for Pedestrian Recognition in Urban Traffic

Clasificador basado en Máquina de Vectores de Soporte Ajustada para Reconocimiento de Peatones en Tráfico Urbano

Henry Roncancio
Researcher
USP – EESC – SEM LabRoM
roncanciovl@usp.br

André Carmona Hernandes
Researcher
USP – EESC – SEM LabRoM
Andre.hernandes@usp.br

John Faber Archila
Researcher
USP – EESC – LabRoM - UIS
John.faber@sc.usp.br
jfarchid@uis.edu.co

Marcelo Becker
Professor - Researcher
USP – EESC – SEM LabRoM
Becker@sc.usp.br

Resumen

La búsqueda de autonomía en desplazamiento automotriz está en auge. En dicha búsqueda existen diversos problemas a resolver; uno de ellos es el reconocimiento de características en ambientes urbanos que permitan a un vehículo tomar decisiones autónomas. Uno de los aspectos más relevantes corresponde al reconocimiento de transeúntes, una tecnología que se espera pueda salvar millones de vidas en accidentes de atropellamientos. En este trabajo de investigación se propone el reconocimiento de transeúntes en ambientes urbanos por medio de un clasificador basado en una Máquina de Vectores de Soporte ajustada; se consideraron hasta 5000 imágenes de la base de datos la INRIA con el fin de entrenar el clasificador y validar su precisión por medio del método de validación cruzada.

Palabras clave: reconocimiento de transeúntes, clasificador SVM ajustado.

Abstract

The need for autonomy and intelligent decision-making in automobile flow is booming. For this purpose there are a number of interesting problems related to recognition of features in urban environments. One of the main relevant aspects in this subject is the recognition of pedestrians, a technology that is expected to save millions of lives avoiding or decreasing the rates of pedestrian run away. In this paper we propose the recognition of pedestrians in urban environments using a classifier based on a Support Vector Machine. We used up to 5000 images from the INRIA database to train the classifier and validate its accuracy through the cross-validation method.

Key words: Pedestrians recognition, Tuned SVM classifier

1. Introduction

Car accidents cause the death of nearly 1.3 million people every year. In addition, injuries due these accidents cost most countries 1-3% of their gross national product. This problem is even worse if we consider the vulnerable road users (i.e. pedestrians, cyclists, and motorized riders). Today, they represent almost half of the number of victims (about 46%) [1].

For example, in some countries of Central and South America, the deaths of vulnerable users become as high as 70% [1]. The last traffic safety report of the Panamerican Health Organization shows that, for example, when it comes to pedestrian fatalities the Northern Andes zone has the biggest in America (43%), Table I shows the percentages for those categories [2].

	Passengers & Drivers	Riders Motorized	Cyclists	Pedestrian	Other
Canada USA	74%	9 %	2 %	12%	2%
Dominican Republic, Cuba, Puerto Rico	51%	12%	1 %	23%	8%
British Virgin Island, Guyana, Suriname, Saint Lucia, Trinidad and Tobago, Bahamas, Jamaica, Barbados	61%	11%	8 %	23%	1%
Paraguay, Brazil, Chile, Argentina, Uruguay	20%	14%	5 %	31%	31%
Mexico, Belize, Costa Rica, Guatemala, Nicaragua, Honduras, Panama, El salvador	49%	6 %	6 %	32%	8%
Venezuela, Peru, Bolivia, Colombia, Ecuador	36%	8 %	2 %	43%	10%
Americas Region	47%	12%	4 %	23%	14%

Table I. Deaths for road user category [2].

According to the World Health Organization (WHO), if this death trend continues, it is predicted that the road traffic injures will be the 5th cause of death in the world by 2030 [3]. The WHO has named this decade: "The Decade of Action for Road Safety", after its plan to change this trend; it is expected that over 5 million lives will be saved (before 2020) just by strengthening the road safe management, making safer roads and safer vehicles.

Regarding vehicles and roads, many efforts are being addressed to design autonomous vehicles able to deal with real traffic conditions. Because most car accidents are attributed to human error [1][4], making not only autonomous vehicles but also intelligent transportation systems will hopefully bring reduction in road fatalities. Even though worldwide, large scale implementation of such technology is still a plan, in the meanwhile the current development has focus on designing pre-emptive systems to help the driver in avoiding collisions.

