Foreground Object Detection by Motion-based Grouping of Object Parts

Zhipeng Wang · Jinshi Cui · Hongbin Zha ·
Masataka Kagesawa · Shintaro Ono · Katsushi Ikeuchi

Received: 2 August 2012 / Revised: 24 September 2013 / Accepted: 4 December 2013 / Published online: 30 January 2014
© The Author(s) 2014. This article is published with open access at Springerlink.com

Abstract Effective video-based detection methods are of great importance to intelligent transportation systems (ITS), and here we propose a method to localize and label objects. The method is able to detect pedestrians and bicycle riders in a complex scene. Our method is inspired by the common fate principle, which is a mechanism of visual perception in human beings, and which states tokens moving or functioning in a similar manner tend to be perceived as one unit. Our method embeds the principle in an Implicit Shape Model (ISM). In our method, keypoint-based object parts are firstly detected and then grouped by their motion patterns. Based on the grouping results, when the object parts vote for object centers and labels, each vote belonging to the same object part is assigned a weight according to its consistency with the votes of other object parts in the same motion group. Afterwards, the peaks, which correspond to

detection hypotheses on the Hough image formed by summing up all weighted votes, become easier to find. Thus our method performs better in both position and label estimations. Experiments show the effectiveness of our method in terms of detection accuracy.

Keywords Object detection · Motion grouping · Common fate

1 Introduction

In ITS areas, detection methods using cameras can be used for navigation, safe driving, surveillance, and sustaining results from other sensors. In traditional ITS applications, vehicles are the main targets. Currently pedestrians are also considered as important subjects of ITS applications, and bicycles also are becoming very popular for environmental and economical reasons. In Japan, the number of traffic accidents among bicycles and pedestrians is very large. Thus we tackle an issue of detecting freely moving bicycle riders and pedestrians from the data collected by a camera which keeps them under surveillance from the top. These situations can be observed in parks, university campuses, station squares, tourist spots, etc. Here we focus on techniques from the area of computer vision for detection under surveillance scenarios. It is also assumed by the method that target objects captured by the surveillance cameras do not change much in scale.

Most state-of-the-art visual detection methods fall into two main categories: sliding-window methods and Hough transform based methods. The methods [10, 27] based on a sliding window schema perform detection in a typical machinery way. In these methods, decisions of whether a target object exists or not are made for part of or all of the

---

Z. Wang (✉) · M. Kagesawa · S. Ono · K. Ikeuchi
Institute of Industrial Science, The University of Tokyo,
Tokyo, Japan
e-mail: wangzp@cvt.iis.u-tokyo.ac.jp

M. Kagesawa
e-mail: kagesawa@cvt.iis.u-tokyo.ac.jp

S. Ono
e-mail: onoshin@cvt.iis.u-tokyo.ac.jp

K. Ikeuchi
e-mail: ki@cvt.iis.u-tokyo.ac.jp

J. Cui · H. Zha
Key Lab of Machine Perception (Ministry of Education),
Peking University, Peking, China

J. Cui
e-mail: cjs@cis.pku.edu.cn

H. Zha
e-mail: zha@cis.pku.edu.cn

sub-images in a test image. Besides the attractive performance and the extendibility of combining various kernels, these methods are favorable because they consider each object as a whole during detection. However, they share limited aspects with visual perception in human beings, and their efficiency heavily relies on the size of the test images.

The other methods [5, 6, 13, 18] detect objects based on the generalized Hough transform [1]. Object parts are detected, and the object parts provide confidence of the locations being the potential objects' centers. Locations of objects are decided according to the converged confidence. These methods are favorable for their robustness to partial deformation and ease of training. To human beings, this kind of method seems to be more natural. And in our work, we combine a mechanism of visual perception in humans, with the ISM [13] to demonstrate this natural property.

A typical Hough transform based method contains two steps: training and detection. During training, a codebook of object parts is built from a set of well annotated images. Each code in the codebook contains information about the appearance of the object part, the relative position to the object center, and the class label. Each object part's appearance is given in the form of keypoint descriptors [13], image patches [7, 19], or image regions [8]. Each code not only encodes one object part's appearance, but also its offset to the object center and the class label. During the detection step, object parts are detected on each test image. Then every object part is matched against the codebook, and several codes nearest in appearance are activated. The offset and class label encoded in each activated code will act as a vote. All the votes from the object parts are added up to form a Hough image. The peaks of the Hough image are considered detection hypotheses with the height of each peak as the confidence for the corresponding hypothesis.

Two challenging issues for detection methods are how to separate near objects and how to separate similar different-class objects. The target objects, in the case of ITS applications, are pedestrians, bicycle riders, and automobiles. In the schema of sliding window, usually non-maximum suppression is needed for post-processing, and a mechanism in [10] works by excluding from the feature pool the features which belong to each successive detection response. In Hough transform based methods, a similar mechanism is also employed in [2], however, this effort is employed after the forming of a Hough image. During the forming of a Hough image, two kinds of votes make detection challenging: (1) votes cast by object parts from near objects make the peaks corresponding to different objects mixed up, and (2) votes cast by similar different-class object parts lead to tough decisions on the class label of the peaks. See Fig. 2d. Before the forming of Hough images, problems also arise from the pollution of the training images' background part to the codebook. During training a very clean codebook can

be built with the foreground marked, which requires manual efforts. Otherwise, a large amount of training examples are needed for the effectiveness of the codebook, and this decreases efficiency.

