offsets used for CoHD. A set of three pixels that consist of an origin and a pair of points located at two opposite positions with respect to the origin is used to calculate color difference.

features described in this paper are computed by unweighted voting. Since voting weights for strong step-edges are larger than those for weak ones, CoHED mainly captures shape information rather than texture information. We use 1, 3, 6, and 9 as the distance s from p_0 to $p_1(p_2)$ in the experiments. Thus, the dimension of CoHED is 4 (edge directions) $\times 8$ (directions of color differences) $\times 3$ (color planes) $\times 4$ (scales) $= 384$.

3.3 CoHD

Since color-CoHOG captures shape and texture information and CoHED captures shape information, it's expected that features mainly capturing texture information work complementarily to color-CoHOG and CoHED. Therefore, based on the similar idea as CoHOG, we propose a feature CoHD (Co-occurrence Histograms of color Differences) that simply captures texture information. CoHD represents changes of color values of three pixels located on a given line in an image (shown in Fig. 3b). Color differences are calculated between the centered pixel and the one of the other two pixels, respectively. Calculation of CoHD is described in Algorithm 3. Color differences in Cb-Cr plane are quantized to 4 directions. Eight offsets (shown in Fig. 3b) are used to calculate color differences of pairs of pixels. Thus, the dimension of CoHD is 4 (directions of color differences) $\times 4$ (directions of color differences) $\times 8$ (offsets) $= 128$.

3.4 Color Histogram

The above three features use relative color information. However, absolute color information is also useful for object classification [10,16]. In this paper, we use a

Algorithm 3. CoHD calculation

Input: U: Cb-plane image, V: Cr-plane image, $\{(p_i, q_i)\}_{i=1}^{m}$: m offsets, $\{D_i\}_{i=1}^{n}$: n small regions in the image

initialize co-occurrence histograms H with zeros
for $i = 1$ to m **do**
  for $j = 1$ to n **do**
    for all $(x, y) \in D_j$ **do**
      if $(x + p_i, y + q_i)$ and $(x - p_i, y - q_i)$ are inside of the image **then**
        $u_1 \leftarrow U(x + p_i, y + q_i) - U(x, y)$
        $v_1 \leftarrow V(x + p_i, y + q_i) - V(x, y)$
        $u_2 \leftarrow U(x - p_i, y - q_i) - U(x, y)$
        $v_2 \leftarrow V(x - p_i, y - q_i) - V(x, y)$
        $c_1 \leftarrow (u_1 > 0) + 2 \times (v_1 > 0)$ // quantization into 4 directions
        $c_2 \leftarrow (u_2 > 0) + 2 \times (v_2 > 0)$ // quantization into 4 directions
        $H(c_1, c_2, j, i) \leftarrow H(c_1, c_2, j, i) + 1$
      end if
    end for
  end for
end for
return H

simple color histogram that consists of 17 bins shown in Fig. 2a. Since we use a linear classifier in the experiments, 2nd order polynomial terms of elements of a color histogram are explicitly generated in order to increase linear separability. Thus the number of elements including the 2nd order terms is 170.

4 Experiment 1. INRIA Person Dataset

In this section, we evaluate the proposed method on the INRIA person dataset [1]. The INRIA person dataset provides positive images cropped 64×128 pixels and negative images of various sizes. Some examples are shown in Fig. 4. The number of positive/negative images are $2,416/1,218$ for training and $1,132/453$ for testing, respectively. Detection performance is evaluated by the same way as described in [1]. We extract features separately from 4×8 non-overlapped small regions that are 16×16 pixel sizes and concatenate them into a single feature vector. Since the dimensionality of the feature vectors is high, we use a linear classifier trained by LIBLINEAR [3] that is applicable to a large scale problem. Each component of features is normalized by its maximum value in the training samples, respectively.

