

Detection and Orientation Estimation for Cyclists by Max Pooled Features

Abstract

In this work we propose a new kind of HOG feature which is built by the max pooling operation over spatial bins and orientation channels in multilevel and can efficiently deal with deformation of objects in images. We demonstrate its invariance against both translation and rotation in feature levels. Experimental results show a great precision gain on detection and orientation estimation for cyclists by applying this new feature on classical cascaded detection frameworks. In combination of the geometric constraint, we also show that our system can achieve a real time performance for simultaneous cyclist detection and its orientation estimation.

Index Terms

Cyclist Detection, Orientation Estimation, Max Pooled Features.

I. INTRODUCTION

Influenced by the environmental friendly and healthy lifestyle, the number of people riding bicycles is increasing nowadays [1]. In the meanwhile, the safety of cyclists is still a hot research topic due to the even more complicated traffic conditions [2]. To solve this problem technically, various detection systems are developed, which can recognize cyclists in critical scenarios and warn the other traffic participants to prevent probable accidents. Among a great number of solutions provided in the market, including radar, lidar, acoustic and inductive sensors [3] [4], vision-based detection system has become a favorite choice of industries because of its discriminative power and installation flexibility.

Despite a great success of vision-based detection approaches achieved for other vulnerable road participants such as pedestrians [5] [6], the task for cyclist detection still remains challenging. E.g. in the KITTI benchmark [7] the number of published detection methods for cyclists is less than half of that for pedestrians (Table I). One of the reasons is that the appearance of cyclists is significantly influenced by their viewpoints, behaviors and the sensor setups. The appearance variation not only results in different aspect ratios but also leads to full or partial translation or rotation of an object in the image (Figure 1). Such kind of deformation is difficult to deal with. Furthermore, in most cases, both the location and the orientation of cyclists are required to estimate their trajectories and further to predict probable accidents. This point is rarely considered in most detection approaches. Last but not least, the detection system should be able to run in real time, which is important for most of the on-road safety systems to react without delays.

To tackle these problems, in this work, we at first propose a new kind of HOG (Histogram of Oriented Gradients) feature which is built by a max pooling operation over spatial bins and orientation channels in multilevel. To calculate this kind of feature efficiently, we implement a framework utilizing a pipeline fashion. The robustness of this new feature against both translation and rotation of objects can be demonstrated in feature levels. In addition to that, we present a framework, which is in combination with this kind of new feature and can achieve good performance on simultaneous detection and orientation estimation for cyclists. With the help of geometric constraints, we show that our detector is capable for real time applications.

II. RELATED WORKS

In the last decade numerous research works have been published about object detection. The feature family meets an explosive growth of its members, e.g. Haar [8], HOG [9], LBP [10], sp-Cov [11], deep features [12] and their extensions [13] [14]. The detector structure is also becoming even more complicated, varying from cascades [15] to DPMs [16] and CNNs [17]. While most of the proposed methods are focused on detecting pedestrians, cyclist detection also becomes an interesting point among the research community.

In early works, several models are proposed to represent bicycles, e.g. by wheel shapes [18], by edges and motion [19] and even by DPMs [20]. Despite high detection rates achieved, without considering the bicycle riders, their classification power

TABLE I

NUMBER OF PUBLISHED DETECTION METHODS FOR EACH OBJECT CLASS IN THE KITTI BENCHMARK TILL AUGUST 2016. THE NUMBER OF PROPOSED METHODS FOR CYCLISTS IS LESS THAN HALF OF THAT FOR PEDESTRIANS, ALTHOUGH BOTH OF THEM BELONG TO THE WEAK TRAFFIC PARTICIPANTS.

Object class	Car	Pedestrian	Cyclist
Number of methods	25	29	14




Fig. 1. KITTI image samples. Image (b) represents a full rotation of an object by 20° in comparison with image (a). The partial rotation of an object is presented by image pairs (c)-(d), i.e. the torso of the cyclist in image (d) is rotated by 15° , comparing with image (c). Image pairs (e)-(f) show the full translation of an object in the side view. Compared to image (g), the head of the cyclist has a lower position in image (h), which corresponds to a partial translation.

can be disturbed in scenarios like bicycle park stations. In comparison, Li et al. uses HOG-LP [21] and Takahashi et al. takes advantage of pedaling movement [22] to detect cyclists mainly in crossing cases, with limited applications. For multi-view detection, an approach combining RGB image with lidar measurements is presented by Gonzalez et al. in [23], yet inappropriate for monocular camera systems. Xiang et al. integrates subcategory detection into region proposals [24], still at a relative high computational cost. In the works of [25] [26], viewpoint specific models are built but the detection suffers from the appearance variation. To deal with object deformation, Zou [27] and Paisitkriangkrai [11] choose spatially pooled features and Pepik [28] introduces 3-D DPM. Although they achieve high precision, their detectors cannot run in real time and the rotation handling is not included. As deep learning becomes a trend, by applying CNN with region proposals, Chen et al. achieves the best results both for cyclist detection and orientation estimation at one time [29]. By integrating the cascade concept, the detection precision and speed are further improved in [30]. But the cost is still a high performance GPU. Although such kind of GPUs can be facilitated with more and more equipments nowadays and run much faster than normal CPUs, the large consumption of memory and even of power make themselves inappropriate for products with limited budgets.