In videos, motion information is also available by simple tracking of object parts. Thus we propose a method for detection which utilizes both appearance and motion information. The method is based on the common fate principle [23]. The principle is one of the visual perception principles as theorized by gestalt psychologists, and it states that for human beings, tokens moving coherently are perceptually grouped. This provides an intuition to group the object parts by their motion patterns, and let them vote afterwards. In our work, the object parts are represented using keypoint descriptors, which are tracked to generate trajectories. The object parts are grouped by the pairwise similarities of their corresponding trajectories. Using the assumption that object parts in the same motion group probably belong to the same object, for each object part, we assign higher weights for the votes of the object parts which are more "agreeable" within the motion group. This results in votes corresponding to true detection responses to be more likely assigned higher weights. And on a Hough image formed by summing up these weighted votes, the peaks are easier to find as shown in Fig. 1d.

Due to the combination of motion analysis results and the Hough transform framework, and by assigning different weights to each object part's votes, the proposed method has several appealing properties:

- The method's ability to estimate object position and label multiple objects from different classes. The existence of three types of objects makes the task challenging: near objects, similar different-class objects, and multi-pose same-class objects.
- Its ability to use a codebook trained by images with cluttered backgrounds.
- The framework used to combine grouping results of object parts is very general, and thus can be easily expanded.

The remaining paper is organized as follows. Section 2 reviews related work. Section 3 formalizes the common fate Hough transform. Section 4 describes inference on the formed Hough images. Section 5 gives experimental results, and Section 6 concludes.

2 Related Work

Our work is most closely related to object detection methods [2, 12–16] based on the Hough transform framework. Recently, such methods are making a lot of progress. The ISM [13, 14] is extended by notifying correspondences

Fig. 1 Merit of the proposed method. **a** Original image. **b** Motion grouping results. Some parts are enlarged to show details. **c** Original Hough image. **d** Hough image formed using our method. The grids in **c** and **d** correspond to the grids in **a**

between the object parts and the hypotheses [2] for the detection of multiple near objects. While in the methods [7, 15, 19], the Hough transform is placed in a discriminative framework for object detection in a way that the codes are assigned different weights by the co-occurrence frequency of their appearance and offset to the object center. Two Hough transform methods consider the grouping of object parts [20, 26]. The method in [20] deals with scale change. Instead of estimating the scale by local features trained from different scaled examples, the votes are considered as voting lines. By considering the difference between the voted centers, local features are first grouped, resulting in a more consistent vote for the object center. In [26], the grouping of object parts, the correspondence between object parts and object, and the decisions on detection hypotheses are optimized in the same energy function. For this method, the problem is that the grouping results don't have meaning or correspond to any real entities.

Our work is also related to object detection methods which use trajectories [3, 4], methods using the weighting of features [25], methods dealing with codebook noise [17], and methods which integrate temporal information [24].

3 Common Fate Hough Transform

Probabilistic standpoints are very appealing because of inference ease. However, as pointed out in [11], placing an Implicit Shape Model (ISM) in a probabilistic framework is not satisfactory. Especially, describing weights of the votes as priors does not make sense. A Hough transform can be simply considered as the transformation from a set of object parts, $\{e\}$, to a confidence space of object hypotheses, $C(x, l)$. Where x is the coordinate of the object center, and l the label. Terms described as priors of the votes in the ISM are actually weights, and the likelihood terms are actually blurring functions to convert discrete votes into continuous space. This section describes how a Hough image for the estimation of object centers and labels is formed from object parts observed on an image.

Let e denote an object part observed on the current image. The appearance of e is matched against the codebook, and e activates N best matched codes from the trained codebook. Each code contains the appearance, its offset to the object center, and the class label. According to the N matched codes, e casts N votes. Each vote V_e is about the object center that generates e . The position of the object center casted by a vote, V , is denoted by x_V , while the class label is l_V .

Based on the N votes of $\mathbf{e}$, the confidence that a position $\tilde{\mathbf{x}}$ is the center of an object with class label $\tilde{l}$ is given by,

$$C(\tilde{\mathbf{x}}, \tilde{l}; \mathbf{e}) = \sum_{i=1}^N B(\tilde{\mathbf{x}}, \tilde{l}; V_{\mathbf{e}}^i) w(V_{\mathbf{e}}^i). \quad (1)$$

Here $B(\tilde{\mathbf{x}}, \tilde{l}; V_{\mathbf{e}}^i)$ is the blurring function. And $w(V_{\mathbf{e}}^i)$ is the weight of $V_{\mathbf{e}}^i$.

The idea of the proposed method is that the weight term, $w(V_{\mathbf{e}}^i)$, is defined by the motion grouping results of all the object parts.

