Blind License Plate Detection

Max Guise
Department of Electrical Engineering
Stanford University
Stanford, United States
mguise@stanford.edu

Stephan Goupille
Department of Electrical Engineering
Stanford University
Stanford, United States
stephang@stanford.edu

***Abstract*—In this paper we explore a method for license plate detection in the presence of background objects such as trees and buildings. The initial stages are edge-based methods that utilize a matched filter, while the subsequent false-positive rejection and region refinement stages rely on mean-shift segmentation for color-aided performance enhancement.**

***Index Terms*—**License plate detection, mean-shift segmentation

I. Introduction

Commercial and government automatic license plate recognition (ALPR) systems have been used reliably in a variety of applications. Though the accuracy of typical systems is often quite high, the processing pipeline usually relies on several important factors to produce accurate results. In particular, minimal affine transformation of the plate, a powerful flash, lack of background clutter, and relatively short distances from the camera to the target all significantly aid the detection process. Moreover, the core technique of many ALPR systems is often a classification algorithm dependent on a large database and significant initial training overhead.

The goal of this work is to explore the detection problem without using a classification algorithm (i.e. no prior model of the plates, hence "blind" detection) in a set of images that include background objects such as trees, buildings, street lights, etc. Techniques from several of the referenced papers are modified and combined to produce the final detection system.

II. Related Work

As noted above, this paper concerns the license plate detection problem without the use of a classifier, so only previous works that do not include such a stage shall be discussed. A number of blind methods, either edge-based or color-based, have implemented edge-based and color-based methods for plate detection. Reference [4] outlines a solely edge-density-based method that, while reasonably effective, most certainly relied on a controlled imaging scenario.

Reference [6] approached the problem from an initially color-based standpoint. In particular, the authors first perform a mean-shift segmentation, and subsequently use edge-density and rectangularity criteria to choose license plate candidates from the original segmentation output regions. This is in sharp contrast to Abolghaesmi and Ahmadyfard's work, in which a matched filtering technique is employed to extract candidate regions.

In the following section, our modifications to and combinations of the above approaches shall be described, and its benefits and pitfalls will be discussed in Section IV.

III. Description of Methods

A. *Preprocessing*

To speed the implementation, the image is first downsampled by a factor of 2, and subsequent edge detection and filtering performed at this scale. In addition, several of the previous works discussed above rely on very specific functions for image enhancement, and perhaps with good reason. Here our only preprocessing stage is an optional adaptive histogram equalization step, the results of which will be discussed in Section IV. In our implementation, the dimensions of the tiles for the adaptive equalization are chosen to match those of a typical plate in the target images.

B. *Vertical Edge Detection and Matched Filtering*

A key insight into the license plate detection problem is that the presence of a plate results in *vertical* edges in a rectangular region of a particular aspect ratio [5]. As a result, as is done in [5], the first stage in the detection pipeline consists of taking a simple centered horizontal image gradient, followed by a 21 by 50 pixel Gaussian weighting with standard deviation equal to 10 pixels. An example of a weighted edge map obtained from a car image is shown in Figure 1 below.

Figure 1: Original image and blurred vertical edge map.

The next detection stage consists of choosing candidate regions from the edge map obtained above. Reference [5] accomplishes this with the filter shown below in Figure 2.

Such a filter accomplishes the task reliably in the situation in which it was used: detecting Iranian license plates. Virtually none of these plates were surrounded by license plate holders, while a majority of our test images were California plates surrounded by license plate holders from dealerships, schools, etc. Edges detected in the wording on these license plate holders could be picked up by the extended negative

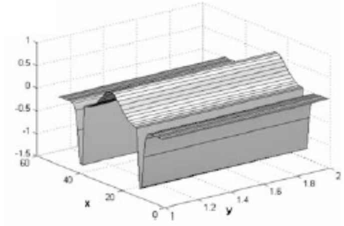

Figure 2: Filter used in Reference [5].

lobes in the filter shown in Figure 2. For this reason, we use instead the filter shown in Figure 3 below and given by:

$$g[x,y]=\begin{cases} -e^{\left(\frac{-(y-5)^2}{100}\right)} & y\in[1,9] \\ -e^{\left(\frac{-(y-36)^2}{100}\right)} & y\in[32,40] \\ e^{\left(\frac{-(y-20.5)^2}{100}\right)} & y\in[10,31] \end{cases}$$

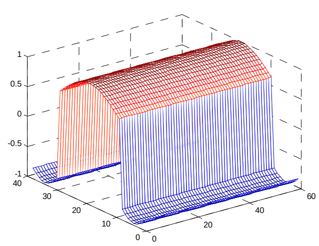

Figure 3: Filter used in core detection stage.