The Mobile Robotics Laboratory (LabRoM) at São Carlos School of Engineering - University of São Paulo (EESC-USP) is currently developing the SENA Project (Sistema Embarcado de Navegação Autônoma)[5], which provides the vehicle with a combination of various sensors, such as laser sensor, camera, IMU (Inertial Measurement Unit) and GPS units, to perceive the environment that surrounds the vehicle and to control its motion; a schematic of the SENA project is shown in Figure 1.
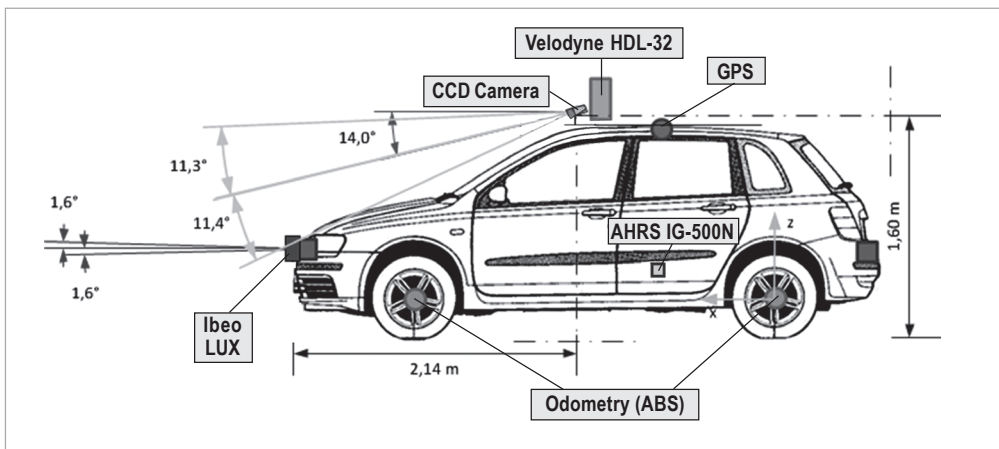

Figure 1. Sensor disposal at SENA project's platform.

One of the main components in this vehicle is the object detection module, namely the pedestrian and car detection module, which uses images captured by a camera and data from laser sensors to sense the outdoor environment. The SENA project is currently in the sensors set-up and calibration phase; for this reason, most data used in this paper was taken from research-purpose databases. One activity of the next stage of the project will be setting up an in-house database.

When it comes to pedestrian detection, there are several approaches in the literature; they are mainly focused on computer vision with either mono or stereo cameras [6], analysis of laser sensor data [7], and even data fusion of both laser and camera [8]. In this paper we present a pedestrian recognition algorithm that uses a classifier based on a computer vision technique. The images used to train and to validate the classifier are from the INRIA database (Institute National de Recherché en Informatique et en Automatique) [9]. In addition, a few images from LabRoM database were also used for testing the algorithm.

The following section shows the methodology used to generate the classifier's model, it also explains the cross-validation method used for model validation and the metrics used for evaluating it. The F_1-score metric is used as a model performance estimator, and in addition, the repeated random sub-sampling validation is used as the cross-validation method, these both procedures show a different approach to the problem in relation to other related papers [9],[6],[10],[11]. Previous papers implemented mainly an SVM with linear kernel and validated the model performance through the recognition rate, i.e. the accuracy of the prediction. In contrast, in this paper we tune an RBF-kernel SVM and take advantage of using the F_1-score metric for improving the tuning, when used simultaneously with a cross-validation procedure. Finally, the classifier performance is discussed through the analysis of tables of results.

2. Methodology

The following steps are carried out in this approach to pedestrian recognition (see Figure 2):
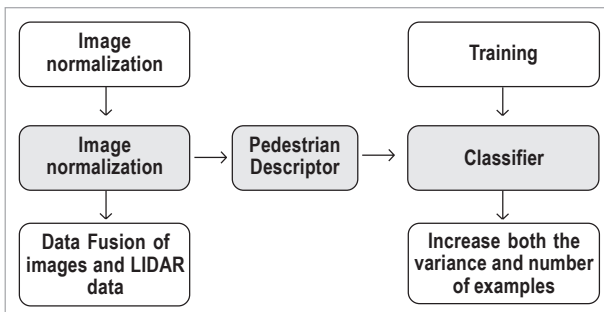

Figure 2. Methodology

1. Image normalization.

2. Extraction of both pedestrian and non-pedestrians descriptors using Histogram of Oriented Gradients (HOG).

3. Creation of a model that represents such descriptors, i.e., training the tuned support vector machine (SVM).

4. Model validation.

2.1. Image Normalization

This normalization is usually done manually. Most of the INRIA images were already normalized manually. However, for future work we propose a method intend to normalize then automatically. The normalization aims to accomplish an image resizing in order to make the pedestrian most of the image, in addition to this; it tries to make the pedestrian the center of the image. A better classifier performance is achieved using normalization, because the positive descriptors become more similar.

