A Multisensor System for Road Surface Identification

Gonzalo Safont[1,*], Addisson Salazar[1], Alberto Rodríguez[2], Luis Vergara[1]

[1]Institute of Telecommunications and Multimedia Applications, Universitat Politècnica de València, Valencia, Spain

[2] Universidad Miguel Hernández de Elche, Departamento de Ingeniería de Comunicaciones, Elche, Spain

gonsaar@upvnet.upv.es; asalazar@dcom.upv.es

Abstract—**This work introduces a multisensor road surface identification system that considers features from four different kind of sensors: microphones, accelerometers, speed signals, and handwheel signals. Features are extracted separately from each sensor, joined together, and then filtered using feature selection before classification. The proposed system was tested on a set of signals extracted from a specially-converted passenger car driving on a closed course. Three types of road surfaces were considered: smooth flat asphalt, cobblestones, and stripes. Three classifiers were considered: linear discriminant analysis, support vector machines, and random forests. All the considered classifiers reached over 90% accuracy, with a maximum accuracy of 96.52% for RDF. These results show the potential of the proposed system for road surface identification.**

Keywords—classification, road surface identification, feature selection, audio signals, self-driving-vehicles

Type of submission—Short Paper

Symposium—Signal & Image Processing, Computer Vision & Pattern Recognition (CSCI-ISPC)

I. INTRODUCTION

The development of autonomous or semi-autonomous car technology is a promising area that is attracting considerable research effort in recent years. Such technologies could improve traffic flow, assist the driver, and reduce the potential for traffic accidents. One of the key parts in this technology is surface identification. Depending on the type of terrain, adjustments could be made to the car to improve the safety and comfort of the driver and passengers.

Existing works on road surface identification can be broadly classified into three partially-overlapping groups: estimation of road roughness profiles, detection of weather conditions hazardous for driving, and road type detection. These three groups typically have diverging goals. Road roughness profiles are typically studied to obtain cost-effective solutions for supervising and planning road maintenance, such as estimating road profiles from multisensor data [1] and detecting road damage [2]. Weather conditions are typically studied to increase the safety of the driver and passengers, such as detecting wet asphalt [3] or classifying between five different weather conditions [4]. Road type detection studies have goals that depend on the considered classes, such as detecting road traffic [5] or autonomous car technology [6-8]. Our work is comprised in this last category.

Previous works have considered one or more sensors for road type detection. Masino et al. classify five types of pavement using the sound in the tire cavity [5]. They reached 85.3% accuracy using SVM, which they increased to 91.8% by post-processing the outputs of the classifier to consider that road types are usually constant for several meters at least. Bystrov et al. ([6-8]) have considered sonar, ultrasound, and radar to perform classification of the road surface in front of the car, reaching over 90% accuracy in multiple classes (typically, asphalt, gravel, grass and sand).

This work presents a combination of four different sensors for road surface identification: (i) an accelerometer on the intermediate shaft; (ii) three microphones, two directional microphones (left/right) on the driver's head, pointing forward, and one microphone pointing to the upper side of the electric power steering (EPS) system column; (iii) the speed of the left and right wheels of the vehicle; and (iv) the torque and position of the handwheel. The last two sensors are already included in the vehicle's EPS system, while the former were fitted in an specially-converted passenger car.