In this paper we show that with the help of max pooling, even low level visual features, i.e. HOG, can efficiently deal with both translation and rotation of objects in images. We also give a demonstration in feature levels. Moreover, we integrate these new features into the framework [26] and extend it with an orientation estimator, which runs almost simultaneously with the object detector. We conduct experiments on the KITTI dataset and show that the max pooled features contribute to a significant performance gain both on detection and orientation estimation for cyclists.

III. DEFORMATION INVARIANT FEATURES BY MAX POOLING

The max pooling operation arises from the concept of CNNs to construct the deep network layers. [31] [32] The principle is to apply a filter mask over a feature map and to search the maximum value inside it. As long as the maximum is not changed, the output value is not influenced by its location within the mask. Inspired by that, we can also apply max pooling on low level visual features to obtain the deformation invariance. Here we choose the HOG features from [16] as the base, not only due to their strong discriminative power but also because both spatial and orientation information are included. In fact, the original HOG features are not translation or rotation invariant. Although interpolation is recommended on both spatial and orientation bins, it can only deal with small appearance variations. Regarding this, we introduce the max pooling with different scales to build deformation invariant HOG features.

Note that max pooling is also applied in other features such as the sp-Cov [11] and Regionlets [27], but our work differs from them in 3 points.

- Unlike the covariance patterns utilized in [11], in this work, we only focus on the HOG features and their performance improvement by additional max pooling procedure.
- Instead of random feature selection in [27], we apply a dense max filtering with different scales.
- We additionally conduct max pooling on orientation histograms to handle object rotations, which makes another significant difference.

A. Translation Invariant HOG Features

For translation invariance we follow the idea in [11] and apply the max filters directly on each feature channel. Here we carefully choose 4 filter mask sizes: 1×1 , 2×2 , 3×3 and 4×4 pixels, so that none of them is greater than the smallest dimension of our detector models. The striding step for each filter is set to 1 pixel and the input channels are padded with an appropriate number of border pixels so that the output feature channels will have the same size as the inputs.

Instead of performing max filtering in each scale separately, here we use an iterative approach for computational efficiency. The calculation is done in a pipeline. Each stage is a max pooling operation with a mask size of 2×2 pixels. The filtered features from the previous stage are reused as the inputs for the next stage. This is equivalent to a decomposition of bigger




Fig. 2. Calculating translation invariant HOG features. On the top are max filters with 4 mask sizes: 1x1, 2x2, 3x3 and 4x4. The first size corresponds to the naive HOG features. On the left side are two samples with full translation and an overlap ratio of 0.7. The HOG features are calculated for each image and illustrated in temperature map. The max filters are applied on each channel in a pipeline fashion.




Fig. 3. Correlation coefficient values for each filter size are illustrated by the boxplot. The median value is indicated by a red short line. The blue box represents 50 percentile of the data. The red point denotes the outlier.

filters into several smaller ones. E.g. the output after a two stage filtering is the same as that filtered by a mask of 3x3 pixels (Figure 2). And this output can be reused to obtain the filtered feature map by the mask of 4x4 pixels with the help of only one more stage filtering in the pipeline. In this way, the processing time for each filter size is constant and only dependent on the number of utilized filters.

To explore the translation invariance of filtered channels, we take the sample pair of full translation from Figure 1 to execute experiments. As in Figure 2, these two images have an overlap ratio of 0.7, which can be resulted from an improper striding step of the detection window or from an inappropriate scaling factor of the image. Here we also omit the discussion about partial translation because it can be divided into fully translated and unchanged parts. As for experiment, we calculate the HOG features for each image and filter them by max pooling in different scales sequentially. To measure the similarity between the features of both images, we calculate the correlation coefficient for each filtered channel between the two images and sort the results into 4 groups according to the size of filter masks. Distributions of coefficient values for each filter size are illustrated by the boxplot in Figure 3. As can be seen, the high values of correlation coefficients are obtained by big filters, which means the HOG features of both images filtered by bigger masks are more similar to each other despite that the same object is translated in the second image. Thus, the spatially max pooled features become more robust against translations, which is in accordance with the conclusion from [27].

