Improving pan-European speed-limit signs recognition with a new “global number segmentation” before digit recognition

Alexandre Bargeton, Fabien Moutarde, Fawzi Nashashibi, Benazouz Bradai

Improving pan-European speed-limit signs recognition with a new "global number segmentation" before digit recognition

Alexandre Bargeton, Fabien Moutarde, Fawzi Nashashibi, and Benazouz Bradai

Abstract— **In this paper, we present an improved European speed-limit sign recognition system based on an original "global number segmentation" (inside detected circles) before digit segmentation and recognition. The global speed-limit sign detection and correct recognition rate, currently evaluated on videos recorded on a mix of French and German roads, is around 94 %, with a misclassification rate below 1%, and not a single validated false alarm in several hours of recorded videos. Our greyscale-based system is intrinsically insensitive to colour variability and quite robust to illumination variations, as shown by an on-road evaluation under bad weather conditions (cloudy and rainy) which yielded 84% good detection and recognition rate, and by a first night-time on-road evaluation with 75% correct detection rate. Due to recognition occurring at digit level, our system has the potential to be very easily extended to handle properly all variants of speed-limit signs from various European countries. Regarding computation load, videos with images of 640x480 pixels can be processed in real-time at ~20frames/s on a standard 2.13GHz dual-core laptop.**

I. INTRODUCTION

Car automation increases by progressive integration of more and more advanced driving assistance systems. For instance, most current GPS navigators now include a function to inform the driver of the supposed current speed-limit, a feature increasingly motivating drivers as automated speed-limit enforcement gets more common. Furthermore a desired evolution for Adaptive Cruise Control (ACC) would be the development of smarter ACC able to automatically tune target cruising speed depending on current speed-limit.

However speed-limit information extracted from GPS cartographic data is neither always complete nor systematically up-to-date. Moreover, temporary speed limits for road works (see example on figure 1), as well as variable speed limits, are by definition not included in pre-defined digital cartographic data. Therefore a *visual* real-time speed-limit sign detection and recognition system is a mandatory complement to GPS systems for designing high-level advanced driving assistance systems such as Speed Limit Support (SLS) and smart ACC.

Manuscript received Januar 7, 2008.

Alexandre Bargeton, Fabien Moutarde and Fawzi Nashashibi are with the Robotics Laboratory (CAOR), Ecole des Mines de Paris (ParisTech), 60 Bd Saint-Michel, F-75272 Paris cedex06, FRANCE, Phone: (33) 1-40.51.92.92, Fax: (33) 1-43.26.10.51
(e-mail: Alexandre.Bargeton@ensmp.fr, Fabien.Moutarde@ensmp.fr, Fawzi.Nashashibi@ensmp.fr)

Benazouz Bradai is with Valeo Driving Assistance Domain, 34 rue St-André, ZI des Vignes, F-93012 Bobigny, FRANCE, Phone: (33) 1-49.42.60.95 (e-mail: benazouz.bradai@valeo.com).

Fig. 1. Example of roadwork temporary speed-limit sign, whose information can obviously not be present in GPS cartographic data, and therefore has to be visually detected and recognized for a Speed Limit Support system.

II. RELATED AND PREVIOUS WORK

A Traffic Sign Recognition (TSR) system usually involves two main steps: 1/ detection of potential traffic signs in the image, based on the common shape/color design of sought traffic signs; 2/ classification of the selected regions of interest (ROIs) for identifying the exact type of sign, or rejecting the ROI.

Most recently published TSR approaches make use of color information for the detection step (see e.g. [1], [2], [3], [4] or [10]), which makes it easier, but less robust. In contrast with that, our TSR system, whose initial version was already presented in [9] and [13], uses a shape-based detection working on grayscale images. As was already advocated by Gavrila in [5], by Barnes and Zelinsky in [6], and confirmed by García-Garrido et al. in [8], using a shape-based detection increases robustness for detection of signs with colors faded away by time, and makes it possible to work properly even with difficult illumination conditions, such as glare by background sun or light, in the dark, or even at night.

For the classification step, nearly all published works on speed-limit sign recognition use a "holistic" approach in which the whole sign (in fact most of the time a set of features extracted from it) is fed into a classifier. Various kinds of classifiers are used: Bayesian Maximum Likelihood approach after Linear Discriminant Analysis (LDA) [1], Uncorrelated Fisher Discriminant Analysis (UFDA) [12] or after Image matrix based Discriminant analysis (IFDA) [12], ART1 (Adaptive Resonance Theory) neural network in [2], normalized correlation-based pattern matching in [3] and something similar in [6], Radial Basis Function (RBF) network in [5], a fuzzy set approach applied to color and shape features in [7], backpropagation neural network applied to global normalized sign image in [8] and [10], and more recently new learning methods such as ensemble learning in [11].