The blurring function is defined as,

$$B(\tilde{\mathbf{x}}, \tilde{l}; V) = \begin{cases} 0 & \text{if } l_V \neq \tilde{l} \text{ or } |\tilde{\mathbf{x}} - \mathbf{x}_V| > d \\ G(\tilde{\mathbf{x}}; \mathbf{x}_V, \sigma) & \text{otherwise} \end{cases}. \quad (2)$$

Here $G(\tilde{\mathbf{x}}; \mathbf{x}_V, \sigma)$ is a Gaussian function that fixes the spatial gap between $\tilde{\mathbf{x}}$ and $\mathbf{x}_V$.

Let M be the total number of object parts on the image, then by summing up over all the object parts, the confidence of $\tilde{\mathbf{x}}$ being the center of an $\tilde{l}$ -class object is given by,

$$\begin{aligned} C(\tilde{\mathbf{x}}, \tilde{l}) &= \sum_{j=1}^M C(\tilde{\mathbf{x}}, \tilde{l}; \mathbf{e}_j) w(\mathbf{e}_j) \\ &= \sum_{j=1}^M \sum_{i=1}^N B(\tilde{\mathbf{x}}, \tilde{l}; V_{\mathbf{e}_j}^i) w(V_{\mathbf{e}_j}^i) w(\mathbf{e}_j). \end{aligned} \quad (3)$$

A uniform weight is assumed for each object part, and $w(\mathbf{e}_j) = \frac{1}{M}$. By considering $C(\tilde{\mathbf{x}}, \tilde{l})$ as the evaluation score of the Hough space $(\tilde{\mathbf{x}}, \tilde{l})$, the task of estimating object centers and labels converts to finding, and then validating, the local maxima of the Hough image.

3.1 Common Fate Weights

To meet the challenges of separating near objects, separating similar different-class objects, and using a noisy codebook, different weights are assigned to the votes of each object part by considering the motion grouping results of the object parts. In this sub-section, when given some grouping results, how the results are combined into a Hough transform framework is introduced.

Let $\gamma = \{\mathbf{g}\}$ denote the grouping results, where $\mathbf{g}$ is a group of object parts. Assume $\mathbf{e}_m \in \mathbf{g}$ and $\mathbf{e}_n \in \mathbf{g}$. Those votes of $\mathbf{e}_m$ which are more “agreeable” than the votes of the other objects in $\mathbf{g}$ are assigned larger weights.

Towards this end, the relationship between the votes of $\mathbf{e}_m$ and the votes of $\mathbf{e}_n$ needs to be given in advance. This relationship is named support. The support from $V_{\mathbf{e}_n}$ to $V_{\mathbf{e}_m}$ is defined based on $V_{\mathbf{e}_n}$ and the confidence that $V_{\mathbf{e}_m}$ ’s voted center is correct, as,

$$S(V_{\mathbf{e}_n} \rightarrow V_{\mathbf{e}_m}) = B(\mathbf{x}_{V_{\mathbf{e}_m}}, l_{V_{\mathbf{e}_m}}; V_{\mathbf{e}_n}), n \neq m.$$

Here $B(\mathbf{x}_{V_{\mathbf{e}_m}}, l_{V_{\mathbf{e}_m}}; V_{\mathbf{e}_n})$ is defined in Eq. 2. This measures the coherence of the two votes from different object parts.

Then, the support from $\mathbf{e}_n$ to $V_{\mathbf{e}_m}$ is defined based on $\mathbf{e}_n$, and the confidence that $V_{\mathbf{e}_m}$ ’s voted center is correct, as,

$$\begin{aligned} S(\mathbf{e}_n \rightarrow V_{\mathbf{e}_m}) &= C(\mathbf{x}_{V_{\mathbf{e}_m}}, l_{V_{\mathbf{e}_m}}; \mathbf{e}_n) \\ &= \sum_{i=1}^N S(V_{\mathbf{e}_n}^i \rightarrow V_{\mathbf{e}_m}) w(V_{\mathbf{e}_n}^i), n \neq m. \end{aligned}$$

And the support from $\mathbf{g}$ to $V_{\mathbf{e}_m}$ is defined by the confidence that $V_{\mathbf{e}_m}$ ’s voted center is correct based on the votes of all the other object parts excluding its belonging object part in $\mathbf{g}$, as,

$$\begin{aligned} S(\mathbf{g} \rightarrow V_{\mathbf{e}_m}) &= \sum_{\mathbf{e}_i \in \mathbf{g} - \{\mathbf{e}_m\}} C(\mathbf{x}_{V_{\mathbf{e}_i}}, l_{V_{\mathbf{e}_i}}; \mathbf{e}_i) w(\mathbf{e}_i) \\ &= \frac{1}{M} \sum_{\mathbf{e}_i \in \mathbf{g} - \{\mathbf{e}_m\}} S(\mathbf{e}_i \rightarrow V_{\mathbf{e}_m}). \end{aligned}$$