To model translations in different scales, we concatenate the filtered channels by all the above masks (including the filter mask 1x1) to each other as new features. We do it in this way because of two points. On one hand, we would like to keep the discriminative power of the naive HOG features and to avoid the drop of detection accuracy by only using blurred features from max filtering. On the other hand, we expect the improvement of detection on translated images, especially with the help of the new features, which are filtered by bigger masks. In this case, the channel number of HOG features is quadrupled. Here we call the spatially pooled HOG features as maxHOG_T.




Fig. 4. Calculating rotation invariant HOG features. On the top are max filters with 3 mask sizes: 1x1, 2x1 and 3x1. The first size corresponds to the naive HOG features. On the left side are two samples with full rotation and the gap of roll angle is 20°. The HOG features are calculated for each image. The max filters are applied on each histogram in a pipeline.

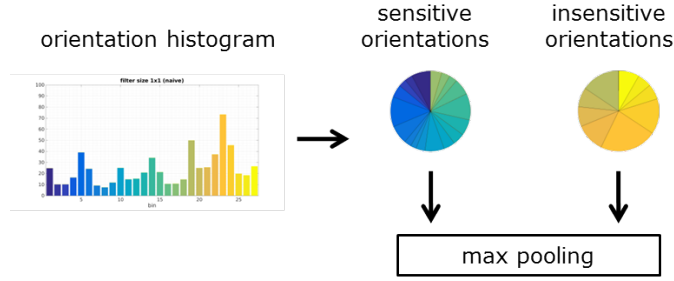


Fig. 5. Orientation histogram is divided into 2 groups, which represent the sensitive and the insensitive orientations respectively. For each of them we link the first and the last orientation bin to build a ring shape. These new histograms will be processed by max pooling separately.

B. Rotation Invariant HOG Features

As gradient orientations are discretized in equidistant histogram bins, the rotation of an object has influence mainly on their distribution not on the bin sums. Here we have totally 27 bins (18 sensitive and 9 insensitive orientation bins, same as [16]). We also apply max filtering on the histogram with 3 mask sizes: 1x1, 2x1 and 3x1. The striding step is set to 1 orientation bin. As the last filter size corresponds to a maximal rotation of 60°, which is acceptable for cyclists in most cases, we do not use even bigger masks. The filtered channels are also calculated in a pipeline fashion by decomposing the biggest filter mask into two smaller ones, each with a size of 2x1 pixels, as illustrated in Figure 4.

Aware that the orientation bins are actually located in a circular form, the max filtering should also be done in the same way. Here we divide the orientation histogram into 2 groups, which represent the sensitive and the insensitive orientations respectively. For each of them we link the first and the last orientation bin to build a ring shape (Figure 5). These new histograms will be processed by max pooling operation separately.

To verify the rotation invariance of filtered features, we also take the sample pair of full rotation from Figure 1. With the same reason, we omit the discussion about partial rotations. We calculate the orientation histograms for each image and perform the max filtering in 3 scales as in Figure 4. Then we compute the correlation coefficients for each filtered histogram between both images. The coefficient value of each filter size is registered in Table II. Obviously, the similarity between histograms of both images increases with bigger filter masks, which infers the HOG features become more robust against rotations.

In the same way, we concatenate the filtered feature channels by different masks (including the filter mask 1x1). This new kind of feature is called as maxHOG_R and utilized in combination with maxHOG_T for cyclist detection.

For an intuitive comparison of the classification power between the naive HOG features and the proposed new ones, we conduct principal component analysis (PCA) on both of them. Here we randomly select 2000 positive and negative samples for the front view of cyclists. We calculate HOG features for each sample and perform max filtering both on spatial and orientation

TABLE II
CORRELATION COEFFICIENT OF EACH FILTER SIZE APPLIED ON ORIENTATION HISTOGRAMS.