The output of the filtering stage is shown below in Figure 4. We threshold this result, using a value chosen experimentally to be any pixels that fall into the 99.5 percentile or higher. This choice often results in a large number of candidate regions, but is necessary largely due to the presence of multiple license plates and, at times, smaller plates with less well-defined edges.

After thresholding, region labeling is performed. We extract rectangular candidate regions from the returned set, centered at the mean of each cluster. The region dimensions are chosen proportional to the width of the cluster (a margin is provided for slightly distorted plates).

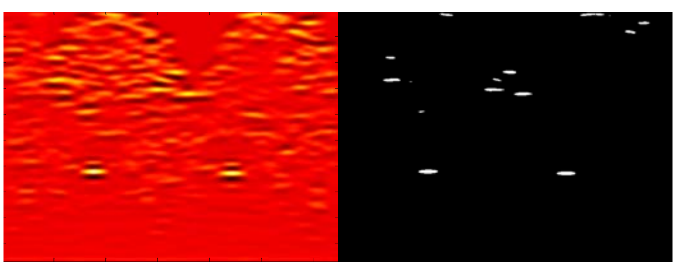

Figure 4: Filtering output before and after thresholding.

C. False Positive Rejection

As shown in Figure 4, simply returning all of the candidate regions output by the filtering stage would not be an acceptable detector. We thus require a method of rejecting the regions that likely do not contain license plates. We considered discarding small regions, or performing morphological erosion with a rectangular structuring element of the approximate aspect ratio of a license plate and then discarding small regions. Though these methods did drop many of the false positives, tuning the parameters to avoid also discarding license plate regions proved rather difficult across our set of images. Consequently, we employ two methods, described below.

1) "Contextual" Rejection

Although it is clearly possible for a license plate to be present in an image without a car, plates in the situations in which we are interested in performing license plate detection are usually attached to vehicles. Thus, we can reject false positive candidate windows by searching for characteristics of cars in the local region. However, even though various car detection methods have been successful (e.g. Histograms of Oriented Gradients), they typically use a training database and classifier.

We suggest the very simple approach of using the Hough transform to search above and below the candidate window for prominent horizontal lines, which, for an actual license plate region, often result from the edges of the plate itself, the license plate holder, or various portions of the bumper. This last distinction is important because the edges of the plate itself are often not particularly well defined in the edge map; dirt, glare, license plate holder and/or bumper color can all severely interfere with the detection of the plate border. This contextual rejection is most effective for image patches that contain, for example, tree branches or leaves against a sky or building background. In many cases, the branches result in a rectangular region of dense edge pixels that can easily be rejected via this method.

In our implementation, the outer pixels of the original candidate rectangle output by the matched filtering stage, as well as a small margin around it, are searched for horizontal lines of at least half the length of the original window.

2) Mean-Shift Segmentation Based Rejection

Because the license plate region consists mostly of a solid color background, we'd like to segment the image region containing a plate candidate window in order to reject additional false positives. An algorithm such as *k*-means is not suitable because of the vast differences in output based on the choice of *k*, as well as the fact that at a minimum, the lettering on the plate will be segmented separately from the plate background.

The mean-shift algorithm, however, detects modes of a dataset, which, for images, can be done both in the spatial and range domains. In particular, this allows for a spatial- and color-based segmentation, which will achieve results much closer to the desired output. If it were perfect, however, we could have just used this algorithm as the basis for the detector, as in [6].

Instead, we utilize the algorithm to further reject windows that passed through the previous stage. To achieve a good segmentation result, we set the minimum region area parameter proportional to the region size selected by the filtering and thresholding stage. This often results in a favorable segmentation output that can be used for false positive rejection, or even region refinement. These two applications, however, were not fully explored as more development is needed in the analysis of the segmented regions. Currently, the image is downsampled by a factor of 8, mean-shift is segmentation is performed, followed by a morphological erosion is performed using a 2 by 5 (height x width) structuring element, and then regions that are sufficiently small or large are rejected. Additionally, if the rectangle enclosing all of the region's pixels is sufficiently far from the aspect ratio of a plate, the region is rejected. The parameters were chosen experimentally. Note also that the regions needed to be upsampled to return a region in the units of the original image.