By assuming all object parts in the same motion group are from the same object, which means motion grouping gives good results, the estimations for center position and class label given by every object part should be consistent with that given by the motion group. Thus for a particular vote of $\mathbf{e}_m$, i.e., $\tilde{V}_{\mathbf{e}_m}$, a weight is assigned to it by considering its consistence with $\mathbf{g}$ and the consistence of $\mathbf{e}_m$ ’s other votes with $\mathbf{g}$, as:

$$\begin{aligned} w(\tilde{V}_{\mathbf{e}_m}) &= \frac{S(\mathbf{g} \rightarrow \tilde{V}_{\mathbf{e}_m}) + \frac{\Delta}{N}}{\sum_{i=1}^N S(\mathbf{g} \rightarrow V_{\mathbf{e}_m}^i) + \Delta} \\ &= \frac{\sum_{\mathbf{e}_j \in \mathbf{g} - \{\mathbf{e}_m\}} \sum_{k=1}^N S(V_{\mathbf{e}_j}^k \rightarrow \tilde{V}_{\mathbf{e}_m}) w(V_{\mathbf{e}_j}^k) + \frac{M\Delta}{N}}{\sum_{i=1}^N \sum_{\mathbf{e}_j \in \mathbf{g} - \{\mathbf{e}_m\}} \sum_{k=1}^N S(V_{\mathbf{e}_j}^k \rightarrow V_{\mathbf{e}_m}^i) w(V_{\mathbf{e}_j}^k) + M\Delta}. \end{aligned} \quad (4)$$

Here, Δ is a small constant for preventing zeros. Notice $w(\tilde{V}_{\mathbf{e}_m})$ is defined using $w(V_{\mathbf{e}_j}^k)$ - the weights of the votes of the other object parts in $\mathbf{g}$. In order to determine $w(\tilde{V}_{\mathbf{e}_m})$, uniform weights are firstly assigned to the votes of each object part in $\mathbf{g}$, i.e., $w(V_{\mathbf{e}_j}^k) = \frac{1}{N}$. Then new weights are calculated based on the uniformly assigned weights. The weights of votes used to form the Hough image are the iteratively converged weights.

The grouping result $\gamma = \{\mathbf{g}\}$, can be replaced by grouping results based on other information, for example our method utilizes motion to group the voting elements. The manner of extending the Hough transform is very general, and the extended Hough transform with motion grouping results is called the common fate Hough transform. The votes given by the best matched codes and the votes with higher defined weights are shown in Fig. 2.

Fig. 2 Effect of the proposed weight. **a** Motion groups, different colors mark different motion groups. **b** Voted centers given by the 7 best matched codes. **c** Voted centers with the highest defined weights. **d** Voted centers with weights higher than a threshold

3.2 Motion Grouping

In this subsection, how to group the object parts by their motion patterns is introduced. Basically the object parts are tracked, and clustered by their motion patterns. The object parts are tracked through frames before and after the current frame, to generate trajectories. Then the object parts are grouped by their corresponding trajectories' pairwise motion similarities.

The object parts in this method are in the form of keypoint descriptors. The Harris Corner [9] feature is chosen,

for robustness, to represent each object part, while for appearance, the region covariance [22] feature of the image patch around each keypoint is used. The image feature is chosen because of its flexibility to combine multiple channels of information, and also for its capability of handling scale changes in a certain range.

For each object part, a trajectory is generated by tracking its corresponding Harris Corner by the KLT tracker [21]. To group the trajectories, two pairwise similarities are defined.

Fig. 3 **a** Training images. Note some keypoints fall on the background. **b** The manner how a 9×9 image patch is used to generate six region covariances, and red rectangles indicate the pixels used for each covariance

Fig. 4 Motion grouping results

Let T_{e_m} and T_{e_n} denote two trajectories corresponding to e_m and e_n . The first similarity between two trajectories is defined as,

$$D_1(T_{e_m}, T_{e_n}) = \max_{i=1 \dots L} \left(\left| \mathbf{x}_{T_{e_m}}^i - \mathbf{x}_{T_{e_n}}^i \right| \right).$$

Here, i is the frame index, and L is the number of frames in which both trajectories exist.

To define the second similarity, the i th directional vector of T is firstly defined as, $\mathbf{d}_T^i = \mathbf{x}_T^{i+3} - \mathbf{x}_T^i$. Let $\mathbf{a}_i = \mathbf{d}_{T_{e_m}}^i$, $\mathbf{b}_i = \mathbf{d}_{T_{e_n}}^i$, $a_i = \frac{\mathbf{a}_i \cdot \mathbf{b}_i}{\|\mathbf{a}_i\| \|\mathbf{b}_i\|}$, and $b_i = \frac{\mathbf{a}_i \cdot \mathbf{b}_i}{\|\mathbf{a}_i\| \|\mathbf{b}_i\|}$. Then the second similarity is defined as,

$$D_2(T_{e_m}, T_{e_n}) = \max_{i=1 \dots L-3} (\max(|\mathbf{a}_i - a_i \mathbf{a}_i|, |\mathbf{b}_i - b_i \mathbf{b}_i|)).$$

Before grouping the trajectories, the static points are excluded. The defined D_1 is calculated for all pairs of

Fig. 5 **a** Precision-recall curves (red: the proposed method, blue: the benchmark method). **b** Confusion matrices (upper: the proposed method, lower: the benchmark method)

b

pedestrian	487	2	21
bicycle	0	108	15
none	64	72	0
	pedestrian	bicycle	none

pedestrian	437	2	71
bicycle	0	93	30
none	82	56	0
	pedestrian	bicycle	none

Fig. 6 Results. Red rectangles and blue rectangles mark correctly detected pedestrians and bicycle riders. Yellow rectangles mark missed detections. White rectangles mark correctly detected but not correctly

labeled objects. Green rectangles mark false alarms. Black rectangles mark static objects which are beyond the verification of this method