II. ROAD SURFACE IDENTIFICATION SYSTEM

The proposed surface identification system is shown in Fig. 1. Features are extracted separately from the four considered sensors, joined together, and filtered using feature selection before classification. The output of this classifier
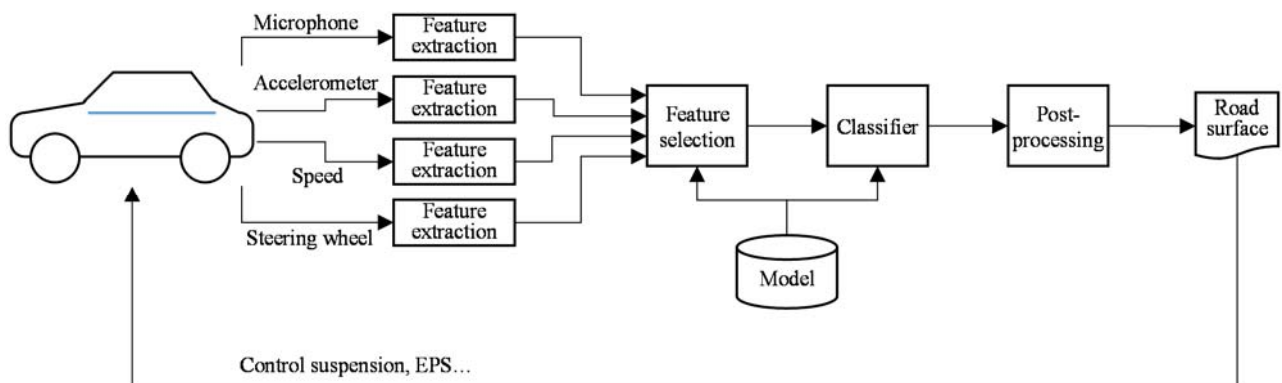

Fig. 1. Diagram of the proposed system.

could then be used as feedback for the car, controlling the suspension and other car systems in order to improve the safety of the driver.

A. Feature extraction and selection

Given that most of the channels are audio channels or vibration channels, we considered the following features that are commonly used for audio classification [9]:

- Average power across all frequency bands
- Centroid frequency
- Maximum frequency
- Spectral contrast: ratio between the minimum and maximum spectral values in each octave.
- Spectral slope: first-order polynomial trend of the power spectrum, assuming that the spectrum follows a power law of the frequency.
- Spectral flatness: ratio of the geometric and arithmetic means of the power spectrum.

We also considered the following high order statistics:

- third-order autocorrelation
- time reversibility

This resulted in a total of 56 features extracted from each of the 10 available channels, i.e., a total of 560 features for classification.

B. Feature selection

Given the high dimensionality of the data, feature selection was performed before classification. For this task, a feature ranking [10] method was chosen. Feature ranking methods do not transform the original features but assign a score or importance to each feature and then rank the features by score. This ranked list can then used to perform feature selection, e.g., by choosing the Q best-ranked features. The rest of the features are never used and, in a trained system, could be omitted from the feature extraction stage.

In the feature ranking method used in this work, the score of the ith feature was computed as the average sub-score of the feature for each class j, $j = 1 \dots K$, with K being the number of classes of the problem. This sub-score was computed as the informedness [11] of a simple classifier fit to the 1-vs-all binary problem (class j versus every other class) using only feature i, where

$$informedness = specificity + recall - 1$$

This indicator is very fast to compute and is robust with respect to differently-scaled variables, outliers, and mislabeling errors.

C. Classifiers

Although the road surface identification could be modeled as a change detection problem [12] followed by classification, in this work, we only considered classification. The following classifiers were selected: linear discriminant analysis (LDA), support vector machines with linear kernel (SVM), and random forests with 50 trees (RDF). These classifiers were chosen because of their widespread application in machine learning problems. Furthermore, while SVM and RDF might be computationally expensive to train, the trained models are fast to evaluate, making them appropriate for the task.

D. Post-processing

In an actual setting, the system will not be required to classify epochs in isolation; more likely, the road surface type will be classified continuously and the type of surface will change, at most, once every few seconds. Thus, the results of the classification were post-processed to consider time dependencies in the result. Thus, any change in road surface that was sustained for less than 1 full second was removed and replaced with the previous road surface. This post-processing produced more stable outputs that considered the temporal dependencies in the data.