2.2. Histogram of Oriented Gradients

This human-shape descriptor created by Dalal [9] is invariant to local geometric and photometric transformations in pedestrian imagery. The HOG is widely used in pedestrian recognition because it presents a better performance when compared with other descriptors [12]. This algorithm extracts local features from the image by calculating the distribution of local intensity gradients or edge directions; it also does a local contrast normalization. In this implementation, the resulting

descriptor for any image size is a vector in $\Re^{81}$ [13]. The descriptor has the class attributes or features that will be subsequently used for training the classifier.

2.3. Metrics for Model Validation

The performance of a binary classifier is quantified by calculating the number of correct, missing and unexpected predictions, as presented in Table II and Equations (1-4).

Table II. Confusion matrix. Relationships between the predicted and actual class.

	Actual Class (expectation)	
Predicted Class (Observation)	TP (True Positive) Correct Result	FP (False Positive) Unexpected result
	F N (False Negative) Missing Result	TN (True Negative) Correct absent of result

$$Accuracy = \frac{True\ Positives + True\ Negatives}{Number\ of\ Samples} \quad (1)$$

$$Precision = \frac{True\ Positives}{No.\ of\ Predicted\ positive} \quad (2)$$

$$Recall = \frac{True\ Positives}{No.\ of\ Actual\ positive} \quad (3)$$

$$F_1\ score = 2 * \frac{Precision * Recall}{Precision + Recall} \quad (4)$$

2.4. Tuned Support Vector Machine

An SVM is an algorithm that makes predictions on new data by creating a model of known data (training data). When the data set is labeled to a class, the SVM becomes a trained classifier. Basically, the SVM does a mathematical optimization to find the model that best separates the classes. This procedure is carried out by minimizing the cost function that represents the error between the data and the model; this cost is shown in (5). As a result, the best hyperplane of separation (hypothesis) is determined. In some classification problems the relation between class labels and features is predominantly non-linear, so that, in order to achieve the separation an especial non-linear function is used to map the examples to a higher dimension, in which the features are more easily separated. The function that defines how the examples are mapped is known as the "kernel", and it can be either non-linear or linear. The radial-based function (RBF) shown in (6) is a non-linear function widely used as kernel. It defines a similarity function between all the examples.

The approach proposed in this paper uses a *tuned* SVM: a search for the best kernel-based SVM parameters that achieve the greatest F_1 score in the prediction is conducted. For the SVM with RBF kernel, the model parameters are C, known as regularization in (5), and σ, which is the RBF width, as in (6) [14]. Having a number of examples many times bigger than the number of features allows the model with RBF kernel to obtain better predictions than the linear kernel. The following steps are carried out to tune the model [15] (see Figure 3):
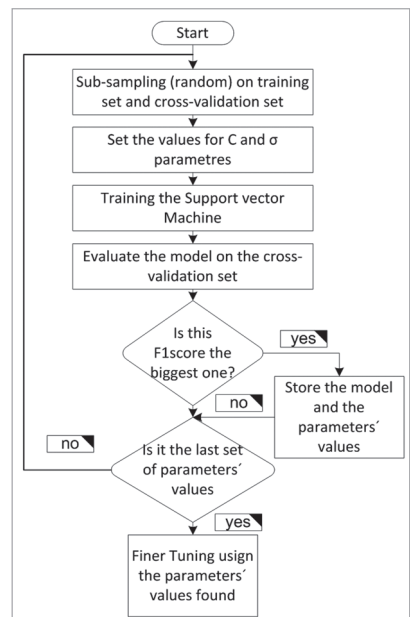

Figure 3. Flowchart of the SVM tuning process.

- *Feature scaling.* In order for the SVM with RBF kernel to avoid features in greater numeric ranges dominating those in smaller ranges, a feature scaling is carried out.