Contrary to all those works, and as already described in more details in [9] and [13], our Speed Limit Support system prototype relies on a digit extraction and recognition scheme, as in Torresen et al. [4], but more general and robust because we use separate recognition of each of the 2 or 3 digits, and employ a digit segmentation that is "orientation insensitive" so as to be able to recognize properly slightly tilted signs.

The recognition of extracted digits can be made, as described in [14], by any of commonly used algorithms from pattern recognition domain (kNN, MLP, RBF, SVM,…), which all can reach more than 98.5% in character recognition accuracy. In our algorithm, digits are recognized by a multi-layer neural network (MLP) trained on digit examples extracted from real videos.

In our system, a temporal integration on several frames is also used to confirm (or infirm) the recognition of the sign into the video, and provide a confidence level.

Fig. 2. Example of E.U. speed-limit sign detection and recognition with our initial system. All detected circles are shown (in red or green), and the candidate digits segmented inside them are outlined in red. The speed signs recognized on the current image are shown on the top black zone, with their associated confidence, and the currently validated speed limit is superimposed on top-center of image.

In our first version, the digit segmentation consisted in connected components extraction inside binarized detected circles. This technique proved its robustness and computation efficiency. However, it suffers in some cases of non-segmentation when two digits "touch" each other on the binarized image (see examples on fig. 3, 4, 6 or 7), thus resulting in non-recognition of the speed limit sign. These cases, which mostly happen for small and/or distant signs, 3-digits signs whose digits are not well-separated, or tagged signs, are not so common but significantly impaired performance, notably on German signs whose 3-digits signs are often small. We therefore decided to design a complementary technique to solve this problem, which is presented in this paper.

We also present here a new systematic evaluation of our improved system, on a mix of French and German roads, which are very encouraging for a future pan-European Speed Limit Support system, and results of recent on-road evaluations under bad weather conditions, and at night.

III. Improved Digit Character Segmentation

The Optical Character Recognition (OCR) domain provides three mains techniques in order to segment characters [15]: (1) via the image histogram, (2) via over-segmentation, (3) and via word by word segmentation. The first one fails with overlapped characters and depends on the precision of circles detection and image quality (see figure 3).

Fig. 3. Histograms of binarized speed limit signs

The second one is not applicable for real-time detection and can imply a lot of false recognition. We therefore turned to some hybrid of the third kind of method and our initial technique: we first try to globally segment the "word" inside the sign, which is in fact a 2- or 3-digits number in this case. The new proposed algorithm to segment digits inside a speed limit sign thus consists in two steps:

- Find the number / word into the circular sign
- Segment digits into the obtained rectangular zone

A. "Global number segmentation" into the sign

We further divide the number segmentation in two successive steps:

- finding upper/lower limits around the number
- determining left/right limits of the number

The two algorithms are detailed below.

A.1. Search upper/lower limits: The general idea is some kind of pixel-by-pixel "guided propagation". For each sub-quarter of the image, starting from each pixel on a small central horizontal segment (Fig 4a – blue segment), browse the image pixel by pixel choosing one of the 3 neighbours pixels (Fig 4c) firstly attracted by the vertical center and secondly by the top (resp. bottom) of the image (Fig 4a). The vertical limits are the highest (resp. lowest) end point (Fig 4b). To prevent the propagation on the surrounding black circle, a rule is added to block the way exceeding a vertical segment / limit (i.e. stop the propagation to the outside if it's to too far from the vertical center). A last rule is that if the height of a propagation from a starting pixel is too close to the height of the sign (ie propagated on the outer circle) or too small (ie no propagation), this propagation is considered as failed, and not included in the computation of the upper / lower limit.

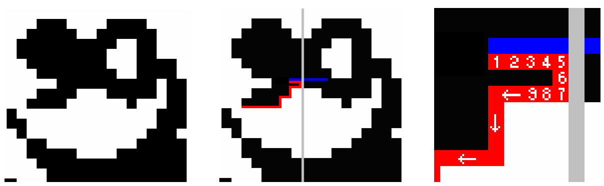

Fig. 4a. Searching vertical limits

(left) Initial binarized image – (center) a bottom left subquarter propagation way in red starting from the blue centered segment – (right) a zoomed part of this way from start with indexed pixels and direction of the propagation

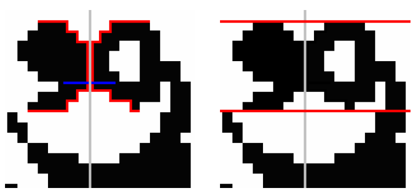

Fig. 4b. Searching vertical limits

(left) The propagation in red for each sub-quarter – (right) the result is the highest way for each half part (top – bottom) of the image.