III. EXPERIMENT

The data analyzed in this work were obtained using a specially-converted passenger car, driving on a closed course over three different types of surfaces: smooth flat asphalt, cobblestones, and stripes. A total of 63 files were taken, with an average length of 14.59 seconds. Each file corresponded to a different configuration of vehicle speed (within 0 to 30 kph) and road surface. The vehicle drove in a straight line during the capture of the data. All four sensors were sampled at a rate of 48 kHz. Features were extracted for each window of size 1.5 seconds, with an overlap of 1.4 seconds between each pair of consecutive windows, resulting in a total of 8309 samples. One of the captured files is shown in Fig. 2. Around 6.50 seconds, there was a transition from cobblestones to smooth flat asphalt. This transition was noticeable in the microphone channels and the handwheel channels, which are much noisier for cobblestones than they are for smooth flat asphalt.
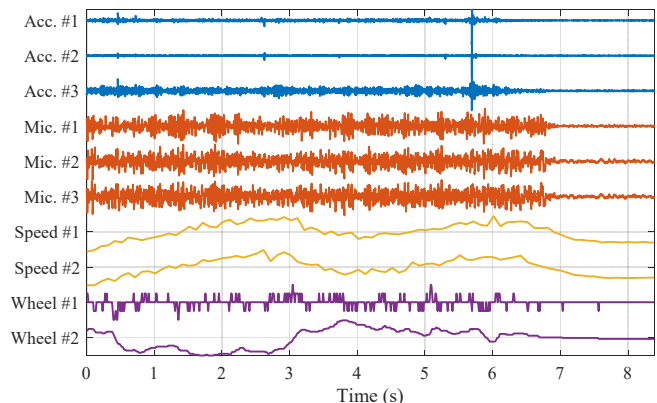

Fig. 2. Example of the captured signals from four sensors: three accelerometers (Acc.), three microphones (Mic.), two speed channels (Speed), and two handwheel signals (Wheel).

The proposed system was tested through a series of Monte Carlo experiments. For each iteration of the experiment, the samples were randomly split 50/50 into training and testing. During these splits, it was ensured that the proportion of samples of each type of road surface was approximately the same for the training and testing sets. To avoid overfitting, samples from the same file could only appear in the training set or in the testing set, but not both. Feature ranking was estimated using the training set, and the number of chosen features was optimized separately for each method using cross-validation: 9 features for LDA, 4 features for SVM, and all features for RDF. An example of the results of the feature ranking process is shown in Fig. 3, which shows the four best-ranking features extracted from the data displayed in Fig. 2.
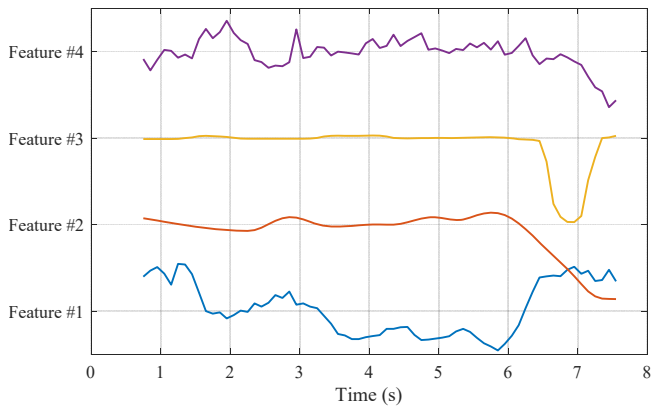
Fig. 3. Four best-ranking features extracted from the data shown in Fig. 2. The transition at 6.50 seconds is marked in all four features.

Finally, after feature selection, the considered classifiers were trained using the training set and tested on the testing set. Performance was estimated as the average of 250 iterations of the Monte Carlo experiment. The average performance of each of the considered classifiers is shown in Table I. These result exceed those of previous works in the literature [5-8]. Out of the considered classifiers, RDF was the method that yielded the best accuracy and stability, reaching over 96% accuracy with the smallest deviation. The improvement of RDF over the other classfiers might have owed to the fact that RDF is essentially a fusion of multiple decision trees, thus allowing for better stability and performance.

TABLE I.   AVERAGE ACCURACY OF THE PROPOSED SYSTEM.

Classifier	LDA	SVM	RDF
Average accuracy (%)	93.87	92.46	96.52
Standard error (%)	0.25	0.31	0.14

