

OPTIMIZED HOG FOR ON-ROAD VIDEO BASED VEHICLE VERIFICATION

Gonzalo Ballesteros¹, Luis Salgado^{1,2}

¹Visual Processing and Understanding Lab, Universidad Autónoma de Madrid

²Grupo de Tratamiento de Imágenes, Universidad Politécnica de Madrid

ABSTRACT

Vision-based object detection from a moving platform becomes particularly challenging in the field of advanced driver assistance systems (ADAS). In this context, on-board vision-based vehicle verification strategies become critical, facing challenges derived from the variability of vehicles appearance, illumination, and vehicle speed.

In this paper, an optimized HOG configuration for on-board vehicle verification is proposed which not only considers its spatial and orientation resolution, but descriptor processing strategies and classification. An in-depth analysis of the optimal settings for HOG for on-board vehicle verification is presented, in the context of SVM classification with different kernels. In contrast to many existing approaches, the evaluation is realized in a public and heterogeneous database of vehicle and non-vehicle images in different areas of the road, rendering excellent verification rates that outperform other similar approaches in the literature.

Index Terms— HOG, feature extraction, feature classification, video-based vehicle verification, O-HOG.

1. INTRODUCTION

Regarding in-vehicle video-based systems to identify the possibility of collision in the context of ADAS, the detection of vehicles is presented as the fundamental element in the processing chain. The methodology mostly used for on-board video-based detection of vehicles, so that it can perform in real time, consists of two stages. First, the whole image is analyzed to identify regions potentially containing vehicles. For this vehicle hypothesis generation stage, knowledge-based methods keep track of specific features present in the vehicle and its environment [1][2][3], while others resort to stereovision [4] or motion analysis [5]. Second, the hypothesized candidates are typically verified using features relating to their appearance. Template or model-based verification methods have been used, although recently learning-based methods are being adopted, in which the characteristics of vehicles are learned from a training set, and new candidates are dynamically classified according to the learned patterns.

Learning-based hypothesis verification is typically approached as a two-class classification problem: feature vector is extracted from the image, and the sample is classified as vehicle or non-vehicle. The election of the feature space is critical to achieve a successful classification. Explicit features such as symmetry, edges or

shadows, although intuitive, fast to compute and classify, show limited discriminating capabilities that bound the accuracy of the results. On the contrary, implicit features do not make assumptions about vehicles appearance or shape, demanding higher computational resources but rendering significantly better classification results. Among the most commonly used implicit features [6][7], Histograms of Oriented Gradients (HOG) stand out for their excellent results in the object verification.

Dalal and Triggs [8] proposed HOG for human detection obtaining excellent results, and explore their application to different objects [9]. In recent studies, HOG is gaining attention for vehicle detection and classification, being used on images taken under different camera settings. In [10] vehicle top views taken from low-attitude air bone videos are considered, while most works consider more lateral and frontal views, either from wide area monitoring imaging [11] [15] or from on-board-like camera views [13][14][16][18]. Proposals range from straightforward applications of standard HOG and its variations, customized to the particular applications in combination with different classifiers, to combinations with other features as Haar-like [11][18] to improve detection performance. The reduction of the computation requirements for HOG-based drives many contributions, as in [12], where the use of masks adapted to the vehicle shape is proposed to speed up classification. Other works use an AdaBoost classifier on HOG features, but as in [10] they only deal with vehicle top views, or in [13] where only rear views of vehicles in front of the own vehicle (in the same lane) are considered. In other works such as [14], larger databases are considered, proposing cost-effective HOG descriptors where computational cost is alleviated considering only horizontal or vertical cells.

In almost all existing studies that use HOG, the configuration of the descriptor is not fully explored. Direct assumptions of the parameterization for people and other objects detection are applied. Moreover, the variability of the camera settings and databases used bound the impact of the solutions proposed as no direct comparison is possible. Particularly, except for [14], databases are mostly ad-hoc generated, non-public or very limited both, in the number of positive and negative samples and their heterogeneity. HOG configurations in terms of spatial and binning resolution have been analyzed [14]. However, the impact analysis of choices regarding the descriptor computation and the classification strategy is still required. Therefore, extensive quantitative experiments are lacking on the effectiveness of HOG descriptors for vehicle verification.

