AllGoVision Achieves high Performance Optimization for its ANPR Solution with OpenVINO™ Toolkit

Version 3.9

Migrating to OpenVINO™ framework help its deep learning Automatic Number Plate Recognition (ANPR) model achieve high accuracy in real time processing of up to 3 lanes in one core of CPU

Revision Date: November 2018

AllGoVision Technologies Pvt Ltd

Email: contact@allgovision.com
Website: www.allgovision.com

Contents

COPYRIGHT INFORMATION 2
INTRODUCTION 3
Challenging Scenario for ANPR 3
Solution - DL based ANPR from AllGoVision 4
Optimizing the DL Model 4
Proven Solution for the World on a CPU-Based System 5

COPYRIGHT INFORMATION

Introduction

License plate recognition poses many challenges, especially in places like India where the compliance with license plate specification is very low. The difficulties are not just in recognition, but also in detection and localization. AllGoVision has developed a solution that solves the problem effectively using Deep Learning.

The network architectures have been optimized for efficiency. Further, by using OpenVINO™ toolkit, the networks have been optimized to for Intel platforms with AVX2 that yielded about 4X overall improvement in speed. This helped to achieve high accuracy and real time performance using a single core of the CPU for monitoring multiple lanes of moderately fast traffic.

Challenging Scenario for ANPR

Licence plate recognition systems need to be able to deal with several challenges on the ground, especially in places like India. The main difficulties are:

1. Non-standard license plate: This includes deviations in layout of characters and sizes, fonts, plate type, plate location and dirt on the plate. This is illustrated in Figure 2, which shows the specification for 1 and 2 row plates on the left, and commonly encountered deviant plates on the right. These deviant plates can constitute a significant fraction of vehicles , often over 25%, especially in lower tier towns.
2. Plates are often not distinct on the vehicle: Sometimes the numbers are just hand-painted on the back of the vehicles such as auto-rickshaws (3-wheelers). The plates also co-exist with other decorations and banners, which makes it difficult to detect the plates
3. Lack of lane discipline: Vehicles frequently change lines thus presenting unfavourable angles to the camera.
4. Dense traffic: This causes partial occlusion of the plates by neighbouring vehicles leading to fewer observations of the full plate

Figure 2: Plate variations

An effective solution must detect nearly all license plates and achieve a high degree of accuracy for all the variations that are encountered.

Solution - DL based ANPR from AllGoVision

AllGoVision ANPR solution uses custom built deep CNN network with detector, localizer and recognizer.

Plate detection scans the region of interest in the input image. The network is tuned for high sensitivity to detect plates even if they are in unexpected locations, or if they do not stand out from the background motif. The side effect of this that there are several false plates detected. Next, the detected plates are localized with fine resolution. Then the attributes of the block such as type, colour and orientation are determined. After correcting for any perspective change, the characters are recognized.

False plates increase the computational burden of the system. They are eliminated at each stage after detection using various criteria to minimize the overall computational complexity while still retaining the of detection capability.

Multiple observations of the vehicle and plate tracking are used to increase the accuracy of the result.

Our solution is now able to detect over 98% of the plates and achieve around 95% accuracy for recognition among legible plates (plates where humans can unambiguously read the plate with appropriate zoom level)

Optimizing the DL Model

All deep networks are all CNNs with 4 to 7 layers of convolution filters followed by fully connected layers. The initial implementation was based on a GPU using the Caffe deep learning framework. Since system complexity was of paramount importance in our application, our objective was to achieve a solution that could work on low end computers without a GPU. Towards this, a series of optimizations were carried out to reduce the complexity of the deep networks while still retaining the minimum accuracy that was needed. We achieved about 5X reduction in network complexity. With this, we could run our solution using only the CPU, but we needed 3 cores per channel. Our objective was to be able to run the ANPR processing chain on core of a CPU in real time. This is when we migrated our models to OpenVINO™. There were two methods that helped us achieve the required speed-up:

1. Model optimization: The OpenVINO™ model optimizer yielded a 3X improvement in the inferencing time of the deep networks. The summary of each of the blocks is shown in Figure 4.

