Supplemental material for "A CNN-Based Approach for Automatic License Plate Recognition in the Wild"

BMVC 2017 Submission # 176

1 Visualized Detection Results

Fig.1 shows some qualitative results on license plate detection. The predicted corner positions for each detected license plate are marked with red points. Our detector performs surprisingly well in some of the hard cases, where the plates are small, tilted, blurred, in weak light or even truncated.

2 Visualized Recognition Results

We compare our recognizer with the state-of-the-art recognizers [1, 2], both of which take the original license plate patches without rectification as input while ours use rectified patches. For fair comparison, we supplement two experiments, in which all three schemes use the same rectified plates as input. Table 1 shows the results. The performances of both reference recognizers are significantly improved. However, our method still performs the best, demonstrating the benefit of STN and shared-weight classifier.

Table 1: Comparison of recognition performance. The added results are flagged with "*".

	Acc	Acc@1e	Acc@2e	CER
[1]	0.7425	0.9036	0.9615	0.0614
[2]	0.3187	0.6743	0.8205	0.1961
ours	**0.8905**	**0.9510**	**0.9755**	**0.0314**
[1]*	0.8599	0.9475	0.9720	0.0362
[2]*	0.5259	0.7925	0.8958	0.1232

In Fig. 2, we provide eight sets of visualized recognition. The first column contains original patches cropped with bounding boxes, and the second column contains the rectified plates. The third column compares the results of our scheme and two versions of method [1]. Comparing the results of [1] and [1]*, we can find that proper rectification eases the recognition task, e.g. it is easier for the recognizer to differentiate between "4" and "A". Comparing the results of [1] and ours, we see that our recognizer performs better on both Chinese characters and uppercase letters. This is because we have a recognizer dedicated for 31 Chinese characters, and the other shared-weight recognizers make full use of the limited

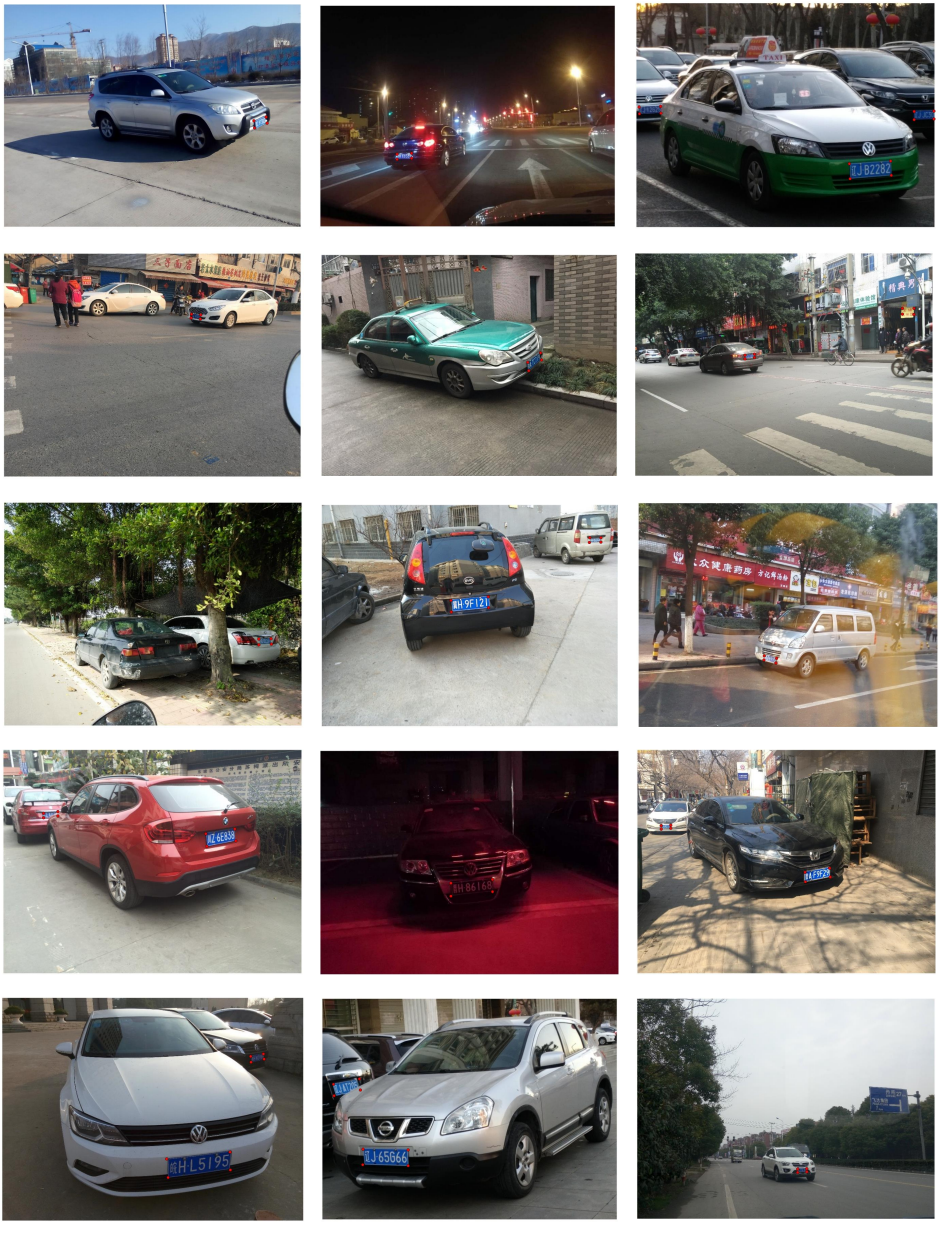

Figure 1: Visualized detection results

training data. As the uppercase letters are rarer than numerals, uppercase letters benefit more from our design.

Figure 2: Recognition results: (a) Original license plate patches (b) Rectified license plate patches (c) The groundtruth labeling, the results of [1] (taking images in (a) as inputs), [1]* (taking images in (b) as inputs) and our recognizer.

3 Failure Cases

Looking into the failure cases, we find two main causes. One is inaccurate corner prediction due to complex scenes. The other is very poor image quality. Fig. 3 shows a few failure cases. The left two columns show the original and rectified plates from inaccurate detection.

The right two columns show the plates which are too blurred to read.

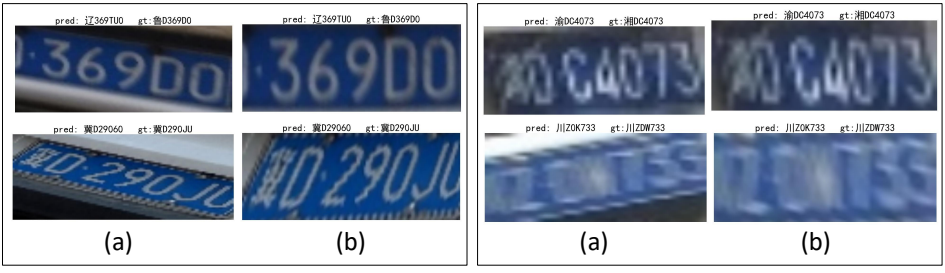

Figure 3: Visualized failure cases (a) Original patches (b) Rectified results. The "pred" are all predicted by our recognizer

References

[1] Max Jaderberg, Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Synthetic data and artificial neural networks for natural scene text recognition. *arXiv preprint arXiv:1406.2227*, 2014.

[2] Hui Li and Chunhua Shen. Reading car license plates using deep convolutional neural networks and lstms. *arXiv preprint arXiv:1601.05610*, 2016.