Robust on-vehicle real-time visual detection of American and European speed limit signs, with a modular Traffic Signs Recognition system

Fabien Moutarde, Alexandre Bargeton, Anne Herbin, and Lowik Chanussot

Abstract— In this paper, we present robust visual speed limit signs detection and recognition systems for American and European signs. Both are variants of the same modular traffic signs recognition architecture, with a sign detection step based only on shape-detection (rectangles or circles), which makes our systems insensitive to color variability and quite robust to illumination variations. Instead of a global recognition, our system classifies (or rejects) the speed-limit sign candidates by segmenting potential digits inside them, and then applying a neural network digit recognition. This helps handling global sign variability, as long as digits are properly recognized. The global sign detection rate is around 90% for both (standard) U.S. and E.U. speed limit signs, with a misclassification rate below 1%, and not a single validated false alarm in >150 minutes of recorded videos. The system processes in real-time videos with images of 640x480 pixels, at ~20frames/s on a standard 2.13GHz dual-core laptop.

I. INTRODUCTION

The current path to future really intelligent vehicles is the progressive integration in new car models of smart advanced driving assistance systems. Traffic signs along the road provide many crucial information for safe and efficient driving. However, some of them may escape the driver's attention, which could be compensated by some kind of driver "companion" able to remind him latest warning or regulation traffic signs. Also some smart functions such as adaptive cruise control (ACC) could be further automated and improved by taking into account traffic sign information. Traffic sign visual detection and recognition (TSR) is therefore a key module for various driving assistance functions, as it is a requirement for the necessary level of traffic scene understanding. For example a robust visual real-time TSR system is a pre-requisite for developing a system reminding the driver (or informing the ACC) what is the current speed limit. Some of the traffic sign information may sometime be extracted from the GPS navigation data, but it is neither always complete nor systematically up-to-date. Moreover, temporary speed limits for road works, as well as variable speed limits, are by definition not included in pre-defined digital cartographic data. Therefore a visual real-time TSR system is a mandatory complement to GPS systems for designing high-level advanced driving assistance systems.

II. RELATED WORK

A 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. As noted by Bahlman et al. in [1], the majority of recently published TSR approaches make use of color information (see e.g. [2], [3] or [4]), which makes the detection step easier. In contrast with that, the TSR system presented in the present paper uses a shape-based detection working on grayscale images. This puts less constraint on the required sensor, which may help meet global costs requirements. Moreover and most importantly, as was already pointed out and advocated by Gavrila in [5], by Barnes and Zelinsky in [6], and confirmed by García-Garrido et al. in [8], being grayscale- and shape-based also simplifies correct operation in dark or night condition, or robustness for detection of signs with colors faded away by time.

Regarding the classification step, the majority of published works on speed-limit sign recognition have adopted 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 in [1], 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]. Contrary to all those works, and because digits variability is lower than that of global signs (particularly in the U.S.), we decided *not* to use a global "holistic" recognition approach, but instead to rely on a digit extraction and recognition

scheme, as in Torresen et al. [4]. However, we use separate recognition of the 2 or 3 digits, while [4] rely on recognition of only the first digit, making their system unable to work on 3-digits or "ending-by-a-5" speed limit values. Also, we made sure to use a digit segmentation that is "orientation insensitive" so as to be able to recognize properly slightly tilted signs.

Finally, TSR system are generally developed for only one particular country regulation, while one of the originality of our system is to be modular enough to be easily adapted to very different traffic signs designs: we present here promising results for detection and recognition of both American (U.S.) rectangular speed limit signs, and European (E.U.) circular signs.

III. MODULAR SYSTEM ARCHITECTURE

A. Architecture overview

The system presented here is implemented using the ^{RT}MAPS® software for real-time multi-sensor applications prototyping, distributed by the Intempora company (www.intempora.com). This prototyping environment has already been adopted by many French car manufacturers for the development of on-vehicle real-time applications.
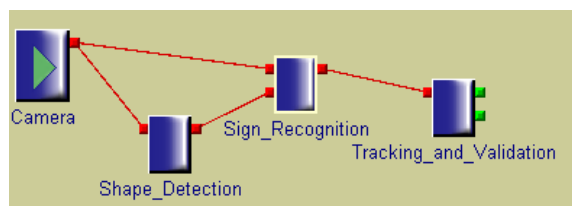

Fig. 1. Global traffic sign recognition modular architecture

We take full advantage of the intrinsic modularity offered by the graphical programming paradigm of ^{RT}MAPS®, by using separate modules for the main successive steps of TSR:

- detection of potential traffic signs
- recognition
- tracking and validation

The modularity of our system allows to easily adapt it to different types of speed-limit signs (e.g. U.S. vs E.U. speed limit signs), or even detection of other kind of traffic sign.