In this paper, an optimal configuration for HOG is proposed which not only considers spatial and orientation resolution, but descriptor processing strategies and classification. A deep analysis of the optimal settings for HOG for on-board vehicle verification is presented, in the

² This work has been partially supported by the Ministerio de Economía y Competitividad of the Spanish Government under project TEC2010-20412 (Enhanced 3DTV).




Fig. 1. Feature extraction and classification system based on HOG.

context of SVM classification with different kernels. In contrast to many existing approaches, the evaluation is realized in a large and heterogeneous database of vehicle and non-vehicle images in different areas of the road. Results demonstrate a very significant improvement of the verification accuracy, outperforming other strategies in the state of the art [14].

2. HOG DESCRIPTOR

HOGs [8] result from the computation of local histograms of the orientation of the image gradients in a grid. The underlying idea is that objects appearance can be characterized by the local distribution of its edges orientation. A diagram of the HOG feature extraction and classification strategy is shown in Figure 1.

Gradients computation is the first step to generate the HOG descriptor: the application of the selected gradient operator results on edge intensity and orientation values for each pixel. The image is divided into regions called cells, as shown in Figure 2 (a), and the second step involves creating the cell histograms. Orientation binning is applied using an even number of intervals that span the range $[0, \pi]$ or $[0, 2\pi]$. Each pixel in the cell contributes with a vote to the two closest histogram channels, weighted according to the gradient magnitude and the distance to the channel centers. The last step accounts for changes in illumination and contrast through locally normalizing gradient responses. As proposed in [8], cells are grouped into larger structures named blocks as shown in Figure 2 (a). For each block, the non-normalized vector holding all histograms of its cells is normalized using any standard norm such as the L1-norm, L2-norm, etc. Block overlap is also suggested to make this step more robust. The final HOG descriptor is the vector resulting from the concatenation of the normalized vectors of all the blocks.

Rectangular (R-HOG) and circular (C-HOG) cell configurations have been both proposed (Figure 2). However, R-HOG is the most widely used for vehicles as it naturally adapts to the dominant vehicle geometry. Thus it is the one considered in this study.

3. VEHICLE VERIFICATION USING STANDARD HOG

In order to classify the input samples into vehicle or non-vehicle, discriminative approaches largely outnumber generative methods in the literature. Indeed, discriminative models usually give better performance on discriminative tasks than generative models. Among the discriminative approaches, Support Vector Machines (SVMs) have been extensively proposed for HOG-based vehicle verification, providing good generalization [21] and rendering excellent results. Therefore, SVMs is the learning method adopted in this work.

Classification Methodology

The evaluation is carried out in the most extensive public database for vehicle verification proposed in the literature [17]. The GTI vehicle database is a complete data set with 4000 vehicles and 4000 non-vehicles images of size 64×64 pixels. They have been taken from videos acquired with on-board forward looking cameras under different illumination and weather conditions with variable quality. Samples include partial views of vehicle rears, views misaligned with the image, shifted horizontally and/or vertically. Therefore, training and evaluation can be carried out considering a large variability of situations

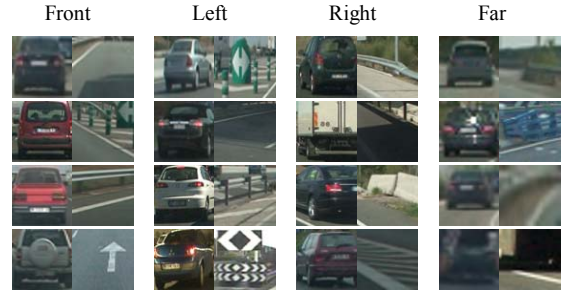


Fig. 3. GTI Database: vehicle (left) and non-vehicle (right) samples for front, left, right and far views.

typical from on-line hypothesis generation systems. Additionally, samples are further classified into four sets according to the relative distance and position to the camera: left, front and right in the close-middle distance, and far distance. This allows exploring classification taken into account the relative pose of the vehicles. Figure 3 shows example images from this database.