Filter size	1x1	2x1	3x1
Correlation coefficient	0.54	0.69	0.82

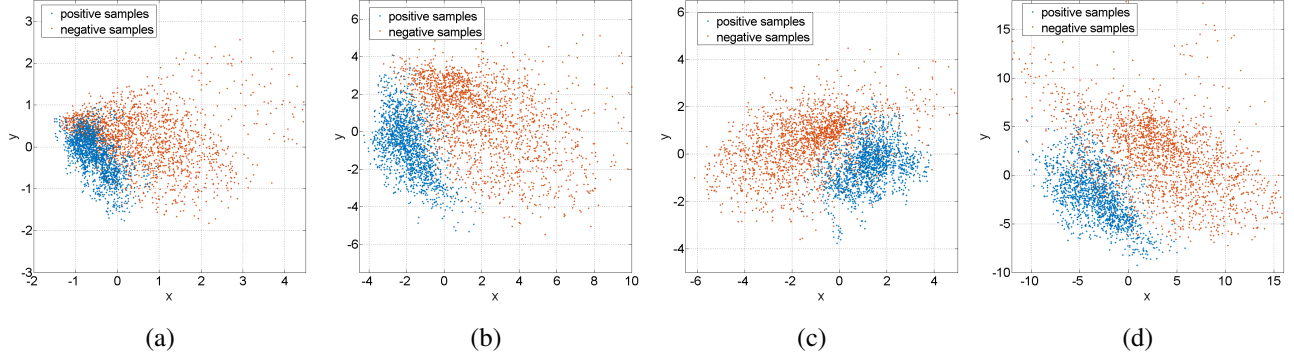


Fig. 6. PCA for 2000 positive and negative samples, which are represented in blue and red respectively. Totally we calculate 4 different HOG feature types for each sample: (a) the naive HOG, (b) the maxHOG_T , (c) the maxHOG_R and (d) the maxHOG ($\text{maxHOG}_T + \text{maxHOG}_R$) features. For each kind of feature, we project them onto the first two principal components of PCA.

bins. We project the sample features onto the first two principal components from PCA and plot them for each HOG feature type in Figure 6. Obviously, with the help of max filtering either on spatial or on orientation bins, the positive and negative samples can be better distinguished from each other than only with the naive HOG features, which is the benefit from dealing with object deformations. For a quantitative analysis of the classification power of each HOG type, please refer to Section V.

IV. DETECTION FRAMEWORK

Here we reuse our previous framework from [26] due to its capability for multi-view detection and computational efficiency. The viewpoints of cyclists are divided in 8 subcategories, each with a range of 45° (Figure 7). For each subcategory a cascaded detector is built, consisting of n decision forests (DF) and 1 support vector machine (SVM_C) at the last stage (Figure 8). The number of n is chosen between 2 to 3, which is a trade-off between detection precision and processing time [26]. The detectors are applied in sliding window fashion with 3 aspect ratios: 0.5 for viewpoint II and VI, 1.0 for IV and VIII, and 0.75 for the others. The minimal window height is set to 80 pixels.

A. Feature Processing and Orientation Estimation

The processing of HOG features is divided into two steps: max pooling over spatial and orientation bins. In the second step, we apply max filters not only on the naive HOG features but also on the spatially filtered ones. By doing this, we are able to deal with simultaneous translation and rotation of objects. Note that there are 4 channels of summed gradients from [16], which do not contain the orientation information, we only perform spatial max filtering on those channels. We call the processed HOG features as maxHOG . Since only bitwise operations are performed, the whole max pooling procedure can be done very fast. E.g. for an KITTI image, it takes only about 0.02 seconds (Table III). Such new features are calculated for samples from the KITTI dataset to train our detectors. The decision trees have a depth of 2 and are trained by AdaBoost. The detector from a previous cascade stage will be tested on predefined image set to gather false positives, which are used as negative samples to train the classifier in the next stage. For further details about training and implementation please refer to [26].

Since the orientation of cyclists are continuous values, a direct approximation by the discrete viewpoints can increase the estimation errors. Instead, we add another support vector machine SVM_R for regression. As shown in Figure 8, we reuse the detection results of the cascaded detector as inputs for SVM_R . Regarding that the number of processed samples at the last stage is small due to the strong filtering of frontal DFs, the additional computational cost for the orientation estimation is also little, i.e. about 0.005 seconds per image (Table III). Therefore, both the cyclist detection and orientation estimation can be done almost simultaneously. To train the regression stage, we use the same samples as for training the detectors.

For further improving processing efficiency, at the beginning of the detection, we only apply the detectors on the given image. The final object hypotheses are obtained by performing non-maximum-suppression on the outputs of all individual detectors. The orientation estimation is only performed on the detection results with the highest scores.

B. Selective Detection by Geometric Constraints

Given a fixed striding step, the number of searching locations grows with the square of the image size and so is the processing time. This fact makes some applications with high resolution images inefficient, if further details on objects like head orientations or body poses are interested. To solve this problem, one of the common approaches is to extract region of interest (ROI) in images, where objects can appear.



Fig. 7. Viewpoints of cyclists are divided into 8 equidistant subcategories I to VIII.




Fig. 8. Cascaded detector, consisting of n decision forests DF_1 to DF_n and one stage of support vector machine SVM_C . The additional SVM_R is for orientation estimation. The naive HOG features are preprocessed by max pooling operation and then given into the detector.

Here we prefer the geometric constraint [26] [33], which does not require additional sensor information and the computational burden is mere. The key idea of this approach is to mathematically formulate the dependence between the height h of an object and its ground coordinates $[u, v]$ in the image. This dependence is called geometric constraint. As shown in Figure 9 (a) and (b), given detectors with a height of 80 pixels and assuming a real height of 2 meters, the objects can only appear in a small image region between two horizontal lines. For varying heights in the 3-D world, e.g. from 1 meter to 2 meters, the corresponding ROI will be vertically increased, but only by a few pixels, as in Figure 9 (c). Since detection only takes place in this small region, the processing time can be greatly reduced, by about 2/3 in our case (Table III).