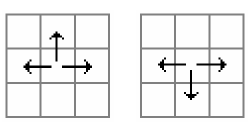

Fig. 4c. Searching vertical limits

The three neighbors used for top propagation (left) and bottom one (right)

A.2. Search left/right limits: Browse column by column from the vertical center to left and right, inside the previously found region. The main idea is that the first column (on each side) from which a black pixel is vertically outside the region or which is a white column (no more digit) is the searched vertical limit (see figure 5). In order to prevent premature stop between two separated digits, we temporally save the first horizontal index and continue our search by browsing columns with the same stop criteria. When there is no more way to stop, we use the last saved index as the limit.

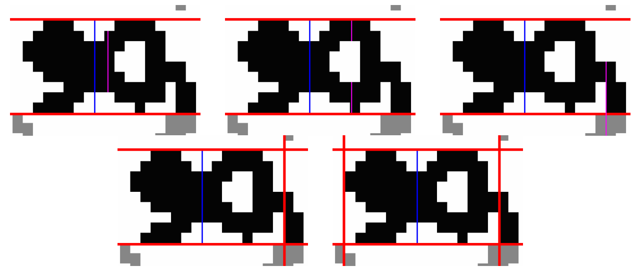

Fig. 5. Searching horizontal limits: exploration column by column (purple segment) for each half part (left – right) and stop on a white column (left) or exceed black (right).

This algorithm has been tested and validated under various conditions: day and night (see figures 6 and 7), with backlight or front light illumination resulting with some poor quality binarized images (see figure 6).

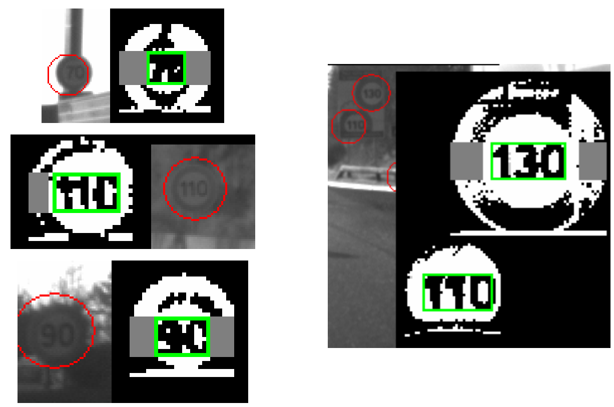

Fig. 6. Global number segmentation

B. Digit segmentation inside the segmented number

We use the same technique as in our initial method, but instead of binarizing the speed limit sign image, we binarize only the rectangular region found by our "number segmentation" algorithm. The main idea is that the binarization of the sub-image is much easier and provides a greater result than in the circle / speed image (see figure 7). In this new image we can easily extract connected components.

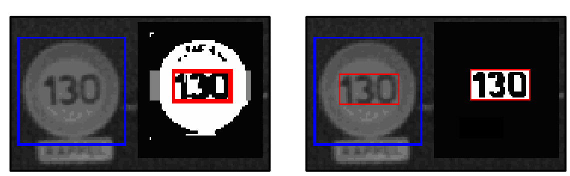

Fig. 7. Image binarization into the new area

There are still some digit-overlap cases, but the practice showed that overlap is done by only one or two pixels between digits in this new type of image. So in parallel we provide a histogram segmentation (see figure 3) which now can work very well.

IV. Experiments and Results

A traffic sign detection and recognition system can be evaluated with various criteria. What is important in our application context is that all sought signs are validated *at least once* during the video segment between its first appearance and its final disappearance. We therefore evaluate only the *global* system performance by comparing type/time/position of all *validated* signs issued by our system to a ground-truth indicating a space-time visibility interval and type for each potentially detectable speed sign. The main resulting measure is the percentage of speed signs correctly detected and validated within their space-time visibility interval. This comparison and measure is of course done on video recordings independent from those used for extracting digits for training.

A. Recognition improvement on "difficult cases"

A first evaluation of the improvement brought by our new "number segmentation" approach has been made by considering only a set of speed-limit signs which were NOT correctly recognized with our initial digit-segmentation-based method. The percentage of those "difficult signs" that are now properly identified with our algorithm including "number segmentation" turned out to be 11/18 = 61%. This means our approach does not solve *all* recognition problems, but nevertheless provides a solution for a quite significant proportion of them, such as the tagged sign on figure 8.