- *Search for the best C and ó in a coarse grid.* Both parameters determine the model variance and bias. High bias means an underfitting model and high variance means an overfitting model. The final model should manage a good trade-off between bias and variance. The RBF is tuned by the parameter σ, small values of ó increase the model variance, the smaller this parameter the more overfitting the model. The C parameter in the cost function regularizes the model by underfitting it; small values decrease the model variance. A good combination of these parameters accomplishes the required trade-off. A grid search with steps 2^x for $x \in [-6,5]$, $x \in Z$, was performed for tuning both parameters, C and σ.

- *Search for the best C and ó in a finer grid.* The finer grid is created by splitting into 10 equally divided parts the region near the parameters found in the previous search. The finer grid becomes different for each tuning value.

$$J(\theta) = \frac{1}{2} \sum_{j=1}^{n} \theta_j^2 + \cdots$$

$$C \sum_{i=1}^{m} [y^{(i)} \, cost_1 \, (\theta^T x^{(i)}) + (1 - y^{(i)}) \, cost_0 \, (\theta^T x^{(i)})] \tag{5}$$

$$f_i = e^{-\frac{||x - l^{(i)}||^2}{2\sigma^2}} \tag{6}$$

Where:

$J(\theta)$ *: Cost function*
f_i *: Similarity function used as Kernel*
θ *: Parameter Vector of the estimated hypothesis*
$l^{(i)}$ *: Landmark based on example i*
$x^{(i)}$ *: Feature vector for example i*
$cost_1$ *: Cost for positive examples*
$cost_0$ *: Cost for negative examples*
n *: Number of features*
m *: Number of examples*

2.5. Cross-Validation Method (CV)

In order for the model to be statistically well validated its performance should be evaluated on "new data" (different from training set) several times, otherwise it may show an overoptimistic result. For this purpose, the data is commonly split into two sets, namely training set and cross-validation set. For one split a simple validation is carried out by calculating an estimator of risk on the cross-validation set, e.g. accuracy, using the model created with the training set. CV consists in averaging several validations corresponding to different-data splits [16]. Even though this method is intended to establish the model statistical performance, in this paper a CV is carried out in order to find the best model for our data. Several models are created by varying both the model parameters and the training set, and its performance is subsequently evaluated in the cross-validation set.

Because the cross-validation set is used to tune the model, there is a need for having another independent data set for further validation. Therefore, the data is split into three sets and the additional set is named test set. A model with high F_1 score, i.e., both precision and recall high, in the cross-validation set is considered a good one. The best model is finally validated on the test set in order to calculate a more realistic estimation. This paper uses this CV evaluation for tuning the classifier´s model just by looking the parameters that make the model reaches the highest F1 score in the prediction.

The type of CV method used is the repeated random sub-sampling validation, i.e., the data is randomly split into training and cross-validation set several times, and for every split both the model accuracy and F_1 score are calculated. The values are averaged in order to get the model statistical performance. Because the model performance depends on not only the model parameter but also on the data, a different performance is achieved for each split. It is expected that a database with both most of the scenarios and many examples generates a better model.

3. Results

3.1. Data

The database used is composed of over 5,000 color images downloaded from INRIA [9]. The images were taken on different weather conditions, such as rain and snow/sleet; besides, with different illumination conditions, including some images taken on night. In addition, the human bodies have different poses (although all are upright), and there are even a few people riding bikes. Several normalized pedestrian images are shown in Figure 4. Due to this normalization, there are two image resolutions in the positive images, namely 96×160 pixels and 70×134 pixels. On the other hand since the negative images were not normalized, they are bigger than the positive ones and even have different sizes.

Figure 4. INRIA´s Normalized positive images [3]

3.2. Preliminar Experiment

The INRIA data set is originally split in two sets, namely train and test set (70% and 30% respectively). Table III shows a few preliminary tests for those original sets. Without normalization, the SVM (linear kernel) predicted with an accuracy greater than the SVM (RBF kernel), despite of the accuracy in the training set is greater using the RBF kernel. The latter indicates that the SVM (RBF kernel) without normalization is over fitting the data.

For normalized images, classification results are better. The SVM (Linear kernel) again obtains a better accuracy; however, the precision is not good enough (71.2%); even though the recall is 100%, the F_1 score (83.2%) is affected by the low precision. In this case, all the actual pedestrians are being predicted, but there are a few non-pedestrian data that is erroneously predicted as pedestrian. These values illustrate that accuracy is not always a good estimator of model perfor-

Table III. Preliminary test using the INRIA two-set original split.