Fig. 7 Effect of the proposed weight assignment. Red circles are voted centers for leopards, while blue ones are voted centers for tigers. On the top are the motion grouping results. In the middle are the voted centers according to the best matched codes. On the bottom are the voted centers voted by votes with the highest weights

trajectories, and a minimum spanning tree is then built using the calculated similarities. The built minimum spanning tree is split by cutting edges larger than a threshold, D_{th}^1 , and this gives a grouping result of the trajectories. For each element in the clustering result, D_2 is used in the same procedure to generate even smaller clusters. This hierarchical procedure ensures that trajectories in the same group have both small D_1 and D_2 . Max operation is used in the definitions of both D_1 and D_2 . This is helpful because very often two trajectories are of different lengths, and under such situations, max operation will have better stability than other operations, e.g., average, that consider only overlapping frames.

Each trajectory corresponds to an object part, and the grouping results of the trajectories correspond to grouping results of the object parts.

3.3 Codebook

For training, Harris corners are extracted from the training images with the object center and the class label annotated. In this method, region covariance is chosen to represent the appearance, which is defined as,

$$\mathbf{r} = \frac{1}{K-1} \sum_{i=1}^K (\mathbf{z}_i - \mu)(\mathbf{z}_i - \mu)^T.$$

Here, K is the number of pixels in the region, and $\mathbf{z}_i$ is a 7-dimensional vector regarding the (x, y) coordinate of the pixel, while μ is the mean of $\mathbf{z}_i$. And $\mathbf{z}(x, y)$ contains the RGB color of the pixel and the intensity gradients of the pixel, as: $r(x, y)$, $g(x, y)$, $b(x, y)$, $\left| \frac{\partial I(x, y)}{\partial x} \right|$, $\left| \frac{\partial I(x, y)}{\partial y} \right|$, $\left| \frac{\partial^2 I(x, y)}{\partial x^2} \right|$, and $\left| \frac{\partial^2 I(x, y)}{\partial y^2} \right|$.

Fig. 8 Example Hough images. On the *top* are the original images. In the *middle* are the Hough images formed by votes with uniform priors. On the *bottom* are the Hough images formed by votes with the proposed weights. *Red* indicates leopards, and *blue* indicates tigers. Note that for the two leopards, there is no peak corresponding to the one on the *right*, on the benchmark Hough image. For the three leopards, there is also no peak corresponding to the leopard behind on the benchmark Hough image

The appearance similarity between $\mathbf{r}_m$ and $\mathbf{r}_n$ is given by,

$$\rho(\mathbf{r}_m, \mathbf{r}_n) = \sqrt{\sum_{i=1}^7 \ln^2 \lambda_i}.$$

Here, λ_i is the generalized eigenvalue obtained by solving the generalized eigenvalue problem, $\lambda_i \mathbf{r}_m \mathbf{u}_i = \mathbf{r}_n \mathbf{u}_i$, $\mathbf{u}_i \neq \mathbf{0}$, with $\mathbf{u}_i$ the eigenvector.

A square image patch around each keypoint is used to represent the appearance of an object part. Six region covariances are generated for each image patch by using the pixels of the top-left, the top-right, the bottom-left, the bottom-right, the central portion, and the entire image patch. Then besides the offset and the class label, a code contains six region covariances. All codes from all training images constitute the codebook. When an object part is matched against the codebook, the similarity between the image patch of the object part and a code is defined by the smallest similarity of the corresponding region covariance. This will handle scale changes of a small range, since the six image patches are not of the same scale. And the method's ability of handling scale changes is limited. So it can only be used in surveillance situations where the scales of target objects change in a limited range.

4 Detection

After forming the Hough image, the detection hypotheses are validated. Let $\mathbf{h} = \{H\}$ be the points in the Hough space

which are evaluated by $C(\mathbf{x}_H, l_H)$ and have $C(\mathbf{x}_H, l_H) > 0$. Inspired by [2], the hypotheses are validated by an optimizing procedure. Let O be the number of the points in $\mathbf{h}$. Let $u_i = 1$ or 0 , indicate H_i as a true object center or not. The problem is:

$$\arg \max_{u_i} \prod_{i=1}^O C^{u_i}(H_i) \iff \arg \max_{u_i} \sum_{i=1}^O u_i \ln(C(H_i)).$$

Let $v_{ij} = 1$ or 0 , indicate $\mathbf{e}_j$ belongs to H_i or not, then

$$\begin{aligned} C(H_i) &= \sum_{j=1}^M C(\mathbf{x}_{H_i}, l_{H_i}; \mathbf{e}_j) w(\mathbf{e}_j) \\ &= \frac{1}{M} \sum_{j=1}^M v_{ij} C(\mathbf{x}_{H_i}, l_{H_i}; \mathbf{e}_j), \end{aligned}$$

and by assuming one object part belongs to and only belongs to one hypothesis, the problem is,

$$\begin{aligned} \arg \max_{u_i, v_{ij}} \sum_{i=1}^O u_i \ln \left(\sum_{j=1}^M v_{ij} C(\mathbf{x}_{H_i}, l_{H_i}; \mathbf{e}_j) \right) \\ \text{s.t. : } u_i = 0 \text{ or } u_i = 1, \forall i; \\ v_{ij} = 0 \text{ or } v_{ij} = 1, \forall i, \forall j; \\ \sum_{i=1}^O v_{ij} = 1, \forall j; \\ \sum_{j=1}^M v_{ij} \leq u_i, \forall i. \end{aligned}$$