An example of the results yielded by the proposed system are shown in Fig. 4. Regardless of the considered classification method, the proposed system yielded a very accurate classification. As seen in Fig. 4, most of the differences between the actual road surface and the estimation took place during the transitions between road surfaces, where the proposed system experienced a slight delay in response.
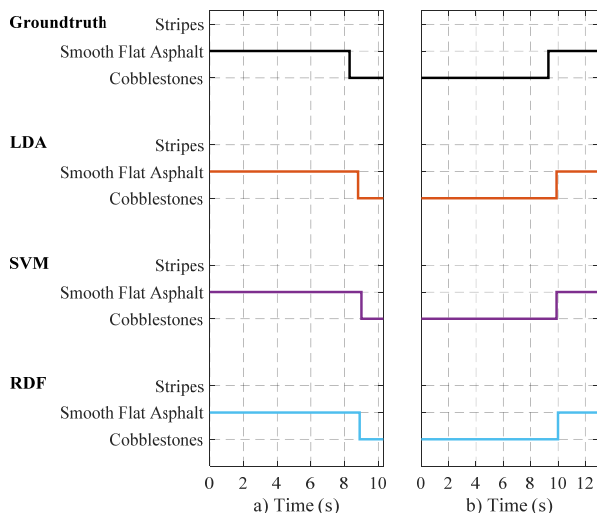

Fig. 4. Classification of two files with transitions between two types of road surface: a) stripes to asphalt; b) cobblestones to asphalt.

IV. CONCLUSIONS AND FUTURE WORK

This work has introduced a multisensor system to perform road surface identification. Four sensors were considered: accelerometers, audio signals, and the speed and handwheel channels of the EPS system. The proposed system extracts features from the four sensors, performs feature selection, and considers three possible classifiers: LDA, SVM, and RDF.

The proposed system was tested on a set of real data captured using a specially-converted passenger car on a closed course over three road surfaces: smooth flat asphalt, cobblestones, and stripes. Over this dataset, the proposed system reached 96.52% accuracy and yielded very stable results. Most of the identification errors were produced by small delays in the transitions between surfaces. These initial results show the potential of the proposed system for road surface identification.

From here, two possible lines of work remain open for future work. Firstly, it would be interesting to consider the dependence of the results on each type of sensor. In this work, four types sensors have been considered, two of which had to be specially fitted in the car. If the contribution of any given sensor is not too high or it can be compensated for with the contribution of the other sensors, it could be removed to reduce the expected cost of the system in view of a future implementation. Knowledge discovery and hierarchical methods could be used to determine these dependences [13-15]. Furthermore, this work has combined the contribution from each sensor by concatenation of their extracted features. In future works, the combination between sensors will be explored as a fusion problem through the use of advanced fusion techniques, such as alpha integration [16,17].

The second line of work consists in the testing of the proposed system on more realistic scenarios. The system has been tested with a set of short and well-behaved files. In order to adapt the system for performance in more realistic settings, field tests of longer duration would be required. Such files will probably be very noisy, resulting in a need to reconstruct missing or noisy samples [18] or remove particular sources of noise using blind source separation [19]. Furthermore, it is likely than such a thorough experimental campaign would result in hours or days of signals that should be carefully labeled by hand. In order to reduce the amount of work, future developments of the system will consider the possibility of semi-supervised learning [20,21]. To the best of our knowledge, semi-supervision has not been used yet on road surface identification, although it has been used in related topics such as road segmentation for image vision [22].

REFERENCES

[1] J.C. Tudon-Martinez, S. Fergani, O. Sename, J.J. Martinez, R. Morales-Menendez, and L. Dugard, "Adaptive Road Profile Estimation in Semiactive Car Suspensions," IEEE Transactions on Control Systems Technology, vol. 23, no. 6, pp. 2293-2305, 2015.

[2] J. Park, K. Min, H. Kim, W. Lee, G. Cho, and K. Huh, "Road surface classification using a deep ensemble network with sensor feature selection," Sensors, vol. 18, art.no. 4342, 2018.

[3] J. Alonso, J.M. López, I. Pavón, M. Recuero, C. Asensio, G. Arcas, and A. Bravo, "On-board wet road surface identification using tyre/road noise and support vector machines," Applied Acoustics, vol. 76, pp. 407-415, 2014.

[4] J. Zhao, H. Wu, and L. Chen, "Road surface state recognition based on SVM optimization and image segmentation processing," Journal of Advanced Transportation, vol. 2017, art. no. 6458495, 2017.

[5] J. Masino, J. Pinay, M. Reischl, and F. Gauterin, "Road surface prediction from acoustical measurements in the tire cavity using support vector machine," Applied Acoustics, vol. 125, pp. 41-48, 2017.