Fig. 2. Cell configurations: rectangular (a) circular (b) HOG.

Experiments are carried out based on 5-fold 50% cross-validation methodology, and results are evaluated as the average percentage of correctly classified samples.

Experiments on HOG

When using HOG, it is necessary to define the descriptor parameterization. Experiments have been performed to obtain the most suitable set of parameters and thus get the best possible classification accuracy following some of the recommendations in [8]. These configurations, referred in

some works as standard HOG (S-HOG), are explored evaluating the impact of different parameter settings.

Configuration parameters include the number of cells, $\eta_1 \times \eta_2$, block size, b , and number of orientation bins, β . Square cells are considered, as they are the most used in the literature and adapt naturally to the square images size (64×64 pixels). Therefore, $\eta_1 = \eta_2 = \eta$ and $\eta = \{2, 4, 8, 16\}$ that correspond to $\{4, 16, 64, 256\}$ cells. The number of orientation bins explored are $\beta = \{8, 12, 16, 32\}$.

For the gradients extraction phase, as suggested in [8], the mask used is $[-1 \ 0 \ 1]$ and the range of angles is $[0, \pi]$ so the sign of gradients is ignored. The L2-norm is used to

Front	$\beta=8$	$\beta=12$	$\beta=16$	$\beta=32$
$\eta=2$	98.64	95.36	98.30	92.78
$\eta=4$	99.48	98.52	99.18	98.28
$\eta=8$	98.68	98.54	99.16	98.40
$\eta=16$	98.62	98.80	98.62	98.76

Left	$\beta=8$	$\beta=12$	$\beta=16$	$\beta=32$
$\eta=2$	93.28	91.68	94.14	90.04
$\eta=4$	97.64	96.96	97.04	95.70
$\eta=8$	97.34	96.56	96.96	96.04
$\eta=16$	96.46	96.26	96.02	95.94

Right	$\beta=8$	$\beta=12$	$\beta=16$	$\beta=32$
$\eta=2$	92.60	90.22	93.40	88.66
$\eta=4$	96.00	95.46	96.22	95.10
$\eta=8$	94.70	95.12	94.54	94.46
$\eta=16$	93.96	95.36	93.72	95.64

Far	$\beta=8$	$\beta=12$	$\beta=16$	$\beta=32$
$\eta=2$	92.10	87.12	91.92	83.28
$\eta=4$	97.76	96.54	97.04	95.56
$\eta=8$	97.28	96.70	96.44	96.18
$\eta=16$	97.66	97.62	96.98	97.54

Table 1. Accuracy of Standard HOG.

normalize square blocks of side $b=2$, (2×2 cells). Finally, to classify the HOG vectors, a linear SVM is trained.

Experiments were conducted training individual classifiers for each relative pose of the vehicles. The classification rates (%) are presented in Table 1 where the highest scores are highlighted in bold. We can draw some conclusions from these results, starting first with the close-middle front view. As expected, this is the view that gets the best results in classification. Hypothesis generated for vehicles located in front of the own vehicle show well defined and quite stable geometrical patterns that adapt perfectly to the HOG topology. Analyzing the cells' granularity, $\eta=2$, *i.e.* four cells, results in unstable and poor classification results. However with $\eta=4$ we get the same or even better results than for higher granularities. This means that 16 cells get sufficient detail of the edges orientation distribution for optimal classification: a higher number of cells do not add relevant information for classification, although that information overload does not significantly underperform. Regarding the number of orientation bins, it can also be seen that with the lowest number, $\beta=8$, is enough for a good rating, reaching the maxima of 99.48%, and that higher values do not guarantee better results. For the right and left views, $\eta=2$ underperforms dramatically. The best results, significantly lower than for those for the front view are obtained for $\eta=4$, and there is a clear trend to lose accuracy for higher

number of cells. Lateral vehicle views do show much higher edge pattern variability than vehicle rears, and also the non-vehicle samples in those areas show richer and similarly oriented edge patterns to those of the vehicles (particularly noticeable for right views). Increasing the granularity in cells or orientation bins cause that for some vehicle images, their geometrical structures are diluted in a long descriptor, becoming very similar to that of images where there are no cars but also have a high level of detail.