Fig. 9 Results. Red crosses mark the centers for leopards and blue crosses mark the centers for tigers

Following [2], the optimal result for the problem is given by greedy maximization. As described in Algorithm 1, the largest local maximum of all the local maxima is chosen to be the center of a true object, and then the object parts belonging to the chosen object center are excluded from the object part set. A new Hough image, where new objects are found, is formed using the remaining object parts. And this procedure ends when the object part set is empty, or when the confidence of the chosen object is lower than a given threshold.

Algorithm 1 Greedy Maximization

Let ε be the set of object parts, C_{th} be the low confidence threshold to accept detection responses, and $\hat{\mathbf{h}}$ be the local maxima of $\mathbf{h}$

```

1: while  $\varepsilon \neq \emptyset$  do
2:   Form  $\mathbf{h}$  with  $\varepsilon$ 
3:   Generate  $\hat{\mathbf{h}}$  and select  $H_i \in \hat{\mathbf{h}}$  with the largest
      $C(\mathbf{x}_{H_i}, l_{H_i})$ 
4:   if  $C(\mathbf{x}_{H_i}, l_{H_i}) \geq C_{th}$  then
5:     for  $\mathbf{e}_j \in \varepsilon$  do
6:       if  $\forall H' \in \hat{\mathbf{h}}, C(\mathbf{x}_{H_i}, l_{H_i} | \mathbf{e}_j) \geq$ 
          $C(\mathbf{x}_{H'}, l_{H'} | \mathbf{e}_j)$  then
7:          $\varepsilon \leftarrow \varepsilon - \{\mathbf{e}_j\}$ 
8:       end if
9:     end for
10:  else
11:     $\varepsilon \leftarrow \emptyset$ 
12:  end if
13: end while
14: return  $\{H_i\}$ 

```

---

5 Experimental Results

In our experiments, improvement of the method is verified in terms of detection accuracy. The method is tested on the P-campus dataset with [2] as benchmark, and then tested on a dataset of several animals.

5.1 Campus-scene Detection

Dataset The P-campus dataset contains two primary classes of foreground objects: pedestrians and bicycle riders. The frame size is 720×576. Among all the 401 continuous frames, 633 different-class ground truth bounding boxes are annotated on 79 frames. In this dataset, pedestrians and bicycle riders have in common the upper human body, and pedestrians appear in front, back, and side views.

Implementation Settings For training, 52 bicycle riders and 171 pedestrians are randomly selected from the marked

ground truths. Harris corners are detected on these randomly selected training images, examples are given in Fig. 3a. For appearance, six region covariances are generated for each keypoint using the 9×9 image patch around it as shown in Fig. 3b. The appearance, the offset to the image (object) center, and the label of the training image are encoded into a code, and the code is inserted into a codebook. The final codebook contains 5502 codes. Testing data is formed by the 79 frames, on which the ground truth bounding boxes are marked. Harris corners are detected, and region covariances are generated in the same manner as for the training images. For each Harris corner on one testing image, the corresponding region covariances are matched against the codebook for the most similar codes. Some of the training examples will appear in the test sequences. The emphasis of this experiment is to verify the proposed framework's ability of combining motion information. Both the proposed method and the benchmark method use the same training and testing images, so the comparison is fair and proves the effectiveness of the proposed method.

For motion grouping, each keypoint is tracked through 10 frames before, and 10 frames after the current frame. The similarity of two 21-point trajectories is defined using only the frames in which both trajectories exist. To set the two thresholds for motion grouping, D_1 and D_2 are measured for keypoint pairs of different objects. D_{th}^1 is set so that it is larger than only 10% of the measured D_1 s, and so is D_{th}^2 . By doing this, keypoints belonging to different objects are not likely to be grouped together. So that in one motion group, the keypoints are very likely to belong to the same object, as shown in Fig. 4.

To form the Hough image, 35 best matched codes are chosen from the codebook for each object part. In Eq. 3, d and σ need to be given. The precision-recall curves are based on σ , while d is set to 10. Here σ is the most important parameter.

Comparisons For comparison, detection is done on the Hough images formed with and without motion grouping results. The same codebook and the same parameter settings are used for forming and searching over both Hough images. The votes of each object part are assigned uniform weights in the benchmark method, while weights defined in Eq. 4 are assigned in the proposed method.

The precision-recall curves are shown in Fig. 5a. An object is considered as correctly detected only if the distance from the ground truth to it is less than 10 pixels. In Fig. 5a, the correctly positioned but wrongly labeled objects are considered as true positives, aiming at verifying the positioning ability of the proposed method.