Figure 5: An example of excellent plate separation by mean-shift segmentation. Note that without prior knowledge of candidate region size, the segmentation would not produce such a desirable result in many cases.

IV. EXPERIMENTAL RESULTS

The recognition rate for plates whose original width is 100 pixels or larger was 91.7% on our dataset, although it should be noted that the false positive ratio was an excessive 30% in this case. This latter fact is due to the addition to the dataset of a number of images that contained business signs in the background (taken in the parking lot of a large outdoor mall area). For images without signs in the background, approximately 5-10% of the detected windows were false positives. When smaller plates were included, the overall recognition rate dropped to 67.6%.

V. CONCLUSIONS

Blind, edge-based license plate detection is a difficult problem that requires specific imaging conditions to perform accurately and reliably. The presence of background clutter must be mitigated with some form of false positive rejection. The proposed methods are not stringent enough, seeing as the false positive rate is far higher than would be acceptable in any practical application. Although the horizontal-line search method worked well in parts of our image set, mean-shift segmentation seems most promising as a robust rejection stage. However, the implementation details of this latter method are a major area for potential improvement; proper analysis of the regions as in [6] could result in elimination of the horizontal-line-based false positive rejection stage.

Additionally, the scale-invariance of the proposed methods could be improved. Though we are not particularly interested in plates that are too small to perform reliable optical character recognition (OCR), there is nonetheless significant room for improvement in the scale-invariance of the proposed method for plates well within the reach of OCR.

More importantly, however, is the detector's susceptibility to affine distortions of the plates. Distortions resulting in less than an approximately 20° rotation of the top and bottom edges of the plate do not seem to interfere with the proposed method, but it seems that any greater distortion results in a 0% detection rate. Future work might experiment with rotated edge detection and filters, although an implementation faster than Matlab would likely be desirable at that point.

Another attempt at building the detector would likely focus more on the capabilities of the mean-shift algorithm and perhaps initially using clues from the image to choose appropriate segmentation parameters. Such a second attempt at the overall process of improving detector performance would most definitely be done in a more systematic fashion to alleviate the difficulties that come with the problem.

REFERENCES

[1] D. Comanicu, P. Meer, "Mean shift: A robust approach toward feature space analysis," IEEE Trans. Pattern Anal. Machine Intell., May 2002.

[2] P. Meer, B. Georgescu, "Edge detection with embedded confidence," IEEE Trans. Pattern Anal. Machine Intell., 28, 2001.

[3] C. Christoudias, B. Georgescu, P. Meer, "Synergism in low level vision," 16th International Conference of Pattern Recognition, Track 1 - Computer Vision and Robotics, Quebec City, Canada, August 2001.

[4] D. Zheng, Y. Zhao, and J. Wang, "An efficient method of license plate location," Pattern Recognition Letters, vol. 26, pp. 2431-2438, 2005.

[5] V. Abolghasemi and A. Ahmadyfard, "An edge-based color-aided method for license plate detection," Image and Vision Computing, vol. 27, pp. 1134-1142, 2009.

[6] W. Jia, H. Zhang, and X. He, "Region-based license plate detection," Journal of Network and Computer Applications, vol. 30, pp. 1324-1333, 2007.

APPENDIX

Note that Stephan is only enrolled in Psych 221; Max is enrolled in both 221 and 368 and was thus responsible for a larger share of the project work.

A great deal of work went into several endeavors which are not reflected in the official writeup because they were unfortunately ineffective detection methods, and thus didn't fit into the template for the report. Listed below are contributions to the project.

- Both group members spent time collecting images.
- Stephan experimented a good deal with kmeans segmentation, even though we later opted for mean-shift segmentation.
- Both group members experimented with using the CIELAB color space and several color balancing algorithms as an aid to the first detection stage, although this was replaced with the mean-shift final stage.
- Max experimented extensively with using the Hough transform as the first processing stage, as well as somewhat less time-consuming attempts at using Harris corners or Histograms of Oriented Gradients as the initial step.
- Both Max and Stephan experimented with (adaptive) histogram equalization for detection performance improvement.
- Max set up the Android SDK and an application that displays the viewfinder, takes a picture, and uploads it to a webserver (also set up the webserver). We unfortunately didn't get to use this because the difficulties with detection required our attention.
- Max coded most of the Matlab source used in the final implementation, but both Max and Stephan experimented with parameter choices. Max also coded the infrastructure for labeling the ground truth in the data images and running tests.
- Both group members worked on the Psych 221 presentation, but Max was responsible for the 368 poster and report.