Finally, far rear views are taken from vehicles at longer distances from the camera, thus showing highly softened details and lower quality than those in the close-middle range. Gradients extraction is not optimal and affects classification, but less than may be expected: these far distance views show very similar structure and stability to that in the close-middle front views that suit perfectly R-

$\eta=4$	[-1 0 1]			Sobel		
	$\beta=8$	$\beta=12$	$\beta=16$	$\beta=8$	$\beta=12$	$\beta=16$
Front	99.48	98.52	99.18	98.28	98.54	99.28
Left	97.64	96.96	97.04	95.70	96.22	96.66
Right	96.00	95.46	96.22	95.10	95.36	95.38
Far	97.76	96.54	97.04	95.56	96.34	97.02
$\eta=8$						
Front	98.68	98.54	99.16	98.68	98.34	99.00
Left	97.34	96.56	96.96	96.74	96.04	96.24
Right	94.70	95.12	94.54	94.78	94.76	93.94
Far	97.28	96.70	96.44	96.92	96.48	96.56

Table 2. Verification rate using different detectors.

HOGs. Although the HOG descriptor provides very good results in view-dependent vehicle verification, its potentiality is not fully exploited in the explored standard configuration. In the next section, how to improve the descriptor and classification modifying some of its parameters is studied, trying to find the optimum configuration for vehicle verification.

4. OPTIMIZING HOG BASED VEHICLE VERIFICATION

Classification performance not only depends on HOG configuration, but on the classifier and strategies used to compute the descriptor. Different alternatives are here evaluated following the steps described in Figure 1.

A. Feature extraction phase

1) Gradients extraction

In several works [14], Sobel detector is proposed for vehicle detection and verification. Results shown in Table 2 demonstrate that, overall, it underperforms, showing a trend to narrow the differences with $[-1 \ 0 \ 1]$ when the binning resolution increases. The mask used in this study demonstrates more stability in the results and significantly higher accuracy rates for lower parameter values, thus resulting in shorter but more efficient descriptors.

Expanding the range of angles is now evaluated. For the application to vehicles [9], considering the sign is suggested, although it also significantly impacts the descriptor size and therefore, the cost of classification. This motivates in some works such as [14] the decision to keep the original positive range as good results are still obtained.

	[0, π]			[- π , π]		
	$\beta=8$	$\beta=12$	$\beta=16$	$\beta=8$	$\beta=12$	$\beta=16$
$\eta=4$						
Front	99.48	98.52	99.18	99.00	99.32	99.28
Left	97.64	96.96	97.04	98.72	98.82	99.04
Right	96.00	95.46	96.22	98.14	98.34	98.36
Far	97.76	96.54	97.04	97.38	98.72	98.50
$\eta=8$						
Front	98.68	98.54	99.16	99.20	99.52	99.76
Left	97.34	96.56	96.96	98.70	99.00	99.10
Right	94.70	95.12	94.54	98.32	97.98	98.32
Far	97.28	96.70	96.44	97.54	98.34	98.46

Table 3. Verification rate extending the orientation range.

Table 3 shows the comparative results between using $[0, \pi]$ and $[-\pi, \pi]$ ranges. The impact of sign is very significant. In almost all configurations, considering $[-\pi, \pi]$ renders better results for the same number of bins, being particularly relevant in the most critical right and far views. The improved results obtained for the same number of bins when sign is considered suggests that, for vehicles, sign information may be even more relevant than orientation resolution. Although the trend has now changed and better performance is generally obtained for higher values of β , sign information should prevail to orientation resolution if the descriptor size is an issue for the application.