Classifier Settings and Data description	Accuracy in Training Set	Accuracy in Test Set	Recall	Precision	F_1 score
SVM (Linear Kernel) in Images without normalization	85.3%	82.5%			
SVM (RBF Kernel) in Images without normalization	91.9%	74.9%			
SVM (Linear Kernel) in normalized Images	96.0%	91.9%	100%	71.2%	83.2%
Tuned SVM (ó=0.01, C=0.03) in normalized Images	100%	89.7%	85.6%	100%	92.2%
Tuned SVM (ó=5, C=4) in normalized Images	94.6%	89.0%	93.6%	91%	92.3%

mance. A better estimator is the F_1 score, because it takes into account the precision as well as the recall. Accordingly, a better performance is achieved using an RBF kernel because the F_1 score is as high as 92.3%.

Unlike Linear kernel, the RBF kernel obtains better precision rates. However, in some cases the recall is not satisfactory (tuned SVM with σ =0.01, C=0.03). By performing the grid search procedure described above, a model with a better trade-off between precision and recall, such as the tuned SVM with σ =5, C=4 is found (both metrics are higher than 90%).
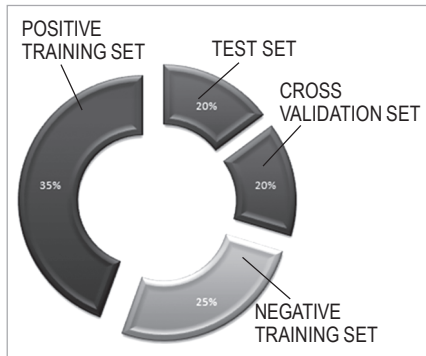

Figure 5. Data Distribution used in CV method.

3.3. CV Experiment

In this experiment, the CV method explained above was used. Data was split into training, cross-validation and test sets, as seen in Figure 5. Positive examples correspond to pedestrian images (70%) and negative ones to city's streets and country images (30%).

Table IV shows the average for both model accuracy and F_1 score for 10 random sub-sampling validations. It can be seen that the performance of the tuned RBF SVM in both metrics is greater than 94% in average in the CV sets, whereas for the test set even higher values are achieved. The model found with the finer grid search tuning also improves the performance rates. We reason that the fact of increasing the variance of the training samples, i.e. adding new different examples, gives the prediction model an improvement on its accuracy, from 89% to 94.7%, and on the F_1 score, from 92.3% to 96.1%, in held-out (unknown) data.

Table IV. Tuned SVM performance evaluation with CV method.						
Classifier Settings (normalized images)	Averaged values in CV		Results in test set using the best model*			
	Accuracy	F_1 Score	Accuracy	Precision	Recall	F_1 Score
Coarse Tuned SVM (*Best model: ó=4, C=16)	94.1%	95.7%	93.9%	97.3%	93.6%	95.4%

In term of processing speed, tuning the SVM model with the finer grid on an image of 640×480 took 1.93s in average; and for an image of 137×70 the time was 0.02s in average. These tests were carried out in MatLab®, with an Intel 2 duo CPU at 3GHz.

3.4 Experiments with images from LabRoM

These experiments were carried out with the fine-tune SVM classifier (σ=4.2, C=11.4) presented in Table IV. Figure 6 shows a positive and a negative image in the same scenario, correctly predicted by the classifier. This result shows that the background had no effect in the classification.

Figure 6. (a) Non-pedestrian image. (b) Pedestrian image.

Figure 7 shows an assorted set of pedestrian images correctly classified. It is interesting to note that classification is correct even for different poses and illumination conditions. These images are very similar to the positive ones used to train the classifier. Nevertheless, other images in which the pedestrian's posture was not upright or he was wearing bags and backpacks, were misclassified (see Figure 8).

In all cases, the images were manually pre-processed in order to get a normalized pedestrian image. However, we are currently developing

Henry Roncancio • André Carmona Hernandes • John Faber Archila • Marcelo Becker

at LabRoM an automatic method for normalizing the images by applying a data fusion technique of CCD camera (AVT stingray 033C) and a frontal, multilayer laser sensor (Ibeo Lux) embedded in our test vehicle. This laser sensor yields the obstacle's location in the front side. By making a correlation between both data, the obstacle's spatial location found through the laser is subsequently used to get from the entire image (taken from the front side) another small image with only the obstacle. As a result, an obstacle's normalized image is obtained.