The confusion matrices are given in Fig. 5b. For clarity, the proposed method is compared with the benchmark

method when the two methods have a nearly equal number of false alarms. To evaluate the labeling ability, a class of “none” to represent missed detections and false alarms is manually added. For example, in Fig. 5b, 487 pedestrian instances are correctly positioned and labeled by the proposed method; 2 are wrongly labeled to be bicycle riders, and 21 are miss-detected. More results are shown in Fig. 6.

5.2 Wild-scene Detection

Dataset In order to show that our method can be used for general purposes, we test our method on complicated scenes, especially, complicated background. Even in these cases, our method works well, which shows robustness of our method. A mini dataset is built upon leopards and tigers of the family Felidae. Note especially that the image feature used by this method belongs to the type texture, and texture from different positions of the leopards are almost the same. The dataset contains 6 video clips of 9 leopards and 4 tigers. The frame size is 640×480. Both of the animals are in the side view.

Implementation settings Most implementation settings are the same as the settings used for campus object detection. For training, 5 leopards and 2 tigers are used. The size of the image patch around each keypoint is 27×27.

Comparisons In Fig. 7, the motion grouping results, and how the voted centers are affected, are given. Since parts from different positions of the leopard are very similar, the true center of a leopard is difficult to find using the voted centers of the object parts. In Fig. 8, example Hough images are given to show the merit of the proposed prior by the ability to detect leopards. In Fig. 9, the detection results are given. The proposed method successfully localizes and labels all the leopards and tigers, while the benchmark method miss-detects three leopards.

6 Conclusion

The computational ability of human beings is limited, while their ability to detect is far beyond that of machines. Thus, it is very possible that this detection ability benefits from multiple perceptual mechanisms. By using one of these mechanisms, we propose a detection method. By embedding motion grouping results into the voting schema of hough transforms, the method is able to distinguish near objects' positions, distinguish similar objects' labels, and maintain the detection rate with a noisy codebook. The success of our method further demonstrates the advancement of perceptual mechanisms in human beings. And the success of this method will help with detection methods in ITS areas.

Acknowledgments This work was, in part, supported by SCOPE program by the Ministry of Internal Affairs and Communications. Part of the work was done during the first author's master course at Peking University. The first author is sponsored by China Scholarship Council. The authors thank Bo Zheng, Boxin Shi, Patricia Knapp, and Nicole Hajicek for suggestions.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

1. Ballard, D.: Generalizing the hough transform to detect arbitrary shapes. *Pattern Recog.* **13**(2), 111–122 (1981)
2. Barinova, O., Lempitsky, V., Kohli, P.: On detection of multiple object instances using hough transforms. In: *CVPR*, pp. 2233–2240 (2010)
3. Brostow, G., Cipolla, R.: Unsupervised Bayesian detection of independent motion in crowds. In: *CVPR*, vol. I, pp. 594–601 (2006)
4. Brox, T., Malik, J.: Object segmentation by long term analysis of point trajectories. In: *ECCV*, vol. V, pp. 282–295 (2010)
5. Felzenszwalb, P., Huttenlocher, D.: Pictorial structures for object recognition. *IJCV* **61**(1), 55–79 (2005)
6. Fergus, R., Perona, P., Zisserman, A.: Object class recognition by unsupervised scale-invariant learning. In: *CVPR*, vol. II, pp. 264–271 (2003)
7. Gall, J., Lempitsky, V.: Class-specific hough forests for object detection. In: *CVPR*, pp. 1022–1029 (2009)
8. Gu, C., Lim, J., Arbelaez, P., Malik, J.: Recognition using regions. In: *CVPR*, pp. 1030–1037 (2009)
9. Harris, C., Stephens, M.: A combined corner and edge detector. In: *Alvey Vision Conference*, pp. 147–152 (1988)
10. Lampert, C., Blaschko, M., Hofmann, T.: Beyond sliding windows: object localization by efficient subwindow search. In: *CVPR*, pp. 1–8 (2008)
11. Lehmann, A., Leibe, B., Van Gool, L.: Fast prism: branch and bound hough transform for object class detection. *IJCV* **94**(2), 175–197 (2011)
12. Leibe, B., Cornelis, N., Cornelis, K., Van Gool, L.: Dynamic 3D scene analysis from a moving vehicle. In: *CVPR*, pp. 1–8 (2007)
13. Leibe, B., Leonardis, A., Schiele, B.: Robust object detection with interleaved categorization and segmentation. *IJCV* **77**(1–3), 259–289 (2008)
14. Leibe, B., Schiele, B.: Interleaved object categorization and segmentation. In: *BMVC*, pp. 759–768 (2003)
15. Maji, S., Malik, J.: Object detection using a max-margin hough transform. In: *CVPR*, pp. 1038–1045 (2009)
16. Mikolajczyk, K., Leibe, B., Schiele, B.: Multiple object class detection with a generative model. In: *CVPR*, vol. I, pp. 26–36 (2006)
17. Mohottala, S., Ono, S., Kagesawa, M., Ikeuchi, K.: Fusion of a camera and a laser range sensor for vehicle recognition. In: *OTCBVS*, pp. 16–23 (2009)
18. Ohba, K., Ikeuchi, K.: Detectability, uniqueness, and reliability of eigen windows for stable verification of partially occluded objects. *PAMI* **19**(9), 1043–1047 (1997)
19. Okada, R.: Discriminative generalized hough transform for object detection. In: *ICCV*, pp. 2000–2005 (2009)
20. Ommer, B., Malik, J.: Multi-scale object detection by clustering lines. In: *ICCV*, pp. 484–491 (2009)