B. Recognition improvement global evaluation

A more thorough evaluation has been done for the E.U. system, using recordings *on a mix of French and German roads and streets*, under various daytime illumination conditions, and containing ~140 speed-limit signs covering 11 different limit values (30, 40, 50, 60, 70, 80, 90, 100, 110, 120, 130). The corresponding videos have been manually annotated to indicate on each frame the position and type of speed-limit sign, so that an automated comparative evaluation is possible. The outcome is that the global system correct detection rate (SCDR) is 94 % with our improved "global number segmentation" approach, compared to 85% with our initial "digit segmentation" algorithm.

As can be seen on table 1, our new "number segmentation" quite significantly improves the correct detection rate. Also, the misclassification rate (signs for which a wrong speed value has been validated) remains below 1%. Nearly all of the 6% remaining non-correct-detections in our current improved system are just *missed* signs (most of the time because of not contrasted enough edges of the sign preventing a correct circle detection). Most importantly, not a single *validated* false alarm has been noticed in several hours of daytime recording: all spurious signs are efficiently filtered by our tracking and validation module.

Sign recognition method	**Signs detected, recognized and validated with correct type**	**Misclassified signs**
Initial digit-segmentation	**85 %**	**0,7%**
New "global number segmentation" before digit segmentation	**94 %**	**0,7%**

Table 1. Global evaluation of European speed limit sign detection on French+German roads.

Fig. 8. Example of a tagged speed-limit sign correctly recognized thanks to our new global number segmentation

C. On-road evaluation under bad weather conditions

As a complement to our in-lab evaluation done on pre-recorded videos, we hereafter present results of a recent on-road test conducted by Valeo with one of their experimental cars. The on-road evaluation (Table 2) is done *under bad weather condition (cloudy and rainy)* on 50km long run in France. Out of the 13 missed signs, 3 signs are partially occluded by others vehicles or dirty and thus nearly impossible to detect/recognize, 4 signs are with very low contrast inside tunnels, and 3 signs were correctly recognized but not validated as confidence did not pass validation threshold.

D. Preliminary night on-road evaluation

Because our algorithm is greyscale-base, and intrinsicly rather insensitive to luminosity variations, it shows very promising results for night-time operation. This was recently quantified by a first on-road night evaluation done by Valeo: during this one hour live test, 78% of 60 speed limit signs encountered have been detected and correctly recognized.

Speed limit	**Good validated detection and recognition**	**Correctly recognized but not validated**	**Missed signs**
30	1		1
50	10		3
70	21		3
80	6		
90	28	3	3
110	4		
Total	70 **84 %**	3 **4 %**	10 **12 %**

Table 2. On-road evaluation in France, under bad weather conditions (cloudy and rainy).

V. Conclusions, discussion and perspectives

We have presented an improved robust and efficient visual speed-limit signs detection and recognition with ~94 % global correct sign detection rate on French+German roads. According to our evaluation, our new "global number segmentation" approach brings a quite significant 9% increase in detection rate, compared to our initial already performant system. The system requires only greyscale videos, and is able to process in real-time a video flow of images with 640x480 pixels at ~20 frames/s on a standard 2.13GHz dual-core laptop. It has a remarkably low false alarm rate (less than 1 spurious sign in several hours of operation).

The quantitative E.U. system evaluation reported in the present paper was restricted to videos recorded on French and German streets and roads. Evaluations in other E.U countries (in which sign digits are sometimes slightly different), are currently in progress, with very promising results (see example of successful detection of an Italian speed-limit sign on figure 9), as long as the ODR neural network is trained on a database with examples of all required variants of digits.

Our use of digit extraction and recognition instead of global sign recognition clearly facilitates the handling of aspect variability of the same sign across different E.U. countries, and even inside a single country such as France for instance. Even though we presently train a specific neural network for speed sign digits recognition, a final really pan-European system (i.e. properly recognizing all variants of speed-limit signs in all European countries) could probably work with a "universal" digit recognition module, therefore not requiring the previous collect of all European variants of each speed-limit sign digits.

Fig. 9. Correct recognition of a Italian speed-limit sign illustrating promising results for pan-European speed-limit recognition.

Fig. 10. Successful LED speed-limit sign recognition obtained with our system without specific re-training, by just inverting pixel scale inside region of detected circles.

Another illustration of the advantage of our speed-sign recognition being based on grayscale and digit recognition is the ease of adaptation to operate properly on LED signs, which are more and more common, especially for variable speed limits, for which the visual recognition is a mandatory complement to GPS. According to a first quick experiment, by simply inversing the image pixel scale inside the circles of potential signs, and then applying our algorithm, we can obtain, without any specific re-training, correct recognition of LED signs, as illustrated on figure 10.

