SPIE Defense & Security, Visual Information Processing XIX, Proc. SPIE, Vol. 7701, 77010V (2010)

Unconstrained License Plate Detection Using the Hausdorff Distance

M. Lalonde, S. Foucher, L. Gagnon
R&D Department, Computer Research Institute of Montreal (CRIM)
405 Ogilvy Avenue, Suite 101, Montreal, QC, Canada, H3N 1M3
{marc.lalonde, samuel.foucher, langis.gagnon}@crim.ca

ABSTRACT

This paper reports on a new technique for unconstrained license plate detection in a surveillance context. The proposed algorithm quickly finds license plates by performing the following steps. The image is first pre-processed to extract the edges; opening with linear structuring elements ensures that plate sides are enhanced. Multiple scans using the Hausdorff distance are made through the vertical edge map with binary templates representing a pair of vertical lines (with varying gap to account for unknown plate size), so they efficiently pinpoint areas in the image where plates may be located. Inside those areas, the Hausdorff is used again, this time over the gradient image and with a family of templates corresponding to rectangles which have been subjected to geometric transformations (to account for perspective effects). The end result is a set of plate location candidates, each associated to a confidence level that is a function of the quality of match between the image and the template. An additional criterion based on the symmetry of plate shapes also supplies complementary information about each hypothesis that allows rejection of many bad candidates. Examples are given to show the performance of the proposed method.

Keywords: license plate detection, Hausdorff distance, rectangle detection, symmetry.

1. INTRODUCTION

The aim of this paper is to describe a license plate detection module that is expected to operate in unconstrained environments, i.e. without specialized acquisition hardware (camera / illumination) and, above all, without knowledge of license plate position with respect to the camera. This module, as well as a scene understanding module detailed in a companion paper [2], has been developed in the course of a follow-up project on previous work already described in [1]. The goal of the latter project was to provide object tracking capabilities to a commercial surveillance system equipped with Pan-Tilt-Zoom (PTZ) cameras. In the current project, detection of license plates in a wide-angle view from a master camera enables closer snapshots by the slave PTZ camera. This context calls for the need for fast detection of potentially multiple plates at low resolution and low false alarm rate since moving the PTZ camera for the wrong reasons may be costly in terms of inappropriate focus of attention.

Automatic detection of license plates has been an active area of research for at least 25 years, and new algorithms are regularly published even though the field has matured well enough to give rise to a significant number of commercial tools. Plate detection systems that are available on the market usually rely on careful camera positioning as well as controlled illumination in order to maximize detection performance. Detecting plates using off-the-shelf surveillance cameras is more challenging due to uncertainty in plate localization and varying appearance (e.g. perspective effects). As underlined by e.g. Nguyen et al. [4], uncontrolled imaging conditions and scene complexity, among others, are major factors that still affect the performance of automatic license plate detection/recognition systems and stimulate the search for more robust techniques. A good survey of recent techniques [5] shows that there are four classes of techniques applicable to license plate detection: binary image processing (based on contour statistics, e.g. [3]), grayscale image processing (includes techniques such as Hough, wavelets, searching for high contrast areas, symmetry, etc., e.g. [9]), color image processing (more or less reliable because of the uncontrolled illumination), and machine learning (Adaboost [10], SVM, neural networks, genetic algorithms, fuzzy logic, etc.). However, few algorithms are expected to perform well as plate size is expected to be as low as 5% of the image width (about 30 pixels in a 640x480 image). The proposed technique relies on pattern matching with the Hausdorff distance to pinpoint plate candidates, and on a symmetry detector to validate/reject the candidates.

The paper is organized as follows: Section 2 describes the license plate search strategy while Section 3 focuses on hypothesis validation based on the symmetry detector; Section 4 shows some results followed by a discussion.