2) Normalization

In [8] or [14] the use of overlapping blocks of 2×2 cells is proposed to account for changes in illumination.

	2x2			1x1		
	$\beta=8$	$\beta=12$	$\beta=16$	$\beta=8$	$\beta=12$	$\beta=16$
$\eta=4$						
Front	99.00	99.32	99.28	99.30	99.54	99.60
Left	98.72	98.82	99.04	98.60	99.04	99.18
Right	98.14	98.34	98.36	97.90	98.36	98.50
Far	97.38	98.72	98.50	98.04	98.40	98.32
$\eta=8$						
Front	99.20	99.52	99.76	99.50	99.56	99.60
Left	98.70	99.00	99.10	98.50	98.86	98.80
Right	98.32	97.98	98.32	98.24	98.34	98.40
Far	97.54	98.34	98.46	97.54	98.06	98.20

Table 4. Verification rate using different block sizes.

Alternatives are here evaluated, going down to blocks of 1 cell, i.e. take only into account changes in illumination between pixels of the same cell. As shown in Table 4, the performance of 1×1 blocks is very similar to 2×2 overlapped blocks. However, the former renders more stable values with shorter descriptors, reducing the cost of the descriptors computation and that of classification.

Regarding the norm used, experiments (Table 5) confirm L2 outperforming L1 (suggested in [9]) for all situations. The classification accuracy degrades here particularly for the far and right views reaching in some combinations around a 1-2% loss.

B. Classification phase

A linear SVM is used as baseline for classification in this study. However, higher order polynomials Kernels have been also evaluated. As shown in Table 6, the degree-2 polynomial kernel outperforms for all categories and descriptor configurations. The highest gains correspond to the critical right view in all (η, β) configurations but (4, 16),

	L2			L1		
	$\beta=8$	$\beta=12$	$\beta=16$	$\beta=8$	$\beta=12$	$\beta=16$
$\eta=4$						
Front	99.30	99.54	99.60	99.26	99.40	99.50
Left	98.60	99.04	99.18	98.42	98.84	98.62
Right	97.90	98.36	98.50	97.50	97.70	97.48
Far	98.04	98.40	98.32	96.92	97.94	98.24
$\eta=8$						
Front	99.50	99.56	99.60	98.78	99.26	99.52
Left	98.50	98.86	98.80	98.40	98.76	98.78
Right	98.24	98.34	98.40	97.86	98.26	97.82
Far	97.54	98.06	98.20	97.54	95.44	98.20

Table 5. Verification rate for different norms.

followed closely by the left and far ones. The higher flexibility of the degree-2 polynomial kernel adapts to the views showing lower discrimination, thus rendering a more balanced and stable verification rates in all categories.

	linear			degree-2		
	$\beta=8$	$\beta=12$	$\beta=16$	$\beta=8$	$\beta=12$	$\beta=16$
$\eta=4$						
Front	99.30	99.54	99.60	99.82	99.72	99.80
Left	98.60	99.04	99.18	99.18	99.48	99.64
Right	97.90	98.36	98.50	98.58	98.88	98.78
Far	98.04	98.40	98.32	98.16	98.84	98.80
$\eta=8$						
Front	99.50	99.56	99.60	99.68	99.68	99.58
Left	98.50	98.86	98.80	99.14	99.48	99.40
Right	98.24	98.34	98.40	98.94	99.04	99.28
Far	97.54	98.06	98.20	97.92	98.66	98.44

Table 6. Verification rate with SVM kernels.

5. DISCUSSION

The proper configuration of the HOG descriptor and the classifier used demonstrates to be fundamental for on-board vehicle verification. Sign information is critical for classification. Substantially higher scores than unsigned HOGs are obtained even for the same binning resolution, being particularly relevant for the complex lateral vehicle views. The simple $[-1, 0, 1]$ mask for gradients computation render better results than others proposed in the literature at a lower computational cost. L2 normalization outperforms, but regarding overlapped multi-cell approaches (block based), their cost in terms of descriptor length and computation barely justifies the marginal classification improvement achieved only for some descriptor combinations. Therefore the proposed descriptor computation choices are: $[-1, 0, 1]$ filter for gradient extraction; $[-\pi, \pi]$ as orientation range, and individual L2-norm cell normalization (blocks of 1 cell).