2. Batching: Increasing the batch size improves the average time for inferencing per image. We used batching for localization and recognition. Batching for detection inferencing has been deferred to a later time since this required an architecture change in our software. It has been planned for a later release. The time improvement due to batching for recognition of single row plates is shown in Figure 3. This significantly reduced the computational burden due to false plates

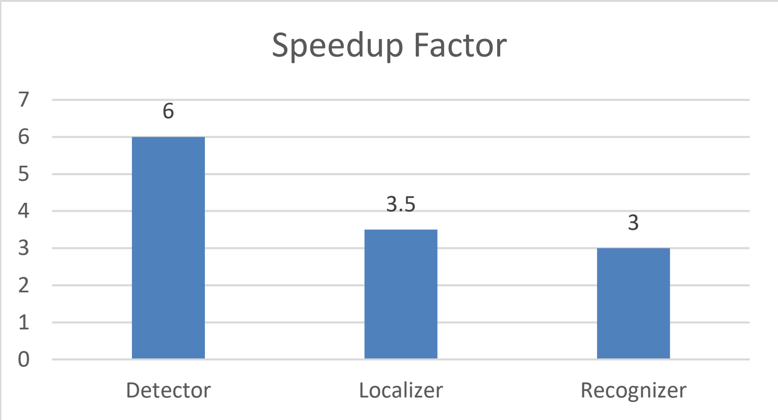

Figure 4: Inference time speedup due to OpenVINO™ model optimizer

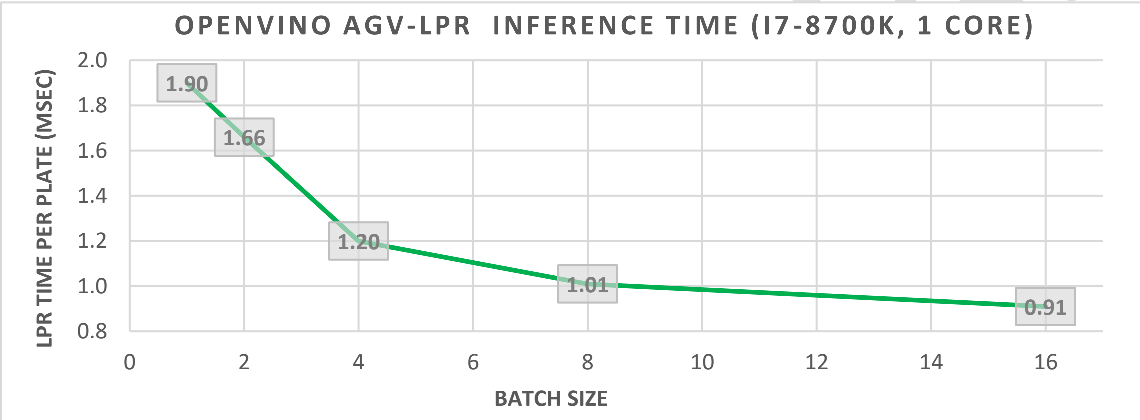

Figure 3: Per plate Inferencing time vs. batch size

Based on the optimizations above, we are able to process a 1080p image with about 4 license plates in the field of view in about 40 milliseconds on one core of an i7 8700K processor. On a Xeon processor running at 2.4 GHz, the processing time is about 60 ms using one core. With allowance for decoding time and some head room, we can comfortably process a 1080p video at 10 fps covering 2-3 lanes of traffic. For fewer lanes, and 720p, higher frame rates can be supported to cater to high speed traffic. For other countries, processing time is lower because of a combination of one or more of: easier detection, lower complexity in the networks, fewer plate types, and fewer number of digits in the plate.

Proven Solution for the World on a CPU-Based System

AllGoVision used the latest advances in Deep Learning and the OpenVINO™ framework to achieve high functional performance at low complexity. We can now perform the ANPR processing chain in real time using just one core of a standard desktop CPU. The solution has been extended to several countries including Mexico, South Africa, Malaysia, Algeria and Turkey. In each case, the accuracy is at least as good as for Indian plates. We are able to add support for any new country within a short amount of time of a couple of weeks. In the future, we plan to take advantage of new architectural advances from Intel including AVX-512 and FPGA and further reduce the number of cores per channel to provide a robust and cost-effective solution to our customers