B. Detection of potential traffic signs

As already mentioned we wanted our system to operate even on grayscale videos. Therefore, our detection modules are based on shape-detection, as for instance in [5], but more general. We have already developed two different detection modules for two common shapes of traffic signs:

- a circular Hough-transform specially adapted and tuned for the application to European Union (E.U.) speed limits signs, which are circular;
- a specially-designed rectangle-detection based on edge detection for United States (U.S.) speed limit signs, which are rectangular.

Fig. 2. Example of successful rectangular sign detection, despite low contrast on the background.

Our rectangle-detection algorithm shall be described elsewhere, after a currently pending patent application has been filed.

The aim of the detection stage is to miss as few real sought signs as possible. It is in particular essential to be able to detect efficiently even in the case of low luminosity and/or contrast of the sign contour on the background, as illustrated on figure 2 and 3.

Fig. 3. Example of successful circular sign detection, despite low contrast and luminosity.

Conversely, false detections at this stage are not a problem, as they will be efficiently filtered by the recognition step. Depending on parameterization of the detection module, and on image content, we typically obtain between several and several tens of candidate zones in each image, as can be seen on figures 2, 3, 4, 8 or 9. There are naturally often more irrelevant rectangles than circles, as there are more rectangular-shaped objects in the environment, especially in urban or semi-urban areas (see figure 4). However, this does

not impair the global performance of our system for U.S. speed-limit signs, for the reason explained above.

Fig. 4. Illustration of typical number of irrelevant rectangles in urban area. It can be noted that most of them do not even contain a single digit candidate (red rectangles), so that they are eliminated right away by the first step of our sign recognition scheme.

C. Recognition of traffic signs

The current version of the recognition part of our system is further subdivided in more modules, as illustrated in figure 5. The first sub-module tries to segment characters inside every potential speed-limit signs resulting from the detection step. Each potential digit is then normalized and fed into a neural network digit recognition module. A final sign classification module issues the recognized sign type and a confidence evaluation in the decision.
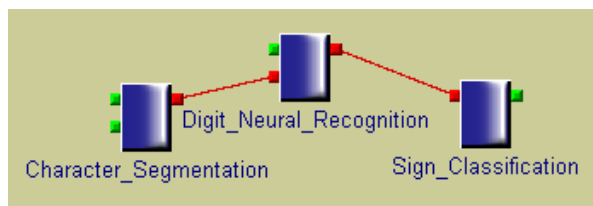

Fig. 5. Detail of the Sign_Recognition block for the E.U. case

Due to real-world signs not always being strictly upright, as well as variation of vehicle roll angle, the sign digits orientation in the image is not always perfectly horizontal. The character segmentation algorithm was therefore chosen to be intrinsically insensitive to orientation variations, as illustrated on figure 7. The segmentation used is a connected-component labeling applied after a binarization obtained by adaptive thresholding on the circular or rectangular area, as illustrated on figure 6. Of course digit segmentation does not *always* work, as in some cases the digits are not well separated on the binarized image. This happens in particular when sign is too small and far, or if some tag or dirt is present in the digits area. The first case

will most of the time just delay the sign recognition, as the digits will eventually appear separated as the car approaches the speed-limit sign. The second case is in practice not so common.

Fig. 6. Illustration of the character segmentation procedure in a rectangle (U.S. case, top line), and in a circle (E.U. case, bottom line): the searched area (left) is first binarized by adaptive thresholding (middle), the connected-component labeling is applied for finding potential digits (right)

The potential digits segmented are then normalized and classified by a neural-network optical digit recognition (ODR) module. The neural network ODR module is a multilayer perceptron with 10 output (1 for each digit value) trained on specially built databases of digits extracted from relevant speed limit signs in videos recorded from on-vehicle camera. For example, our E.U. digits database, constantly enriched with new examples extracted from video recorded in different European countries, currently contains a total of 2762 digits examples and 2789 negative (non-digits) examples.
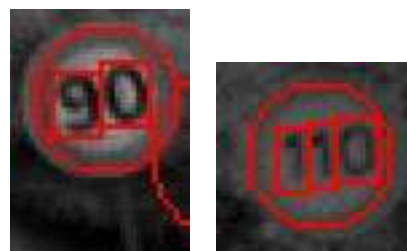

Fig. 7. Illustration of the robustness to speed-limit sign inclination of the chosen character segmentation method.

For the U.S. case, another classifier is also used to analyze the region above the digits in order to verify that the sign is actually a standard speed limit, and not for instance a "truck speed" sign (an example of recognition of such a truck-dedicated speed-limit sign can be seen on figure 8).