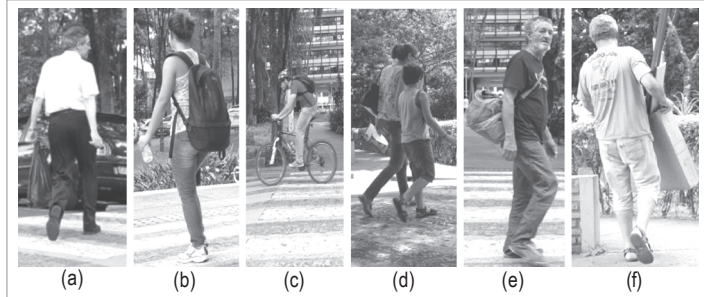**Figure 7.** Pedestrian images classified correctly.

Figure 8. Pedestrians images classified erroneously.

4. Conclusion

The increase in the prediction rates showed in the CV test, indicates that an important factor to improve the classifier performance is to train with a high number of examples, provided that the model presents a good trade-off between bias and variance. On the other hand, the cross-validation procedure used for tuning the model yields a way of avoiding overfitting. Moreover, the model performed well in the test set, which was totally independent. The search for the best F_1 score in the CV also enables the method to find a model with the best trade-off between precision and recall. We conclude that this trade-off is only reached by tuning the model with the F_1 score rather than the accuracy.

Furthermore, tuned RBF-kernel SVM models performed better compared with linear-kernel SVM models, and better prediction rates can be achieved by doing a finer grid-search tuning. The CV test shows that such model is able to obtain an accuracy of 95% and an F_1 score of 96.3% in the pedestrian recognition task, for more than 5,000 images from INRIA dataset.

Even though the classification results were satisfactory for many scenarios, using the HOG and the tuned SVM, there are still pedestrian images that must be included for training in order to make the classifier more robust. Tests showed that in actual road scenarios there are variations with regard to pedestrian's posture, pedestrian's companions, and pedestrian's accessories that introduce confusion factors resulting in misclassifications. It is expected that by adding these scenarios in the training data the classifier robustness is improved. As a final comment we expect that by implementing this kind of computer vision-based recognition system in the vehicle's front bumper, combined with a vehicle control system that is able to avoid collisions, the road fatalities involving pedestrians can be greatly reduced.

Acknowledgment

References

[1] Arlot, Sylvain and Celisse, Alain. A survey of cross-validation procedures for model selection. Statistics Surveys, 4:40—79, 2010.

[2] Bertozzi, M. and Broggi, a. and Rose, M. Del and Felisa, M. and Rakotomamonjy, a. and Suard, F. A Pedestrian Detector Using Histograms of Oriented Gradients and a Support Vector Machine Classifier. 2007 IEEE Intelligent Transportation Systems Conference, :143—148, 2007.

[3] Dalal, Navneet and Triggs, Bill. Histograms of oriented gradients for human detection. Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on, pages 886—893, 2005. Ieee.

[4] Enzweiler, Markus and Gavrila, Dariu M. Monocular pedestrian detection: survey and experiments.. IEEE transactions on pattern analysis and machine intelligence, 31(12):2179—95, 2009.

[5] Haykin, Simon. Neural networks: a comprehensive foundation. Pearson Education, Inc., 1999.

[6] Hsu, C.W. and Chang, C.C. and Lin, C.J. and Others. A practical guide to support vector classification. Technical report, 1, Department of Computer Science National Taiwan University, 2010.

[7] Kidono, Kiyosumi and Miyasaka, Takeo and Watanabe, Akihiro and Naito, Takashi and Miura, Jun. Pedestrian recognition using high-definition LIDAR. Intelligent Vehicles Symposium (IV), 2011 IEEE, number Iv, pages 405—410, 2011. IEEE.

[8] Araújo, R. L. C. and Lacerda, V. M. and Hernandes, A. C. and Mendonça, A. S. B. and Becker, M . Classificação de pedestres usando câmera e sensor LIDAR. In: X Simpósio Brasileiro de Automação Inteligente (SBAI), 2011, São João del Rei, MG. Simpósio Brasileiro de Automação Inteligente (SBAI), 2011.

[9] Ludwig, O. and Delgado, David and Gonçalves, V. and Nunes, Urbano. Trainable classifier-fusion schemes: an application to pedestrian detection. Intelligent Transportation Systems, 2009. ITSC'09. 12th International IEEE Conference on, pages 1—6, 2009. IEEE.