Regarding the number of cells (Table 6), $\eta=4$ provides an excellent performance regardless on the kernel used, outperforming for all vehicle poses but for some (η, β) configurations in the right view. Particularly in this case, higher spatial and binning granularities are required to well overpass 99%. Moving from $\beta=8$ to 12 is justified for the lateral and far views as it improves very significantly the results regardless on the kernel used. Going further benefits verification for vehicles in lateral views, worthy to be considered especially for the right one.

Considering average classification values, i.e. computing for each β the average verification rate of the four classifiers, $\eta=4$ outperforms $\eta=8$ for all values of β ,

confirming it as the best choice for vehicle verification. The best average results are 99.25% for $(\eta, \beta)=(4, 16)$ and a 99.23% for $(\eta, \beta)=(4, 12)$, confirming that the degree-2 polynomial adapts better to vehicle verification. Moreover, one of the advantages of considering independent classifiers for each view is the possibility to adapt the (η, β) operational point to that offering the best and more stable results. In Table 7, verification rates (VR) for the optimal configuration (O-HOG) for each view are summarized.

	S-HOG [14]			O-HOG		
	(η, β)	VR(%)	NCe	(η, β)	VR(%)	NCe
<i>Front</i>	(16,8)	99.18	7200	(4,8)	99.82	128
<i>Left</i>	(8,12)	98.32	2352	(4,16)	99.64	256
<i>Right</i>	(8,12)	97.44	2352	(8,16)	99.28	1024
<i>Far</i>	(8,12)	98.40	2352	(4,12)	98.84	192
	V-HOG [14]			O-HOG		
<i>Front</i>	(4,16)	97.68	64	(2,16)	99.44	64
<i>Left</i>	(4,36)	97.02	144	(2,16)	98.56	64
<i>Right</i>	(4,16)	95.54	64	(2,16)	98.64	64
<i>Far</i>	(4,12)	95.60	48	(2,16)	97.80	64

Table 7. HOG, V-HOG and O-HOG.

Recently, in [14], a new HOG descriptor, V-HOG, was proposed which only considers vertical cells to achieve high computational efficiency. In [14], a throughout study on the standard HOG (S-HOG) configuration is performed; the best results are presented in Table 7 (S-HOG [14]), together with the descriptors size for each view (NCe column). As it can be observed, the O-HOG here proposed, with much shorter descriptors, outperforms S-HOG in all categories. In average, the number of components required for O-HOG is about a 10% of S-HOG (400 vs 3564 components). In the same table, the results of the V-HOG proposed in [14] are compared with O-HOG configurations requiring similar number of components. O-HOG largely outperform V-HOG with gains ranging between 2-3% for all categories. Compared with the results in [20], where average detection rates of 92.2% are achieved, an average 6% gain is obtained using the efficient O-HOG configuration (64 components).

6. CONCLUSION

In this paper, the application of HOG for video-based vehicle verification is explored. The impact of different processing approaches and descriptor configurations is in depth analyzed, proposing an optimized HOG that provides excellent results reducing the computational cost with respect to standard HOG. By selecting an appropriate classifier, results are further improved. Nonlinear kernels on SVMs are the most suitable choice, with the degree-2 polynomial rendering the highest verification rates. Training independent classifiers for the different views allows proposing the combination of spatial and orientation resolution that better adapt to each view, yielding excellent results that outperform other approaches in the literature.