2. SEARCH STRATEGY

The proposed strategy for license plate detection is essentially a shape-based approach, i.e. an efficient search for rectangles in the image. The search is based heavily on the Hausdorff distance [6]. This distance measures the dissimilarity between two sets of points:

$$h(A,B) = \max_{a \in A} \{ \min_{b \in B} \{ d(a,b) \} \}$$

where *a* and *b* are points belonging to sets A and B and *d*(*a*,*b*) is a distance measure (metric). In computer vision, the Hausdorff distance has been used for pattern matching since a pattern can be seen as a collection of points to be compared to another collection of points (an image or a portion of image). The first stage is a pre-processing step whose goal is to enhance the contours of the plate. A Sobel image is computed and thresholded (to eliminate spurious edges) followed by morphological filtering with horizontal and vertical line segments as structuring elements. This stage also produces a mask image that contains white pixels in the area of vertical and horizontal edges. The mask image (Figure 1, bottom left) is designed to block areas where the plate is not likely to be found. The assumptions are similar to [3], i.e. the plate is roughly horizontal.

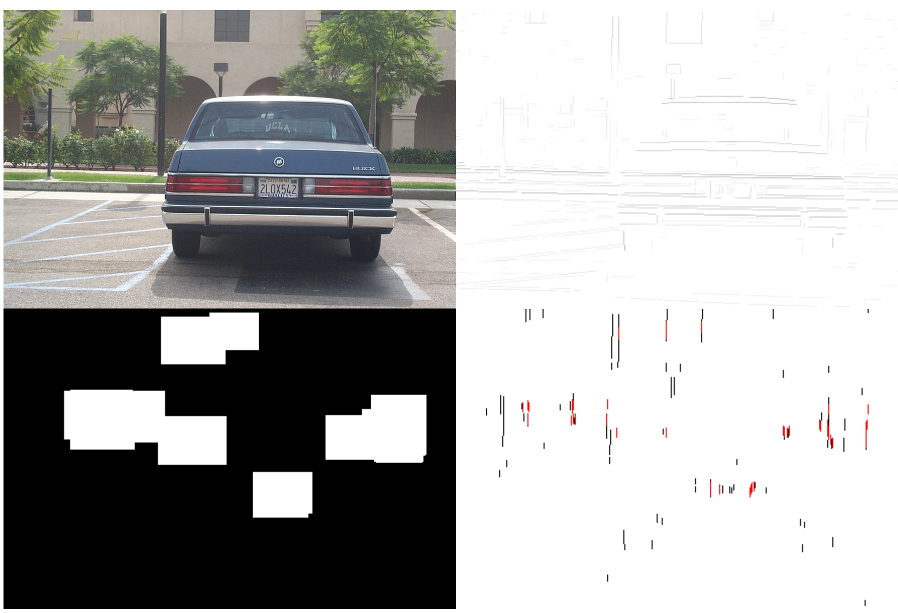

Figure 1 – Pre-processing stage. Top left: original image; top right: Sobel + morphological filtering; bottom left: mask image; bottom right: areas where the templates containing vertical lines match the edges.

The second stage is a region search that is designed to pinpoint candidate regions in the image where a plate could be located. A set of binary templates containing two vertical line segments of varying empty space between them are generated and used as input patterns to the Hausdorff-based search algorithm over the vertical component of the Sobel image (also using the mask produced earlier). Many templates are used to account for the unknown license plate size.

The regions R_i where the Hausdorff distance between a given template and the corresponding sub-image is lower than a threshold are stored in memory. Also note that the gap between the line segments of the matching templates is stored since it represents a good approximation of the license plate size in pixels. Figure 1 (bottom right) gives an example. The third stage again makes use of the Hausdorff distance to search for the actual plate shape using another collection of templates, this time containing rectangles of various orientations and proportions (see Figure 2). These templates are obviously designed to match the large span of variations in plate pose due to perspective effects as well as the unknown distance between the camera and the license plates being searched. Processing time is kept small thanks to the pre-processing stage:

- The search is limited to the regions R_i previously identified as possibly containing license plates.
- The set of templates used during the search is restricted to those whose size corresponds to the estimate of the license plate at each location in R_i as collected during pre-processing.

The end result is a set of hypotheses along with their confidence, i.e. the degree of match between the template and the underlying pixels in the image.