V. EXPERIMENT RESULTS

In the following experiments, we explore the performance of our detection system by evaluating it with the KITTI dataset, which consists of 7518 test images with a size of 1242×375 pixels. These images are taken in various traffic scenes including

TABLE III
AVERAGE PROCESSING TIME FOR EACH STEP IN THE CASCADE. +ROI DENOTES INTEGRATION OF ROI EXTRACTION BY GEOMETRIC CONSTRAINT INTO THE DETECTION FRAMEWORK.

+ROI	naive HOG	max pooling	detection	orientation estimation
no	0.106 s	0.024 s	0.090 s	0.005 s
yes	0.036 s	0.008 s	0.031 s	0.002 s




Fig. 9. Geometric constraint. (a) illustrates an object with a height of 2 meters in the real world. The corresponding region in the camera image with an object height of 80 pixels is indicated in (b) by the color red. The green region in (c) represents the region of objects with a varying height from 1 to 2 meters. The blue arrow denotes the estimated horizon position.

TABLE IV
AVERAGE PRECISION OF CYCLIST DETECTION WITH DIFFERENT HOG FEATURES. +ROI DENOTES INTEGRATION OF ROI EXTRACTION BY GEOMETRIC CONSTRAINT.

Method	Moderate	Easy	Hard
naive HOG + ROI	27.32%	29.74%	25.47%
maxHOG _T + ROI	38.63%	41.88%	34.24%
maxHOG + ROI	43.58%	48.38%	38.73%

TABLE V
AVERAGE PRECISION OF ORIENTATION ESTIMATION FOR CYCLISTS WITH DIFFERENT HOG FEATURE TYPES. +ROI DENOTES INTEGRATION OF ROI EXTRACTION BY GEOMETRIC CONSTRAINT.

Method	Moderate	Easy	Hard
naive HOG + ROI	22.15%	24.71%	20.93%
maxHOG _T + ROI	32.97%	35.68%	29.11%
maxHOG + ROI	38.28%	41.82%	34.27%

rich number of objects, such as cars, pedestrians and cyclists. As the smallest object evaluated by KITTI benchmark is 25 pixels high and our detector has a minimal height of 80 pixels, we just scale the test images by a factor of 3. For processing efficiency, we also apply the geometric constraint with the help of provided calibration data and turn on the multi-threading. Our test computer is equipped with an Intel core of i7-3740QM and a memory of 8 GB.

The first experiment was carried out to reveal the performance gain of the max pooling operation on HOG features. Here we compare the test results with 3 feature types: the naive HOG features, the spatially pooled maxHOG_T and the maxHOG features, with max filtering both on spatial and orientation bins. The corresponding precisions both on detection and orientation estimation for cyclists are listed in Table IV and V. The precision-recall-curves are illustrated in Figure 10. Apparently, with the help of spatial max pooling in different scales, the average precision of the approach with naive HOG features is increased by about 10% for all difficulty levels, which benefits from a better handling of translation of objects. Additionally, it can be further pushed by about 6%, if the max pooling is also applied on orientation histograms, which proves the ability of our new features in dealing with rotations.

Compared to that, a similar trend of precision gain can also be seen with the orientation estimation for cyclists. In fact, the test results of estimated orientation may not be necessarily consistent with the results of detections. However, in our approach, both of them are strongly correlated with each other. This fact implies that our orientation estimator works well for accurately detected cyclists. It also explains why the average precision of orientation estimation is slightly lower than that of detection, because both the detection and orientation approximation errors are accumulated in this case.

In the next step, we conduct experiments with different scaling factors of KITTI images. Here we apply the max pooling