4.1 Feature Evaluation

In this section, we study the effect of the following three parameters; the threshold for neutral gray, the number of color bins, and the scale of the offsets. The former two parameters are related to color-CoHOG and the last parameter is related to color-CoHOG, CoHED and CoHD, respectively. Since the dimensionality

Fig. 4. Examples in the INRIA person dataset

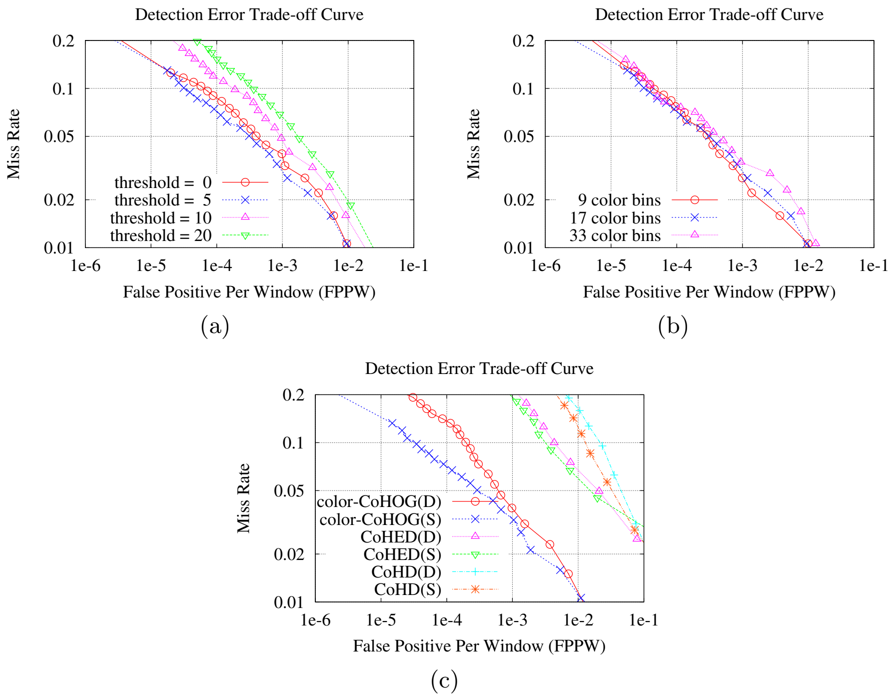

Fig. 5. Feature evaluation on the INRIA person dataset. (a) DET curves of various thresholds for neutral gray. (b) DET curves obtained by changing the number of color bins. (c) DET curves of the sparse setting (denoted by 'S') and the dense setting (denoted by 'D'), respectively.

of features isn't affected by changing the above three parameters, detection performances obtained by changing those parameters can be easily compared with each other. On the other hand, the dimensionality of features is proportional to the square of the number of quantized directions, which is another parameter of the proposed features. In this case, it's difficult to compare the detection performances. Thus we select a practical value for the number of quantized directions and it's used in the experiments in this paper.

The threshold for neutral gray is the parameter that decides whether each pixel is chromatic (labels 0-15 in Fig. 2a) or achromatic (label 16 in Fig. 2a) based on the distance from the origin in Cb-Cr space. Figure 5a shows the DET (detection error trade-off) curves obtained by changing the threshold. The experimental result suggests that a small threshold that classifies most of the pixels as chromatic works well. The setting that classifies all the pixels as chromatic also works as well (threshold = 0 in Fig. 5a). We set the threshold to 5 in other experiments described in this paper.

We also studied the effect of the number of color bins. The experimental result shows that the result of 33 color bins is slightly worse than the other two results but the number of color bins is insensitive to the detection performance (shown in Fig. 5b). We use 17 color bins in other experiments described in the paper.

The scale of the offsets is the parameter that decides the distance between the center pixel and the pixel with an offset. We tested two cases; one is a sparse setting and the other is a dense setting. The sparse setting uses four scales 1, 3, 6 and 9 as the distances between pixels in Figs. 2b and 3 while the dense setting uses 1, 2, 3 and 4. The results of color-CoHOG and CoHD show that the sparse setting is better than the dense one and the result of CoHED shows that the sparse setting is a little bit better than the dense one (shown in Fig. 5c). This suggests that capturing less redundant information is more important to improve classification performance. Therefore, the sparse setting is used in other experiments in the paper.