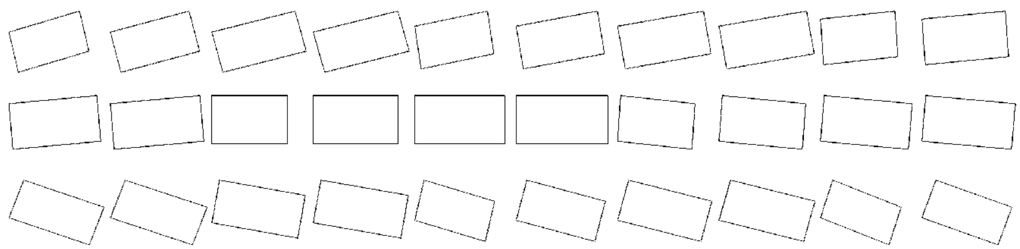

Figure 2 - Collection of license plate templates

3. HYPOTHESIS VALIDATION

Urban scenes are likely to contain linear structures that may negatively affect the search for rectangles, especially if the license plates are expected to be small. In order to reduce the number of false alarms, each hypothesis is validated using a "polygonality criterion". This criterion is based on a symmetry measure proposed in [7] and also used in [8] in the context of road sign recognition. The symmetry measure is the result of a vote by all pixels belonging to the hypothesis, where the voting scheme is designed to yield peaks in the vote image when the symmetry is particularly strong. The process is illustrated in Figure 3 with the contribution of a point *p*: black areas in the vote image receive negative votes while white areas receive positive votes. These areas are defined with respect to the local gradient orientation at *p* as well as the expected object size (width of $4r$ and height of $2r$, because the size of license plates in North-America is standardized to 2:1 width/height ratio). Concretely, given a hypothesis to be validated, pixels belonging to both the template and the underlying edge map take part in the voting scheme:

- A rectangle having width and height of $4r$ and $2r$ respectively accumulates a maximum of votes at its centroid because of the contribution of all the pixels along the boundary.
- Larger or smaller rectangles lead to more than one maximum in the vote image.
- A degraded rectangle leads to smaller peaks in the vote image due to a decrease in the number of voting pixels.
- A noisy cluster of pixels that vaguely look like a rectangle will generate multiple maxima with small amplitude.

Obviously, the higher the peak at the centroid of the hypothetical pattern, the higher the confidence that the pattern matches a license plate.

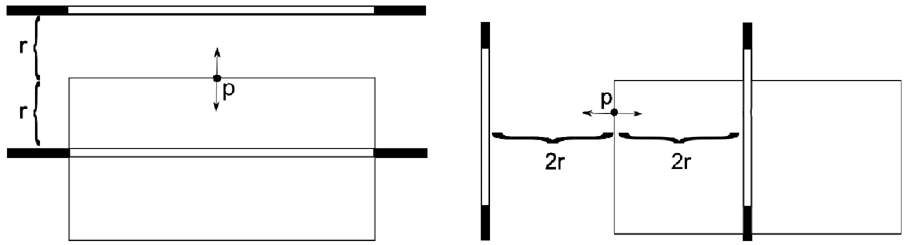

Figure 3 - Voting scheme for symmetry validation

As an example, consider Figure 4 where Hausdorff detection found 10 plate candidates labeled A, B, ..., J. Overlapping hypotheses A and B correspond to one plate while hypothesis J corresponds to the other visible plate in the image. Those three hypotheses get the highest Hausdorff confidences as well as the highest voting results, so both measurements contribute to the decrease in false alarms.

Hypothesis	Hausdorff confidence	Voting result
A	19	408
B	21	419
C	13	102
D	16	65
E	15	82
F	16	35
G	14	75
H	14	104
I	14	56
J	19	145

Figure 4 - Some license plate candidates.

4. RESULTS AND DISCUSSION

In order to evaluate the approach, a dataset of car images made publicly available by the Computational Vision group at Caltech* was used as benchmark dataset. (The availability of this dataset was the primary reason for choosing it.) It contains 126 images of cars taken from the rear, with the license plate visible (the vast majority of cars were registered in California). For the evaluation, images were downsized by a factor of two to simulate low resolution (down to 448x296 pixels). The average plate size was 48x24 pixels after downsizing.

Ground-truth comparison, without hypothesis validation, found that the method correctly located 119 of the 126 license plates (i.e. 94% detection success, with mean overlap accuracy of 77%) while missing 7 plates and committing 297 false alarms (avg. 2.4 false alarms per image). Some results are shown in Figure 6. The distribution of Hausdorff and symmetry confidence measures with respect to good/false detection is shown in Figure 5. Inspection of the plot indicates that using hypothesis validation with thresholds on these confidence measures could eliminate a significant number of false alarms. Indeed, if the thresholds corresponding to dashed lines are used, the detection rate drops slightly to 91% (5 good detections lost) but the number of false alarms falls to 155 (1.2 false alarm per image). A few remarks about the results can be made:

- The number of false alarms can be explained by the complexity of the scenes, with many man-made structures (e.g. building windows); car lights are sometimes viewed as a license plate, which underlines the importance of developing an additional criterion for hypothesis validation, possibly textured-based.