REFERENCES

- [1] J. Hwang, K. Huh, and D. Lee. Vision based vehicle detection and tracking algorithm design. *Optical Engineering*, 48(12):127201–127201–10, 2009.
- [2] G. Y. Song, K. Y. Lee, and J W. Lee. Vehicle detection by edge-based candidate generation and appearance-based classification. In *Intelligent Vehicles Symposium*, 2008 IEEE, pages 428–433, 2008.
- [3] T. Schamm, C. Carlowitz, and J.M. Zollner. On-road vehicle detection during dusk and at night. *IEEE Intelligent Vehicles Symposium (IV)*, 2010., pp. 418–423, 2010.
- [4] F. Oniga and S. Nedeveschi. Processing dense stereo data using elevation maps: Road surface, traffic isle, and obstacle detection. *Vehicular Technology, IEEE Transactions on*, 59(3):1172–1182, 2010.
- [5] F. Woelk and R. Koch. Robust monocular detection of independent motion by a moving observer. In *Proceedings of the 26th DAGM Symposium*, ser. *Lecture Notes in Computer Science*, pp. 27–35. Springer, 2004.
- [6] J. Zhou, D. Gao, and D. Zhang. Moving vehicle detection for automatic traffic monitoring. *IEEE Transactions on Vehicular Technology*, 56(1):51–59, Jan 2007.
- [7] Z. Sun, G. Bebis, and R. Miller. On-road vehicle detection using evolutionary Gabor filter optimization. *Intelligent Transportation Systems, IEEE Transactions on*, 6(2):125–137, 2005.
- [8] N. Dalal and B. Triggs. Histograms of oriented gradients for man detection. In *Computer Vision and Pattern Recognition, CVPR. IEEE Conference on*, vol. 1, pp. 886–893, 2005.
- [9] N. Dalal and B. Triggs. Object detection using histograms of oriented gradients. *Pascal VOC Workshop. ECCV*, 2006.
- [10] X. Cao, C. Wu, P. Yan, and X. Li. Linear svm classification using boosting hog features for vehicle detection in low-altitude airborne videos. In *Image Processing (ICIP), IEEE International Conference on*, pp. 2421–2424, 2011.
- [11] P. Liang, G. Teodoro, H. Ling, E. Blasch, G. Chen, and L. Bai. Multiple kernel learning for vehicle detection in wide area motion imagery. In *Information Fusion (FUSION), International Conference on*, pages 1629–1636, 2012.
- [12] Y. Lv, B. Yao, Y. Wang, and S.-C. Zhu. Reconfigurable templates for robust vehicle detection and classification. In *Applications of Computer Vision (WACV), IEEE Workshop on*, pages 321–328, 2012.
- [13] B. Southall, M. Bansal, and J. Eledath. Real-time vehicle detection for highway driving. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 541–548, 2009.
- [14] J. Arróspide, L. Salgado, and M. Camplani. Image-based on-road vehicle detection using cost-effective histograms of oriented gradients. *Journal of Visual Communication and Image Representation*, 24(7):1182 – 1190, 2013.
- [15] Z. Chen and T. Ellis. Multi-shape Descriptor Vehicle Classification for Urban Traffic, pages 456–461. Institute of Electrical and Electronics Engineers, Dec 2011.
- [16] P. Rybski, D. Huber, D. Morris, R. Hoffman. Visual classification of coarse vehicle orientation using histogram of oriented gradients features. *IEEE Intelligent Vehicles Symposium (IV)*. 921-928. 2010
- [17] GTI vehicle image database, image processing group at upm, <<http://www.gti.ssr.upm.es/data>>, 2011.
- [18] P. Negri, X. Clady, S. M. Hanif, and L. Prevost. A cascade of boosted generative and discriminative classifiers for vehicle detection. *EURASIP Journal on Advances in Signal Processing*, 2008(1):782432.
- [19] J. Jiang and H. Xiong. Fast Pedestrian Detection Based on HOG-PCA and Gentle AdaBoost, pages 1819–1822. Institute of Electrical and Electronics Engineers, Aug 2012.
- [20] D. Balcones, D.F. Llorca, M.A. Sotelo, M. Gavilán, S. Álvarez, I. Parra, and M. Ocaña. Real-Time Vision-Based Vehicle Detection for Rear-End Collision Mitigation Systems. *EUROCAST*, vol. 5717, pp. 320–325, 2009.
- [21] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer Science+Business Media, 2006.