4.2 Comparison with CoHOG

Figure 6 shows the DET curves of CoHOG and color-CoHOG, respectively. We also plotted the result of 3ch-CoHOG as another extension of CoHOG to make use of color information. 3ch-CoHOG is a feature obtained by concatenating CoHOGs extracted separately from each color channel. The offsets that are used for color-CoHOG (shown in Fig. 2b) are used to calculate CoHOG and 3ch-CoHOG for comparison under the same condition. The detection performance of color-CoHOG is superior to that of CoHOG and comparable to that of 3ch-CoHOG while the dimensionality of color-CoHOG $(16,384)$ is half as that of CoHOG $(32,768)$ and only one-sixth of that of 3ch-CoHOG $(98,304)$, respectively. This result means that color-CoHOG makes use of color information efficiently.

4.3 Comparison with Previous Methods

Figure 7 compares the DET curves of the proposed method with those of the previous methods [1,18,21,2,19,16]. Four heterogeneous features, color-CoHOG, CoHED, CoHD, and color histograms, were used for the proposed method. The curves of the previous methods were obtained by tracing the results in the references. The proposed method achieves 8.6%, 5.5% and 2.9% miss rates at 10^{-6}, 10^{-5} and 10^{-4} FPPWs (false positive per window), respectively. This result is comparable to the state-of-the-art method [16] that has achieved 7.9% miss rate at 10^{-6} FPPW and 5.8% miss rate at 10^{-5} FPPW.

We also show the DET curves of single features in Fig. 8. The result of each single feature except color-CoHOG is far inferior to the method of Dalal et al. [1] (shown in Fig. 7) while the method using the concatenated features achieves comparable

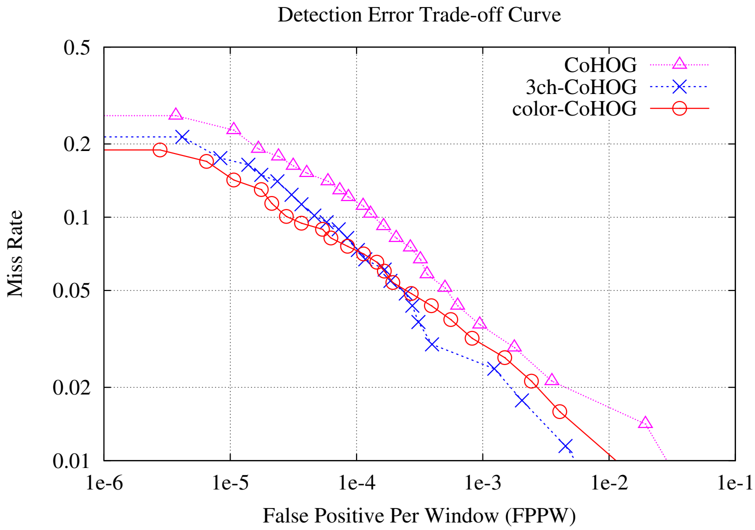

Fig. 6. DET curves of several CoHOGs on the INRIA person dataset. Color-CoHOG (*circle*) is superior to CoHOG (*triangle*) and comparable to 3ch-CoHOG (*cross*) while the dimensionality of color-CoHOG is half as that of CoHOG and only one-sixth of that of 3ch-CoHOG.

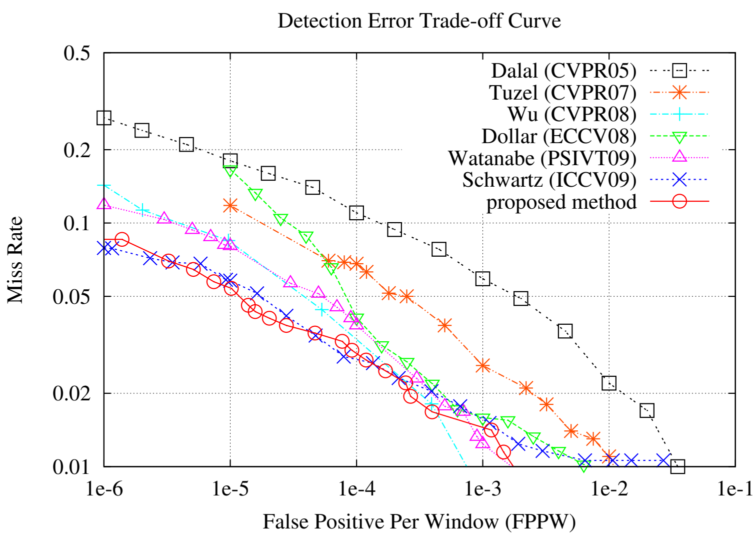

Fig. 7. DET curves of the proposed method and several previous methods on the INRIA person dataset. This figure shows that the proposed method (*circle*) is comparable to the state-of-the-art method (*cross*).