* http://www.vision.caltech.edu/Image_Datasets/cars_markus/cars_markus.tar

- Some plates are missed because of bad image quality, white plate on white car, etc.
- The overlap accuracy is not critical since the final goal is to help control a slave PTZ camera whose purpose is to get a better image of the plate for OCR.

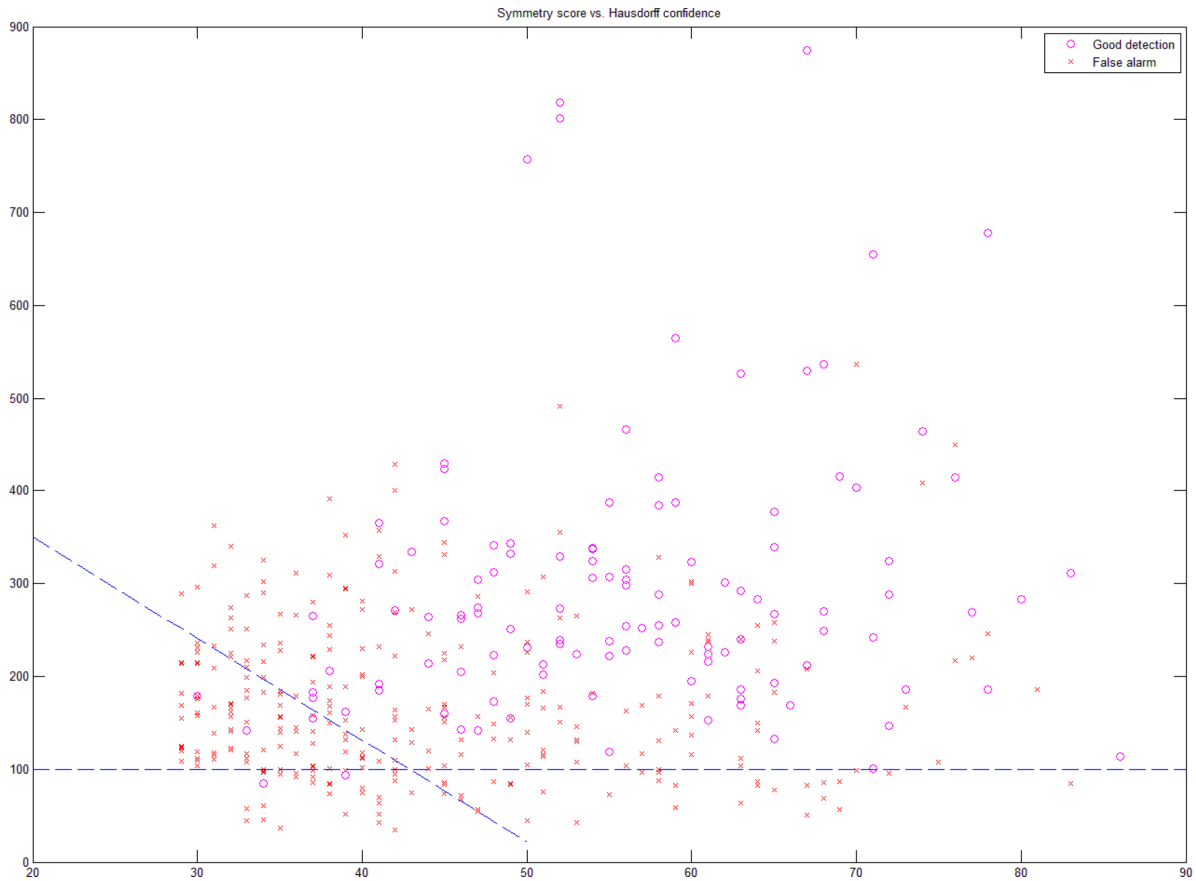

Figure 5 - Symmetry score vs. Hausdorff confidence for good detections and false alarms**.**