21. Tomasi, C., Kanade, T.: Detection and Tracking of Point Features. Tech. rep., Carnegie Mellon University (1991)
22. Tuzel, O., Porikli, F., Meer, P.: Region covariance: A fast descriptor for detection and classification. In: ECCV, vol. II, pp. 589–600 (2006)
23. Wertheimer, M.: Laws of organization in perceptual forms (partial translation). In: Ellis, W.B. (ed) A Sourcebook of Gestalt Psychology, pp. 71–88. Harcourt, Brace (1938)
24. Wojek, C., Roth, S., Schindler, K., Schiele, B.: Monocular 3D scene modeling and inference: Understanding multi-object traffic scenes. In: ECCV, vol. IV, pp. 467–481 (2010)
25. Yang, L., Zheng, N., Chen, M., Yang, Y., Yang, J.: Categorization of multiple objects in a scene without semantic segmentation. In: ACCV (2009)
26. Yarlagadda, P., Monroy, A., Ommer, B.: Voting by grouping dependent parts. In: ECCV, vol. V, pp. 197–210 (2010)
27. Yeh, T., Lee, J., Darrell, T.: Fast concurrent object localization and recognition. In: CVPR, pp. 280–287 (2009)

Zhipeng Wang received his B.S. degree in Industrial Engineering from Tsinghua University, China, in 2007. He received his M.S. degree in Computer Applied Technology from Peking University, 2010 and is currently a PhD student in the Department of Computer Science, the University of Tokyo. His research interests are in the areas of computer vision and machine learning.

Jinshi Cui received the B.S. and PhD degrees in computer science from Tsinghua University, Beijing, China, in 1999 and 2004, respectively. In 2004, she joined the School of Electronics Engineering and Computer Science, Peking University, Beijing, as an Assistant Professor. She was promoted to Associate Professor in 2007. Her research interests include computer vision and robotics.

Speech, and Signal Processing, Surrey University, Surrey, U.K. Since 2000, he has been a Professor with the Center for Information Science, School of Electronics Engineering and Computer Science, Peking University, Beijing, China. He has published more than 200 technical publications in journals, books, and international conference proceedings. His research interests include computer vision, digital geometry processing, and robotics. Dr. Zha was the recipient of the Franklin V. Taylor Award from the IEEE Systems, Man, and Cybernetics Society in 1999.

Hongbin Zha received the B.E. degree in electrical engineering from Hefei University of Technology, Hefei, China, in 1983 and the M.S. and Ph.D. degrees in electrical engineering from Kyushu University, Fukuoka, Japan, in 1987 and 1990 respectively. After being a Research Associate with Kyushu Institute of Technology, he joined Kyushu University in 1991 as an Associate Professor. In 1999, he was also a Visiting Professor with the Centre for Vision,

Masataka Kagesawa received the BS degree in mathematics in 1986 from Chiba University, Japan, the MS degree in mathematics from Tokyo Metropolitan University in 1988. He was a doctor course student at Tokyo Metropolitan University from 1988 to 1990. From 1990 to 1994, he was a technical associate at the Institute of Industrial Science, the university of Tokyo. He is now a research associate at the same institute.

He received the PhD degree in electronic engineering in 2003 from the university of Tokyo. His research interests include traffic simulation with dynamic information, traffic management systems and sensing systems on Intelligent Road Traffic Systems.

Shintaro Ono received the BE degree in 2001 and PhD degree in 2006 from The University of Tokyo. Currently he is a Project Associate Professor in Advanced Mobility Research Center (ITS Center), The University of Tokyo. His research interests include sensing system and computer vision/graphics for ITS, and digital archiving of cultural heritage objects.

Katsushi Ikeuchi is a Professor at the Institute of Industrial Science, the University of Tokyo, Tokyo, Japan. He received the PhD degree in Information Engineering from the University of Tokyo, Tokyo, Japan, in 1978. After working at the Artificial Intelligence Laboratory, MIT for three years, the Electrotechnical Laboratory, MITI for five years, and the School of Computer Science, Carnegie Mellon University for ten years, he joined the university in 1996.

He has served (or will serve) as the program/general chairman of several international conferences, including 1995 IEEE-IROS (General), 1996 IEEE-CVPR (Program), 1999 IEEE-ITSC (General) and 2003 ICCV (Program). He is an Editor-in-Chief of the International Journal of Computer Vision. He has been a fellow of IEEE since 1998. He was selected as a distinguished lecture of IEEE SP society for the period of 2000–2001.

He has received several awards, including the David Marr Prize in computational vision, IEEE RA society K-S Fu memorial best transaction paper award, Computer Vision Significant Researcher Award, and Medal with Purple Ribbon. In addition, in 1992, his paper, "Numerical Shape from Shading and Occluding Boundaries," was selected as one of the most influential papers to have appeared in Artificial Intelligence Journal within the past ten years.