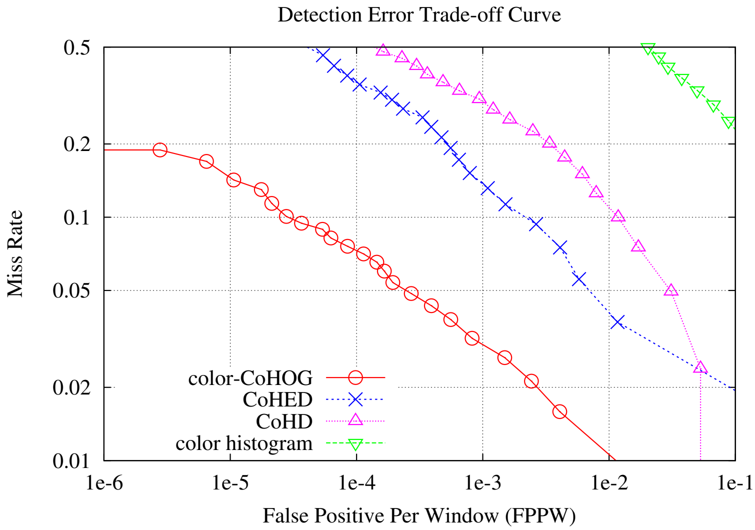

Fig. 8. DET curves of single features on the INRIA person dataset

performance to the state-of-the-art method as described above. This means that our proposed features provide complementary information to each other.

5 Experiment 2. Oxford 17/102 Category Flower Datasets

In this section, we evaluate the proposed method on the Oxford 17/102 category flower datasets [10,11]. The 17 category dataset consists of 80 images per category and the 102 category dataset consists of between 40 and 258 images per category. Some examples are shown in Fig. 9. Classification performance is evaluated by the same manner as described in [10].

In [10], they provide training images, validation images and test images though we don't use validation images since they are not necessary for the proposed method. There are various sizes of images in the datasets, so we crop and resize them into 64×64 pixel size. We extract color-CoHOG, CoHED, CoHD, and a color histogram from the whole region of the resized image and concatenate them into a single feature vector. The dimension of the resulting feature vector is 1,194. In the same manner as described in Sect. 4, linear classifiers trained by LIBLINEAR are used and each component of features is normalized by its maximum value.

Experimental results are shown in Table 1. The proposed method using all features described in Sect. 3 achieves higher classification performance than the state-of-the-art method [11] on both datasets in spite of the simplicity of the proposed method. CoHED achieves the best classification performance among

Fig. 9. Examples of the Oxford 17 category flower dataset

Table 1. Classification performance on the Oxford flower datasets

Method	Performance score [10]	
	17 categories	102 categories
Nilsback [11]	88.33±0.30	72.8
color-CoHOG+CoHED+CoHD+color histogram	94.19±1.22	74.8
color-CoHOG	78.89±1.19	43.4
CoHED	91.54±0.99	64.2
CoHD	84.24±1.07	57.0
color histogram	69.88±2.68	35.6

the four single features on both flower datasets while color-CoHOG achieves the best performance on the INRIA person dataset. This means that effective features are different with respect to object classification tasks. Therefore, a method using homogeneous features, which is effective for a specific object classification task, may fail to achieve high classification performance for another object classification task. In contrast, the proposed method using heterogeneous features can be used as an off-the-shelf method for various object classification tasks.