[10] National Highway Traffic Safety Administration of U.S. Goverment. TRAFFIC SAFETY FACTS 2009. Technical report, NHTSA's National Center for Statistics and Analysis, Washington, 2009.

[11] Organización Panamericana de la Salud. Informe sobre el estado de la seguridad vial en la región de las américas. Technical report, Biblioteca Sede OPS, Washinton, 2009.

[12] Serfling, Matthias and Loehlein, Otto and Schweiger, Roland and Dietmayer, Klaus. Camera and imaging radar feature level sensorfusion for night vision pedestrian recognition. 2009 IEEE Intelligent Vehicles Symposium, :597—603, 2009.

[13] Wang, Zhen-Rui and Jia, Yu-Lan and Huang, Hua and Tang, Shu-Ming. Pedestrian Detection Using Boosted HOG Features. 2008 11th International IEEE Conference on Intelligent Transportation Systems, :1155—1160, 2008.

[14] World Health Organization. Decade of Action for Road Safety 2011-2020. Technical report, World Health Organization, Geneva, 2011.

[15] World Health Organization. Global status report on road safety. Technical report, WHO Library Cataloguing-in-Publication Data, Geneva, 2009.

[16] Zhang, Guangyuan and Gao, F. and Liu, C. and Liu, W. and Yuan, H. A pedestrian detection method based on SVM classifier and optimized Histograms of Oriented Gradients feature. Natural Computation (ICNC), 2010 Sixth International Conference on, number Icnc, pages 3257—3260, 2010. IEEE.

Henry Roncancio

Nació en Bogotá, Colombia. Es ingeniero electrónico de la Universidad Distrital, de Bogotá. Actualmente se encuentra haciendo Maestría en Ingeniería Mecánica en la Universidad de São Paulo, Escuela de Ingeniería de São Carlos, São Carlos, Brasil. Pertenece al grupo de investigación LabRoM donde realiza su proyecto de maestría en detección de obstáculos para vehículos autónomos. Se desempeñó como Ingeniero de diseño en Ultra durante 7 años donde desarrolló sistemas electrónicos para vehículos. e-mail: roncanciovl@usp.br

André Carmona Hernandes

Nació en Sao Paulo, Brasil. Es Ingeniero mecatrónico de la Universidad de Sao Paulo, de Sao Carlos, Brasil. Actualmente cursa su Maestría en ingeniería mecánica en la Universidad de Sao Paulo, Escuela de Ingeniería de Sao Carlos, Brasil. Pertenece al grupo de investigación LabRoM donde realiza su proyecto de maestría con enfoque en robótica cognitiva. e-mail: andre.hernades@usp.br

John Faber Archila

Nació en Vélez, Colombia. Es Ingeniero mecánico de la Universidad Industrial de Santander, Bucaramanga, Colombia. Obtuvo su título de Maestría en Ingeniería Mecánica en el área de robótica en la Universidad Federal de Rio de Janeiro Brasil. Actualmente cursa estudios de PhD en la Universidad de Sao Paulo Brasil. Se desempeñó como investigador en el laboratorio de robótica de la UFRJ durante 2 años. Actualmente se desempeña como profesor en el área de diseño mecatrónico en la Universidad Industrial de Santander de Bucaramanga, Colombia, y pertenece como investigador al grupo GIROD y al LabRoM donde realiza su proyecto de doctorado en el área de robótica móvil y robótica Agrícola. e-mail: jfarchid@uis.edu.co john.faber@sc.usp.br

Marcelo Becker

Nació en Ribeirao Preto, Brasil. Es Ingeniero Mecánico de la Universidad de Sao Paulo, Sao Carlos, Brasil. Obtuvo su título de Maestría en Ingeniería Mecánica en la Universidad Estadual de Campinas, Brasil. Obtuvo su PhD en la Universidad Estadual de Campinas, Brasil. Realizo Posdoctorado en la Escuela politécnica federal de Lausanne Suiza, en el Autonomous System Lab. Actualmente se desempeña como profesor en el área de robótica autónoma en la Universidad de Sao Paulo, Escuela de Ingeniería de Sao Carlos, Brasil y es el director e investigador principal del laboratorio de robótica móvil LabRoM. e-mail: becker@sc.usp.br