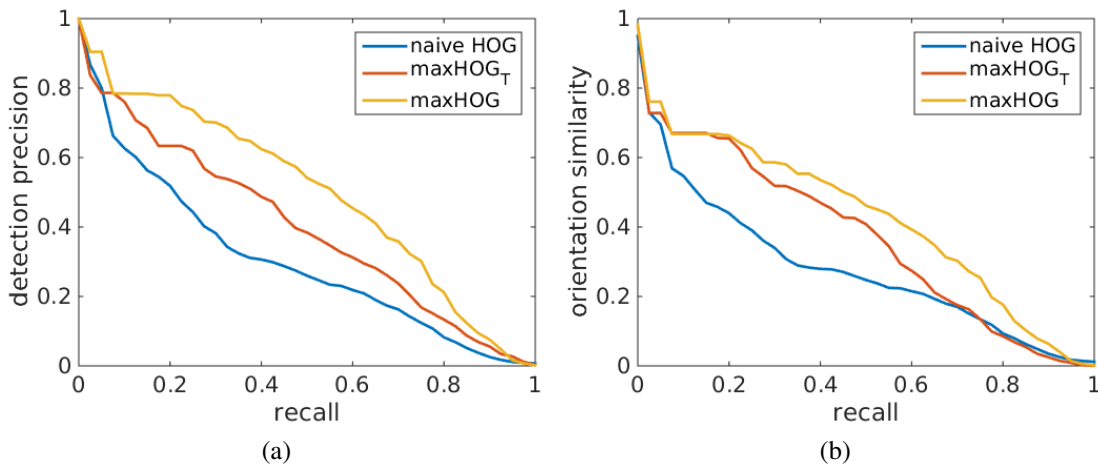


Fig. 10. Precision-recall-curves of cyclist detection and orientation estimation are respectively plotted in (a) and (b). Here we present curves in easy level with 3 feature types: the naive HOG, the spatially pooled maxHOG_T and the maxHOG features with max filtering both on spatial and orientation bins.

TABLE VI
AVERAGE DETECTION PRECISION AND PROCESSING TIME WITH DIFFERENT SCALING FACTORS OF KITTI IMAGE.

Scaling factor	Moderate	Easy	Hard	Runtime
1	33.25%	37.88%	29.69%	0.08 s
2	39.17%	43.14%	34.98%	0.17 s
3	43.58%	48.38%	38.73%	0.25 s

TABLE VII
AVERAGE PRECISION ON CYCLIST DETECTION OF PUBLISHED METHODS ON KITTI BENCHMARK TILL AUGUST 2016. PROPOSED METHOD OF THIS WORK IS MARKED IN BOLD.

Method	Moderate	Easy	Hard	Runtime	Environment
SDP+RPN [30]	73.74%	81.37%	65.31%	0.4 s	GPU@2.5Ghz
SubCNN [24]	71.06%	79.48%	62.68%	2 s	GPU@2.5Ghz
3DOP [29]	68.94%	78.39%	61.37%	3 s	GPU@2.5Ghz
Mono3D [35]	66.36%	76.04%	58.87%	4.2 s	GPU@2.5Ghz
SDP+CRC (ft) [30]	61.31%	74.08%	53.97%	0.6 s	GPU@2.5Ghz
Regionlets [27]	58.72%	70.41%	51.83%	1 s	12 cores@2.5Ghz
maxHOG+ROI	43.58%	48.38%	38.73%	0.25 s	4 cores@2.5Ghz
MV-RGBD-RF [23]	42.61%	52.97%	37.42%	4 s	4 cores@2.5Ghz
pAUCEnsT [11]	38.03%	51.62%	33.38%	60 s	1 core@2.5Ghz
Vote3D [36]	31.24%	41.43%	28.60%	0.5 s	4 cores@2.8Ghz
DPM-VOC+VP [28]	31.08%	42.43%	28.23%	8 s	1 core@2.5Ghz
LSVM-DPM-us [37]	29.88%	38.84%	27.31%	10 s	4 cores@3.0Ghz
DPM-C8B1 [38]	29.04%	43.49%	26.20%	15 s	4 cores@2.5Ghz
LSVM-DPM-sv [37]	27.50%	35.04%	26.21%	10 s	4 cores@3.0Ghz
mBoW [39]	21.62%	28.00%	20.93%	10 s	1 core@2.5Ghz

operation both on spatial and orientation bins of HOG features, as this configuration yields the best results. Corresponding detection precision and processing time are listed in Table VI. Obviously, both their values decrease if the scaling factor shrinks. Note that both the geometric constraint and multi-threading are utilized in our experiments, the processing time is not exactly square proportional to the image size. And for an image without scaling (i.e. the scaling factor equals 1), the processing time is reduced to only 0.08 seconds, yet the precision is about 10% less than the best one. One point to be mentioned is that the evaluation by KITTI is performed while a lot of objects with small image sizes are considered [34]. If we are only interested in objects within a small distance range, then no scaling is necessary and our detector can also yield both a reasonable precision value and a runtime speed of more than 10 frames per second (fps), which is sufficient for most real time requirements.

In another experiment we analyze the performance of our detection system in comparison with other published methods in the KITTI benchmark. For fairness, the average precision of cyclist detection and orientation estimation are compared separately, since not all the proposed methods take simultaneous detection and orientation estimation into count. Their corresponding precision values are listed in Table VII and VIII respectively. Additionally, we give an overview about their processing time with respect to the corresponding hardware environments. As can be seen, our proposed method especially outperforms the DPM-based approaches in the ranking list. One reason is that, although the DPMs can also deal with translation of objects, by including rotation invariant features, our approach is more efficient to handle object deformation.