Other work done in parallel and presented elsewhere includes recognition of end-of-speed-limit signs, and of some supplementary sign placed under signs, those two features being essential for a fully pertinent Speed Limit Support system.

Finally, the final system really makes sense when integrated with GPS information, which can provide “baseline” information for the unavoidable cases of signs occulted by other vehicles. We therefore have begun to develop a framework for fusion of the output of visually-detected speed limits with GPS cartographic speed-limit data, as presented in [16]. Preliminary experiments show quite promising results for a final system that could take into account those two complementary speed-limit information sources.

REFERENCES

[1] Bahlmann C., Zhu Y., Ramesh V., Pellkofer M. and Koehler T., “A System for Traffic Sign Detection, Tracking, and Recognition Using Color, Shape, and Motion Information”, IEEE Intelligent Vehicles Symposium (IV 2005), Las Vegas, June 2005.

[2] de la Escalera A., Armingol J.M. and Mata M., “Traffic sign recognition and analysis for intelligent vehicles”, Image and Vision Computing, 21:247-258, 2003.

[3] Miura J., Kanda T. and Shirai Y., “An active vision for real-time traffic sign recognition”, Proc. IEEE Conf. on Intelligent Transportation Systems, pages 52-57, Dearbom, MI, 2000.

[4] Torresen J., Bakke J.W., and Sekania L., “Efficient recognition of speed limit signs”, Proc. IEEE Conf. on Intelligent Transportation Systems (ITS), Washington DC, 2004.

[5] Gavrila D.M., “Traffic sign recognition revisited”, Proc of 21st DAGM symposium fur Musterekennung, pp. 86-93, Springer-Verlag, 1999

[6] Barnes N. and Zelinsky A., “Real-time radial symmetry for speed sign detection”, Proc. IEEE Intelligent Vehicle Symposium, pages 566-571, Parma, Italy, 2004.

[7] Fang C.-Y., Chen S.-W and Fuh C.-S., “Road-sign detection and tracking”, IEEE Trans. On Vehicular Technology, Vol.52, N°5, September 2003.

[8] GarcÍa-Garrido M. A., Sotelo M. A., and MartÍn-Gorostiza E., “Fast traffic sign detection and recognition under changing lighting conditions”, Proc. IEEE Intelligent Transportation Systems conference, pages 811-816, Toronto, Canada, 2006.

[9] Moutarde F., Bargeton B., Herbin A. and Chanussot C., "Robust on-vehicle real-time visual detection of American and European speed limit signs, with a modular Traffic Signs Recognition system", proc. of IEEE Intelligent Vehicles Symposium, Istanbul, 13-15 juin 2007.

[10] Broggi A., Cerri P., Medici P., Porta P. and Ghisio G., “Real Time Road Signs Recognition”, IEEE-IV2007 , Istanbul, Turkey, 13-15 June, 2007

[11] Kouzani A., “Road-Sign Identification Using Ensemble Learning”, IEEE-IV2007 , Istanbul, Turkey, 13-15 June, 2007

[12] Zhang C., Yang Q., Zhang W and Wang M., "Traffic sign recognition using Discriminant analysis based Algorithms", proc. of 14th World congress on Intelligent Transportation Systems (ITS), Beijing, China, 9-13 October, 2007

[13] Moutarde F., Bargeton B., Herbin A. and Chanussot C., "Modular Traffic Sign Recognition applied to on-vehicle real-time visual detection of American and European speed limit signs", proc. of 14th World congress on Intelligent Transportation Systems (ITS), Beijing, 9-13 octobre 2007.

[14] Liu C.-L., Fujisawa H., “Classification and Learning for Character Recognition: Comparison of Methods and Remaining Problems in Neural Networks and Learning in Document Analysis and Recognition”, First IAPR TC3 NNLDAR Workshop, p1-7, Seoul, Korea, August 29, 2005.

[15] Casey R. G., Lecolinet E., “A Survey of Methods and Strategies in Character Segmentation”, IEEE Transaction on Pattern Analysis and Machine Intelligence, Volume 18, Issue 7, P690-706, July 1996.

[16] Lauffenburger J. Ph., Bradai B., Basset M. and Nashashibi F., “Navigation and Speed Signs Recognition Fusion for Enhanced Vehicle Location”, Proc. 17th IFAC (International Federation of Automatic Control) World Congress, Seoul, Corea, 6-11 July, 2008.