[6] A. Bystrov, E. Hoare, T.Y. Tran, N. Clarke, M. Gashinova, and M. Cherniakov, "Automotive surface identification system," in IEEE International Conference on Vehicular Electronics and Safety (ICVES), pp. 115-120, Vienna (Austria), June 2017.

[7] A. Bystrov, E. Hoare, T.Y. Tran, N. Clarke, M. Gashinova, and M. Cherniakov, "Road surface classification using automotive ultrasonic sensor," Procedia Engineering, vol. 168, pp. 19-22, 2016.

[8] A. Bystrov, M. Abbas, E. Hoare, T.Y. Tran, N. Clarke, M. Gashinova, and M. Cherniakov, "Analysis of classification algorithms applied to road surface recognition," in 2015 IEEE Radar Conference (RadarCon), pp. 907-911, Piscataway, NJ (USA), May 2015.

[9] G. Peeters, A large set of audio features for sound description (similarity and classification) in the CUIDADO project. CUIDADO IST Project Report, vol. 54, pp.1-25, 2004.

[10] H. Lui and H. Motoda, Eds., Computational Methods of Feature Selection. Boca Ratón, FL, USA: CRC Press, 2007.

[11] David M.W. Powers, "Evaluation: From Precision, Recall and F-Measure to ROC, Informedness, Markedness & Correlation," Journal of Machine Learning Technologies, vol. 2, no. 1, pp. 37-63, 2011.

[12] G. Safont, A. Salazar, L. Vergara, E. Gomez, and V. Villanueva, "Probabilistic Distance for Mixtures of Independent Component Analyzers," IEEE Transactions on Neural Networks and Learning Systems, vol. 29, no. 4, pp. 1161-1173, 2018.

[13] A. Salazar, J. Igual, L. Vergara, and A. Serrano, "Learning hierarchies from ICA mixtures," In Proceedings of the IEEE International Joint Conference on Artificial Neural Networks, pp. 2271–2276, 2007.

[14] A. Salazar J. Gosalbez, I. Bosch, R. Miralles, and L. Vergara, "A case study of knowledge discovery on academic achievement, student desertion and student retention," Proceedings of IEEE ITRE 2004 - 2nd International Conference on Information Technology: Research and Education, pp. 150-154, 2004.

[15] A. Salazar, J. Igual, G. Safont, L. Vergara, and A. Vidal, "Image applications of agglomerative clustering using mixtures of non-Gaussian distributions," Proceedings of the International Conference on Computational Science and Computational Intelligence (CSCI), pp. 459-463, Las Vegas, NS (USA), 2015.

[16] A. Soriano, L. Vergara, B. Ahmed, and A. Salazar, "Fusion of scores in a detection context based on alpha integration," Neural Computation, vol. 27, no. 9, pp. 1983-2010, 2015.

[17] G. Safont, A. Salazar, and L. Vergara, "Multiclass alpha integration of scores from multiple classifiers," Neural Computation, vol. 31, no. 4, pp. 806-825, 2019.

[18] G. Safont, A. Salazar, A. Rodriguez, and L. Vergara, "On Recovering missing ground penetrating radar traces by statistical interpolation methods," Remote Sensing, vol. 6, no. 8, pp. 7546-7565, 2014.

[19] R. Llinares, J. Igual, A. Salazar, and A. Camacho, "Semi-blind source extraction of atrial activity by combining statistical and spectral features," Digital Signal Processing: A Review Journal, vol. 21, no. 2, pp. 391-403, 2011.

[20] J. Igual, A. Salazar, G. Safont, and L. Vergara, "Semi-supervised Bayesian classification of materials with impact-echo signals," Sensors, vol. 15, no. 5, pp. 11528-11550, 2015.

[21] A. Salazar and L. Vergara, "ICA mixtures applied to ultrasonic nondestructive classification of archaeological ceramics," EURASIP Journal on Advances in Signal Processing, vol. 2010, art. no. 8, 2010.

[22] X. Han, J. Lu, C. Zhao, S. You, and H. Li, "Semisupervised and Weakly Supervised Road Detection Based on Generative Adversarial Networks," IEEE Signal Processing Letters, vol. 25, no. 4, pp. 551-555, 2018.