License Plate Recognition using Convolutional Neural Networks Trained on Synthetic Images

Abstract

In this thesis, we propose a license plate recognition system and study the feasibility of using synthetic training samples to train convolutional neural networks for a practical application.

First we develop a modular framework for synthetic license plate generation; to generate different license plate types (or other objects) only the first module needs to be adapted. The other modules apply variations to the training samples such as background, occlusions, camera perspective projection, object noise and camera acquisition noise, with the aim to achieve enough variation of the object that the trained networks will also recognize real objects of the same class.

Then we design two convolutional neural networks of low-complexity for license plate detection and character recognition. Both are designed for simultaneous classification and localization by branching the networks into a classification and a regression branch and are trained end-to-end simultaneously over both branches, on only our synthetic training samples.

To recognize real license plates, we design a pipeline for scale invariant license plate detection with a scale pyramid and a fully convolutional application of the license plate detection network in order to detect any number of license plates and of any scale in an image. Before character classification is applied, potential plate regions are un-skewed based on the detected plate location in order to achieve an as optimal representation of the characters as possible. The character classification is also performed with a fully convolutional sweep to simultaneously find all characters at once.

Both the plate and the character stages apply a refinement classification where initial classifications are first centered and rescaled. We show that this simple, yet effective trick greatly improves the accuracy of our classifications, and at a small increase of complexity. To our knowledge, this trick has not been exploited before.

To show the effectiveness of our system we first apply it on a dataset of photos of Italian license plates to evaluate the different stages of our system and which effect the classification thresholds have on the accuracy. We also find robust training parameters and thresholds that are reliable for classification without any need for calibration on a validation set of real annotated samples (which may not always be available) and achieve a balanced precision and recall on the set of Italian license plates, both in excess of 98%.

Finally, to show that our system generalizes to new plate types, we compare our system to two reference system on a dataset of Taiwanese license plates. For this, we only modify the first module of the synthetic plate generation algorithm to produce Taiwanese license plates and adjust parameters regarding plate dimensions, then we train our networks and apply the classification pipeline, using the robust parameters, on the Taiwanese reference dataset. We achieve state-of-the-art performance on plate detection (99.86% precision and 99.1% recall), single character detection (99.6%) and full license reading (98.7%).