CASCADED CLASSIFIER FOR LICENSE PLATE DETECTION

Belarusian State University of Informatics and Radioelectronics
Minsk, Republic of Belarus

Losiukov L.N.

Volkov K.A. – PhD, docent

License Plate Recognition (LPR) has found numerous applications in various areas. It can be used for automatically identifying vehicles in a car park, for vehicle access control in a restricted area and for detecting and verifying stolen vehicles. A LPR system consists of two major components: license plate detection and character recognition. The first step of license plate detection is classifier training at which a six-layer cascade classifier is constructed.

As shown in [1], the basic idea of the detection algorithm is to use a variable scanning window moving around on the input vehicle image. At each position, the image area covered by the scanning window is classified using a pre-trained classifier as either a license-plate area (a positive decision) or a non-license-plate area (a negative decision). The classifier used in this algorithm is a significant extension of Viola and Jones' work shown in [2] to license plate detection.

In this algorithm, a six-layer cascaded classifier is constructed to increase the detection speed, in which the first two layers are based on global features and the last four layers are based on local Haar-like features. The classification process can be taken as a degenerate decision tree containing multi-layer classifiers as shown in Fig. 1. A positive result from the first classifier triggers the evaluation of a second classifier. A positive result from the second classifier triggers a third classifier, and so on. A negative outcome at any layer leads to the immediate rejection of the image region (block). In other words, those image regions that are not rejected by the initial classifier will be processed by a sequence of classifiers. If any classifier rejects a selected image region, no further processing will be performed for this region. It is commonly seen that, for a given vehicle image, the majority of evaluated image regions are negative. Therefore, the cascaded classifier shown in Fig. 1 attempts to reject as many negatives as possible at the earlier stages. As its consequence, this cascaded classifier leads to fast license plate detection. It is also worth to note that the detection algorithm (classifier) acts on the vertical edge maps of input vehicle images rather than on the raw image intensity values. This further enhances the efficiency of the cascaded classifier.

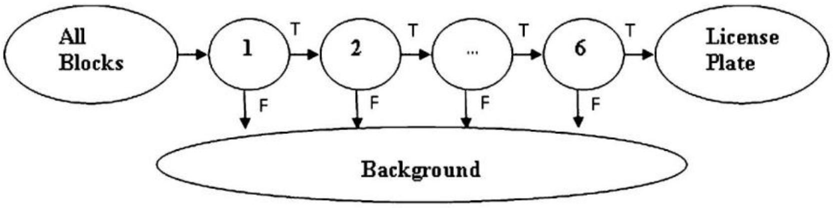

Fig. 1 - Process of constructing a cascaded classifier

In the following, the algorithm is described in two aspects: training and testing. Training is a process that the classifier is learning to make correct decisions using pre-classified samples. Testing is a process to use the pre-trained classifier to classify individual image blocks.

To obtain the cascaded classifier which can make correct decisions, pre-classified positive samples (images containing license plates) and negative samples (images containing non-number-plates) are selected for training.

The individual classifiers that together construct the cascaded classifier are trained independently. Recall that the classifiers on the first two layers are based on global features, and the classifiers on the other four layers are based on local Haar-like features.

To train the classifier on the first layer, the value of the first global feature, called Edge Density, is computed for each input sample. Statistical methods are used to select a threshold that can correctly classify all positive samples (i.e., using the selected threshold, the edge densities of all positive samples are on the "positive" side). Note that the threshold is not unique. Also note that, for a given threshold, some negative samples may be wrongly classified as "positive", i.e., some non-number-plate blocks may be classified as a "license plate". These are referred as false positives. Hence, a threshold which can correctly classify all positive samples and produce the least number of false positives is selected.

For the second layer classifier, another global feature is employed. It is called the Edge Density Variance. All input samples used to train the classifier on the second layer are from "positive" classification outcomes using the first classifier after training. Since the false positive rate of the first classifier is usually non-zero, samples which are classified as "positive" by the first classifier contain both real positive samples and some negative samples. By properly selecting another threshold based on the second global feature, we can classify all positive samples as "positive" and produce minimum false positives. The classifier on the second layer based on the second global feature is thus obtained. Again, the training of the second classifier is implemented using statistical methods similar to those used for the first classifier.

Similarly, the samples used to train the classifier on the third layer are those samples which are filtered as "positive" by the classifiers on the first two layers. Unlike the first two layers, which are both based on global features, the classifiers on the third through sixth layers are all based on local Haar-like features. We will find that, within any image area (region), the total number of Haar-like features is very large and much larger than the total number of pixels within the area. To ensure fast classification, the AdaBoost learning algorithm is used to select best-performing classifiers (called weak classifiers), each based on a Haar-like feature, and to combine these multiple weak classifiers to construct one classifier (referred as strong classifier). This procedure is completed in multiple rounds. In each round, an optimal weak classifier is selected. The AdaBoost algorithm introduces a weight for each sample. Through continuously increasing the weights of “hard” samples in each round and selecting the corresponding best-performing weak classifiers until the constructed strong classifier meets the predefined accuracy requirement, a strong classifier is then constructed and the training on this layer is finished.

Similarly, the samples classified as positive ones by the third layer are input to the fourth layer, and so forth. Finally, a six-layer cascade classifier is constructed.

Since both global and local features in this algorithm are generated from vertical edge maps of input images, the vertical edge maps of images are computed first before any feature is extracted.

References:

1. Huaifeng Zhang, Wenjing Jia, Xiangjian He, Qiang Wu, Learning-based license plate detection using global and local features, in: Proceedings of the 18th International Conference on Pattern Recognition (ICPR 2006), vol. 2, 2006, pp. 1102-1105.

2. Paul Viola, Michael J. Jones, Robust real-time face detection, Int. J. Comput. Vis. 57 (2004) 137-154, http://dx.doi.org/10.1023/B:VISI.0000013087.49260.fb.