As mentioned before, low processing time is an important specification in the context of use of the license plate detection module, and the second stage (region search) of the proposed method plays an important role in the attempt to meet this specification. Without region search, the Hausdorff pattern matching algorithm would have to scan the whole image, and despite the speed-up provided by the use of the distance transform (Voronoi surface) in the evaluation of the Hausdorff distance, the high number of quadrilateral templates required to handle unknown plate size and pose would make computation cost prohibitive. Region search uses simpler templates (short vertical line segments) that are less numerous (no pose information) and contain less pixels, so matching is faster. Also, positive matches provide cues about potential plate size, so the matching pass with the quadrilateral templates of appropriate size can be performed at reasonable speed. Over the 126 images of the dataset, the average processing time is about 0.3 second on a standard PC, broken down approximately as follows: 36 msec. for stage 1, 25 msec. for stage 2, 220 msec. for stage 3 (rectangle matching with Hausdorff distance), 15 msec. for symmetry validation. Note that processing time for stage 3 is scene dependent, i.e. an image with rich horizontal and vertical edges will take longer to process.

One could make the case that the symmetry detector could be run over the whole image and thus pinpoint rectangular areas. Tests have shown that the processing time becomes prohibitive as the number of pixels to be processed increases. Remember also that since the plate size is unknown, many passes with varying *r* would be needed, with an obvious impact on the processing speed (let alone the complexity of isolating the appropriate peaks in the vote image).

Figure 6 - Some results

5. CONCLUSION

A new method for license plate localization based on the Hausdorff distance is proposed. It is capable of locating license plates at low resolution with few misses and a manageable number of false alarms even with perspective distortion, as long as the boundary of the license plate is reasonably well defined. A symmetry criterion helps reject bad plate candidates. Future work includes adding a validation criterion (possibly texture-based) in order to decrease false alarms as well as optimizing Hausdorff matching to decrease computation time.

ACKNOWLEDGEMENTS

This work was supported in part by Precarn Inc. (www.precarn.ca) through the Precarn-CRIM Alliance Financing Program and the “Ministère du Développement Économique de l’Innovation et de l’Exportation (MDEIE)” of “Gouvernement du Québec”.

REFERENCES

[1] Lalonde, M., Foucher, S., Gagnon, L., Pronovost, E., Derenne, M. and Janelle, A. "A System to Automatically Track Humans and Vehicles with a PTZ Camera," Proc. SPIE 6575, (2007).
[2] Foucher, S., Lalonde M., Gagnon L. "Automatic scene modeling for improving object classification," Proc. SPIE 7701 (2010)
[3] Faradji, F., Rezaie, A. H., Ziaratban, M. "A Morphological-Based License Plate Location," Proc. of the Int. Conf. on Image Processing, 57-60 (2007)
[4] Nguyen, C. D., Ardabilian, M., Chen, L. "Real-time license plate localization based on a new scale and rotation invariant texture descriptor," Proc. 11th Int. IEEE Conf. on Intelligent Transportation Systems (ITSC), (2008)
[5] Anagnostopoulos, C. N., Anagnostopoulos, I., Psoroulas, I. D., Loumos, V., Kayafas, E. "License Plate Recognition From Still Images and Video Sequences: A Survey," IEEE Trans. on Intelligent Transportation Systems, 9(3), 377-391 (2008)
[6] Huttenlocher, D., Klanderman, G., Rucklidge, W. "Comparing Images Using the Hausdorff Distance," IEEE Trans. PAMI, 15(9), 850-863 (1993)
[7] Loy, G., Barnes, N. "Fast Shape-based Road Sign Detection for a Driver Assistance System," Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS2004), 70-75 (2004)
[8] Keller, C. G., Sprunk, C., Bahlmann, C., Giebel, J., Baratoff, G. "Real-time recognition of U.S. speed signs," Proc. IEEE Intelligent Vehicles Symposium, 518-523 (2008)
[9] Donoser, M., Arth, C., Bischof, H. "Detecting, Tracking and Recognizing License Plates," Proc. Asian Conf. Computer Vision, 447-456 (2007)
[10] Zhang, H., Jia, W., He, X., Wu, Q. "Learning-Based License Plate Detection Using Global and Local Features," Proc ICPR,2, 1102-1105 (2006)