6 Conclusion

In this paper, we proposed three heterogeneous features based on co-occurrence called color-CoHOG, CoHED, and CoHD, respectively and introduced a color histogram as a complementary feature of those three features. Co-occurrence features are very high dimensional features and highly discriminative, so that a linear classifier is sufficient to achieve high classification performance. Classification performance of each feature was evaluated on the INRIA person dataset and the Oxford 17/102 category flower datasets, respectively. The experimental results show that effective features for the INRIA person dataset are different from those for the Oxford flower datasets. By combining the above four heterogeneous features, the proposed method achieved comparable performance to state-of-the-art methods on the above datasets. The results suggest that the proposed method using heterogeneous features can be used as an off-the-shelf method for various object classification tasks.

References

1. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 1, pp. 886–893 (2005)
2. Dollár, P., Babenko, B., Belongie, S., Perona, P., Tu, Z.: Multiple component learning for object detection. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part II. LNCS, vol. 5303, pp. 211–224. Springer, Heidelberg (2008)
3. Fan, R.E., Chang, K.W., Hsieh, C.J., Wang, X.R., Lin, C.J.: LIBLINEAR: A library for large linear classification. Machine Learning Research 9, 1871–1874 (2008)
4. Gehler, P.V., Nowozin, S.: On feature combination for multiclass object classification. In: Proceedings of the Twelfth IEEE International Conference on Computer Vision (2009)
5. Granlund, G.H.: In search of a general picture processing operator. In: Computer Graphics and Image Processing, pp. 155–173 (1978)
6. Huang, J., Kumar, S.R., Mitra, M., Zhu, W.J., Zabih, R.: Image indexing using color correlograms. In: Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition, p. 762. IEEE Computer Society, Los Alamitos (1997)
7. Koschan, A.: A comparative study on color edge detection. In: Li, S., Teoh, E.-K., Mital, D., Wang, H. (eds.) ACCV 1995. LNCS, vol. 1035, pp. 574–578. Springer, Heidelberg (1995)
8. Kozakaya, T., Ito, S., Kubota, S., Yamaguchi, O.: Cat face detection with two heterogeneous features. In: Proceedings of the 2009 IEEE International Conference on Image Processing (2009)
9. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. Int. Journal of Computer Vision 60(2), 91–110 (2004)
10. Nilsback, M.-E., Zisserman, A.: A visual vocabulary for flower classification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, vol. 2, pp. 1447–1454 (2006)
11. Nilsback, M.E., Zisserman, A.: Automated flower classification over a large number of classes. In: Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing (December 2008)
12. Ott, P., Everingham, M.: Implicit color segmentation features for pedestrian and object detection. In: Proceedings of the Twelfth IEEE International Conference on Computer Vision (2009)
13. Rautkorpi, R., Iivarinen, J.: A novel shape feature for image classification and retrieval. In: Proc. of Int. Conf. on Image Analysis and Recognition, Part I, pp. 753–760 (2004)
14. Ruzon, M.A., Tomasi, C.: Color edge detection with the compass operator. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, vol. 2, pp. 160–166 (1999)
15. Sabzmeydani, P., Mori, G.: Detecting pedestrians by learning shapelet features. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 1–8 (2007)

16. Schwartz, W.R., Kembhavi, A., Harwood, D., Davis, L.S.: Human detection using partial least squares analysis. In: Proceedings of the Twelfth IEEE International Conference on Computer Vision (2009)
17. Swain, M.J., Ballard, D.H.: Color indexing. Int. Journal of Computer Vision 7(1), 11–32 (1991)
18. Tuzel, O., Porikli, F., Meer, P.: Human detection via classification on riemannian manifolds. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2007)
19. Watanabe, T., Ito, S., Yokoi, K.: Co-occurrence histograms of oriented gradients for pedestrian detection. In: The 3rd Pacific Rim Symposium on Advances in Image and Video Technology, pp. 37–47 (2009)
20. Wu, B., Nevatia, R.: Detection of multiple, partially occluded humans in a single image by bayesian combination of edgelet part detectors. In: The Tenth IEEE International Conference on Computer Vision, vol. 1, pp. 90–97. IEEE Computer Society, Washington (2005)
21. Wu, B., Nevatia, R.: Optimizing discrimination-efficiency tradeoff in integrating heterogeneous local features for object detection. In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (2008)