Nevertheless, there is still a gap of up to 30% in accuracy between our proposed method and the top performed ones (actually the top 5 in Table VII). We owe it to the classification power of the feature itself. As all of them use deep features, comparing with the HOG features alone in our approach, they can capture more information from the object and thus their detectors are more discriminative. Therefore, the next step in our future work may be to integrate these deep features into our current framework to further improve the detection performance. However, better accuracy comes at the cost of large computation time and high performance hardware (e.g. GPUs). In comparison, our method yields a runtime performance of about 4 fps on a quad-core CPU. Considering that we have scaled the test image by a factor of 3 to fit the small objects to our detectors, we can also promise an even faster processing ability such as 12 fps (Table VI) in use cases, in which only objects not far away from the camera are considered, e.g. to monitor the cyclists, which are located in the nearby blind field of a truck.

TABLE VIII
AVERAGE PRECISION ON ORIENTATION ESTIMATION FOR CYCLISTS OF PUBLISHED METHODS ON KITTI BENCHMARK TILL AUGUST 2016. PROPOSED METHOD OF THIS WORK IS MARKED IN BOLD.

Method	Moderate	Easy	Hard	Runtime	Environment
SubCNN [24]	63.65%	72.00%	56.32%	2 s	GPU@2.5Ghz
3DOP [29]	58.68%	70.13%	52.35%	3 s	GPU@2.5Ghz
Mono3D [35]	54.97%	65.56%	48.77%	4.2 s	GPU@2.5Ghz
maxHOG+ROI	38.28%	41.82%	34.27%	0.25 s	4 cores@2.5Ghz
DPM-VOC+VP [28]	23.17%	30.52%	21.58%	8 s	1 core@2.5Ghz
LSVM-DPM-sv [37]	22.07%	27.54%	21.45%	10 s	4 cores@3.0Ghz
DPM-C8B1 [38]	19.25%	27.25%	17.95%	15 s	1 core@2.5Ghz

VI. CONCLUSION

Cyclist detection becomes an essential part of the modern intelligent transportation systems, as their safety draws even more attention from the publics. Despite numerous vision-based methods proposed, the state-of-the-art techniques either are troublesome to handle the appearance deformation of image objects or are difficult to provide a real time performance. The contribution of this work is based on two key ideas: firstly, we propose a new kind of feature which can efficiently deal with both translation and rotation of objects in images, based on max pooling over spatial and orientation histogram bins in multilevel. In addition, we present a cascaded framework, which is able to conduct cyclist detection and orientation estimation simultaneously due to shared structure and visual features. The evaluation on the KITTI benchmark demonstrates good accuracy of our approach at a comparable small computation cost. Experiments also have shown that, leveraging the fast implementation of feature preprocessing (i.e. multi-scale max pooling of HOG features) and geometric constraint, our system is able to run in real time (varying from about 4 to 12 fps according to configurations), especially promising for applications on hardwares with limited performance.

REFERENCES

- [1] M. Pape, "Cycling mobility in the EU," *Members' Research Service*, vol. 557013, 2015.
- [2] EU, "Traffic Safety Basic Facts 2015," *European Road Safety Observatory*, 2015.
- [3] R. Dharmaraju, D. A. Noyce, and J. D. Lehman, "An Evaluation of Technologies for Automated Detection and Classification of Pedestrians and Bicyclists," May 2002.
- [4] J. V. Krogmeier and D. M. Bullock, "Inductive Loop Detection of Bicycles and Inductive Loop Signature Processing for Travel Time Estimation," *Statewide Wireless Communications Project*, vol. 2, 2008.
- [5] P. Dollár, C. Wojek, B. Schiele, and P. Perona, "Pedestrian Detection: An Evaluation of the State of the Art," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 34, 2012.
- [6] S. Zhang, R. Benenson, M. Omran, J. H. Hosang, and B. Schiele, "How Far are We from Solving Pedestrian Detection?" *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [7] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [8] P. Viola and M. Jones, "Robust Real-Time Face Detection," *International Journal of Computer Vision*, no. 2, pp. 137–154, 2004.
- [9] N. Dalal and B. Triggs, "Histograms of oriented gradients for human detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2005.
- [10] T. Ojala, M. Pietikainen, and T. Maenpää, "Multiresolution gray-scale and rotation invariant texture classification with local binary patterns," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, Jul 2002.
- [11] S. Paisitkriangkrai, C. Shen, and A. van den Hengel, "Pedestrian Detection with Spatially Pooled Features and Structured Ensemble Learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2015.
- [12] R. Girshick, F. Iandola, T. Darrell, and J. Malik, "Deformable Part Models are Convolutional Neural Networks," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [13] P. Dollár, Z. Tu, P. Perona, and S. Belongie, "Integral Channel Features," in *British Machine Vision Conference (BMVC)*, 2009.
- [14] S. Zhang, R. Benenson, and B. Schiele, "Filtered channel features for pedestrian detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015.
- [15] R. Benenson, M. Mathias, R. Timofte, and L. Van Gool, "Pedestrian detection at 100 frames per second," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2012.
- [16] P. Felzenszwalb, D. McAllester, and D. Ramanan, "A discriminatively trained, multiscale, deformable part model," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2008.
- [17] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional Architecture for Fast Feature Embedding," *ACM international conference on Multimedia*, 2014.
- [18] S. Rogers and N. Papanikolopoulos, "Counting bicycles using computer vision," in *IEEE Conference on Intelligent Transportation Systems*, 2000.
- [19] Z. Qui, D. Yao, Y. Zhang, D. Ma, and X. Liu, "The study of the detection of pedestrian and bicycle using image processing," in *IEEE Conference on Intelligent Transportation Systems*, Oct 2003.
- [20] H. Cho, P. Rybski, and W. Zhang, "Vision-based Bicycle Detection and Tracking using a Deformable Part Model and an EKF Algorithm," in *IEEE Conference on Intelligent Transportation Systems*, July 2010.