D. Tracking and validation of recognized traffic signs

Depending on all the outputs of the neural network ODR module, a confidence measure is computed and assigned to the detected and recognized speed limit signs. This confidence is further increased if the same sign is again

detected and identified at expected image location in subsequent video frames; it is also decreased by a fixed factor for each frame where no recognized sign matches the tracked one. The recognized speed-limit sign is finally validated only when its confidence gets over a validation threshold. The value of this threshold is typically determined so that validation occurs if the sign is recognized with reasonable confidence on at least 2 or 3 nearly successive frames.

IV. EXPERIMENTS AND RESULTS

A traffic sign detection and recognition system can be evaluated with various criteria. For instance one can compare image by image the identified signs with those actually visible. But this "static" performance measure is not the most important, as a given sign shall generally be visible on several successive frames, and it is thus not essential to identify it on *every* image: what is important is that all sought signs are validated *at least once* during the video segment between its first appearance and its final disappearance.
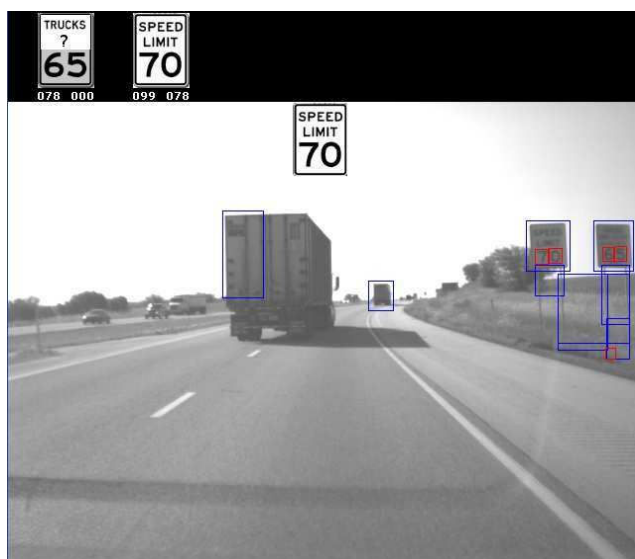

Fig. 8. Example of speed-limit sign detection on a U.S. road. All detected rectangles are visible in blue, and digit candidates segmented inside them are outlined as red sub-rectangle. The speed signs recognized on the current image are shown on the top black zone, with their associated confidence, and the currently validated speed limit is superimposed on top-center of image.

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 our neural ODR

module. Preliminary evaluation of the U.S system showed a global system correct detection rate (SCDR) of ~90%.

A more thorough evaluation has been done for the E.U. system, using ~150 minutes of recordings on French roads and streets, under various daytime illumination conditions, and containing 281 speed-limit signs covering 12 different limit values (10, 20, 30, 40, 45, 50, 60, 70, 80, 90, 110, 130). The global system correct detection rate (SCDR) is ~89%.

Total number of speed signs	Signs detected and validated with correct type	Missed signs (not validated)	Signs detected and validated, but misclassified
281	250	29	2

Table 1. First global evaluation of European speed limit sign detection.

As can be seen on table 1, nearly all of the 11% non-correct-detections are just *missed* signs (sometimes because of not contrasted enough edges of the sign, and most of the time because of noise or partial occlusion impeding digit segmentation). The misclassification rate (signs for which a wrong speed value has been validated) is below 1%.

Most importantly, not a single *validated* false alarm has been noticed in the 150 minutes of daytime recording: all spurious signs are efficiently filtered by our tracking and validation module.

Fig. 9. Example of speed-limit sign detection on a E.U. street. All detected circles are shown (in red), 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.

V. CONCLUSIONS AND PERSPECTIVES

We have presented a robust and efficient visual speed-limit signs detection and recognition system, with 2 variants of the same global architecture working respectively for U.S. and E.U. signs, both with ~90% global correct sign detection rate. The system requires only grayscale 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 150 minutes of operation).

The quantitative E.U. system evaluation reported in the present paper was restricted to videos recorded on French streets and roads. Evaluations in other E.U countries (in which sign digits are sometimes quite different, as illustrated in figure 10) are currently in progress, with very promising results (see example of successful detection of a German speed-limit sign on figure 11), as long as the ODR neural network is trained on a database containing examples of all required variants of digits.

Fig. 10. Illustration of the variability of digits across E.U. countries: German signs (left of each pair) compared to their French equivalent.

Also, the results reported in this paper are only for daytime operation. However, evaluation of the performances of our system at night and in tunnels is currently in progress, and very encouraging.

Future work shall in particular focus on eliminating the remaining ~1% misclassifications, and lowering the current ~10% miss-rate, by adding a complementary sign recognition scheme. In particular, an original and promising way of handling cases where digits are not well separated is currently under development.

Fig. 11. Correct recognition of a German speed-limit sign illustrating promising results for pan-European speed-limit recognition.

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.