- [21] T. Li, X. Cao, and Y. Xu, "An effective crossing cyclist detection on a moving vehicle," in *World Congress on Intelligent Control and Automation (WCICA)*, July 2010.
- [22] K. Takahashi, Y. Kuriya, and T. Morie, "Bicycle detection using pedaling movement by spatiotemporal gabor filtering," in *TENCON 2010 - IEEE Region 10 Conference*, Nov 2010, pp. 918–922.
- [23] A. Gonzalez, G. Villalonga, J. Xu, D. Vazquez, J. Amores, and A. M. Lopez, "Multiview random forest of local experts combining rgb and lidar data for pedestrian detection," in *IEEE Intelligent Vehicles Symposium (IV)*, June 2015.
- [24] Y. Xiang, W. Choi, Y. Lin, and S. Savarese, "Subcategory-aware convolutional neural networks for object proposals and detection," *arXiv:1604.04693*, 2016. [Online]. Available: <http://arxiv.org/abs/1604.04693>
- [25] W. Tian and M. Lauer, "Fast and Robust Cyclist Detection for Monocular Camera Systems," in *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP)*, March 2015.
- [26] —, "Fast Cyclist Detection by Cascaded Detector and Geometric Constraint," *IEEE Conference on Intelligent Transportation Systems*, Sept 2015.
- [27] W. Zou, X. Wang, M. Sun, and Y. Lin, "Generic Object Detection with Dense Neural Patterns and Regionlets," in *British Machine Vision Conference (BMVC)*, 2014.
- [28] B. Pepik, M. Stark, P. Gehler, and B. Schiele, "Multi-View and 3D Deformable Part Models," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, Nov 2015.
- [29] X. Chen, K. Kundu, Y. Zhu, A. G. Berneshawi, H. Ma, S. Fidler, and R. Urtasun, "3D Object Proposals for Accurate Object Class Detection," in *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [30] F. Yang, W. Choi, and Y. Lin, "Exploit All the Layers: Fast and Accurate CNN Object Detector with Scale Dependent Pooling and Cascaded Rejection Classifiers," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [31] W. Xiong, B. Du, L. Zhang, R. Hu, W. Bian, J. Shen, and D. Tao, "R2fp: Rich and robust feature pooling for mining visual data," in *2015 IEEE International Conference on Data Mining (ICDM)*, Nov 2015.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Spatial pyramid pooling in deep convolutional networks for visual recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, Sept 2015.
- [33] P. Sudowe and B. Leibe, "Efficient Use of Geometric Constraints for Sliding-Window Object Detection in Video," in *Computer Vision Systems*. Springer Berlin Heidelberg, 2011, vol. 6962, pp. 11–20.
- [34] R. N. Rajaram, E. Ohn-Bar, and M. M. Trivedi, "An Exploration of Why and When Pedestrian Detection Fails," in *IEEE Conference on Intelligent Transportation Systems*, Sept 2015.
- [35] X. Chen, K. Kundu, Z. Zhang, H. Ma, S. Fidler, and R. Urtasun, "Monocular 3D Object Detection for Autonomous Driving," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [36] D. Z. Wang and I. Posner, "Voting for Voting in Online Point Cloud Object Detection," in *Robotics: Science and Systems*, July 2015.
- [37] P. Felzenszwalb, R. Girshick, D. McAllester, and D. Ramanan, "Object Detection with Discriminatively Trained Part-Based Models," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 32, pp. 1627–1645, Sept 2010.
- [38] J. J. Yebes, L. M. Bergasa, R. Arroyo, and A. Lzaro, "Supervised learning and evaluation of KITTI's cars detector with DPM," in *IEEE Intelligent Vehicles Symposium (IV)*, June 2014.
- [39] J. Behley, V. Steinhage, and A. B. Cremers, "Laser-based segment classification using a mixture of bag-of-words," in *IEEE Conference on Intelligent Robots and Systems*, Nov 2013.