

Performance of Automatic License Plate Recognition Systems on Distorted Images

Nikola Plavac, Seyed Ali Amirshahi, Marius Pedersen, and Sophie Triantaphillidou

Colourlab, Norwegian University of Science and Technology (NTNU)

E-mail: nikola.plavac@ntnu.no

Abstract. Automatic License Plate Recognition (ALPR) systems are essential for various applications, including law enforcement, traffic management, and access control. However, their performance can be significantly affected by image distortion in adverse environmental conditions and the imaging pipeline. Three different ALPR systems were used to evaluate their robustness to different distortions using images from six well-known ALPR datasets. Two groups of distortions were the focus of our study: simulated weather conditions (rain, brightness, fog, frost, and snow), and modeled camera read noise in the simulated imaging pipeline. Results indicate that certain weather distortions drastically reduced the accuracy of ALPR systems, with the accuracy of the systems approaching zero in some cases. Read noise also negatively impacted performance, even at minimal levels. The sensitivity to the introduced distortions varied between different models and datasets. The results underscore the need for robust ALPR system designs that can handle diverse and challenging capturing conditions.

Keywords: automatic license plate recognition, weather distortions, camera read noise

© 2024 Society for Imaging Science and Technology.

[DOI: 10.2352/J.ImagingSci.Technol.2024.68.6.060401]

1. INTRODUCTION AND RELATED WORKS

Automatic License Plate Recognition (ALPR) systems are computer vision systems designed to read license plate symbols automatically [1]. ALPR systems have been a research topic of interest for more than four decades, with the first prototype introduced in 1979 by the UK Police Scientific Development Branch [2]. These systems automate vehicle identification and monitoring for various applications, including parking management, automatic toll collection, law enforcement, border control, and traffic management. ALPR systems reduce the need for manual toll booths, enhance traffic efficiency, improve security and convenience in parking areas and for road users in general.

Typically, ALPR systems consist of three main components: a license plate detection module, a character segmentation module, and a character recognition module [2]. The license plate detection module identifies the region in an image that contains the license plate. The character segmentation module segments each character on the detected license plate. Finally, the character recognition module interprets the segmented characters.

The initial ALPR systems relied on classical image processing methods, which were heavily based on hand-crafted features and rule-based algorithms. Techniques such as edge detection, morphological operations, and template matching were commonly used for plate detection and character recognition and relied on features such as global image information, texture, colour, and character features [3]. To achieve good performance, classical methods required precisely tailored parameters and controlled conditions, and often struggled with variations in lighting, acquisition angle, and plate occlusion [4, 5]. The evolution in the field of deep learning (DL) in the last decade has revolutionised ALPR systems by leveraging the power of convolutional neural networks (CNNs) and other DL architectures. Most modern ALPR systems [4, 6, 7] use end-to-end DL models that integrate all components of the pipeline into a single unified architecture. They handle complex backgrounds, diverse license plate formats, and varying illumination conditions more effectively than classical methods. Models such as You Only Look Once (YOLO) [8], Single Shot MultiBox Detector (SSD) [9], Faster R-CNN [10], and Spatial Transformer Network (STN) [11] have been successfully applied in ALPR systems for license plate detection [6, 12–15]. To complete the ALPR pipeline, the license plate detection component is followed by an optical character recognition (OCR) component. Both template matching-based OCR systems [16] and DL-based OCR systems [17, 18] have been used for this purpose. These models have significantly improved accuracy, reliability, and adaptability of ALPR systems to real-world conditions.

ALPR datasets [19–21] which have been used to develop ALPR systems contain images of good technical quality. However, adverse weather conditions, variable illumination during capture, and distortions within the imaging pipeline, are presently the challenges in the field of ALPR [22]. Understanding how they impact the performance of ALPR systems is essential for developing robust systems. Geirhos et al. [23] demonstrated that deep neural networks tend to show poor robustness to image distortions compared to humans in classification tasks. The authors also noted that, although deep neural networks surpass human performance when trained on distorted images for the same distortion type, they tend to show poor generalisation in the case of other distortion types. Michaelis et al. [24] highlighted how different distortions influence the performance of

Received Aug. 15, 2024; accepted for publication Dec. 4, 2024; published online Dec. 23, 2024. Associate Editor: Steven Simske.

1062-3701/2024/68(6)/060401/16/\$25.00

state-of-the-art (SOTA) object detection networks, showing that reduction in performance ranged from 30% to 60%, compared to the initial performance, for all simulated distortion types. The main reasons why fully autonomous vehicles are not widely used are attributed to poor robustness and generalisation performance under various weather and illumination conditions [25] (i.e. car camera sensors are unreliable at night or in bad weather, and current computer vision systems do not handle such situations well enough). Beyond distortions introduced by weather conditions, the main attribute in the camera pipeline affecting the performance of CNN models is noise. A recent study by Yim and Sohn [26] has shown that camera noise leads to significantly decreased network performance. Various strategies have been proposed to deal with camera noise in images [27–30].

The impact of weather distortions and camera noise on DL architectures has been studied for relatively simple computer vision tasks, such as classification and object detection. Few published works have investigated how these influence ALPR systems. Rio-Alvarez et al. [31] studied the robustness of license plate detection systems when images were distorted by weather conditions, but the dataset was not public and focused only on night captures and rain. It was reported that including challenging weather conditions in the training sets did not improve the accuracy of the license plate detection systems in images affected by these weather conditions. Conversely, the study indicated that including images affected by low illumination significantly increased the accuracy of the systems. In another study by Spanhel et al. [32], a CNN architecture was proposed for holistic license plate recognition, and it was demonstrated that avoiding segmentation of the license plate characters can increase performance when dealing with low-quality and noisy license plate images. Xu et al. [33] proposed a CRNN-based method for performing ALPR on ships. In addition to low spatial resolution and tilted acquisition angle, ship license plates are often affected by foggy weather. The proposed method combined data augmentation and image enhancement to remove fog effects and resample license plate images.

The aim of this study is to evaluate the performance of DL-based ALPR systems on distorted images. Various distortions, including simulated bad weather conditions and camera read noise, are applied across six different datasets. Through systematic assessment, the robustness of different system components is analysed, identifying the most vulnerable components and architectures. Understanding how different distortions impact DL-based ALPR systems can provide insights for further research to develop more resilient ALPR systems capable of maintaining high performance under adverse conditions. The paper is organised as follows. ALPR models used are presented in Section 2.1, followed by an overview of the datasets used (Section 2.2) and the distortions applied (Section 2.3). The metrics for evaluating the performance of ALPR models are defined in Section 2.4. In Section 3, experimental results

are presented and discussed. Finally, conclusions and future work directions are presented in Section 4.

2. METHODOLOGY

2.1 ALPR Models

The architectural details of SOTA ALPR systems are often kept confidential for commercial purposes, and older implementations are often outdated and not maintained. To provide a methodological framework for the assessment of ALPR systems performance, due to the limited availability of implementations, the following three ALPR systems are used in this study:

- (1) ALPR in Unconstrained Scenarios [12]
- (2) HyperLPR [34]
- (3) UltimateALPR-SDK [35]

Training of the models was not part of this work, i.e. pre-trained models were used to examine their robustness. The observed performance variations under different simulated distortion types offer valuable insights for future work aimed at modifying the models to enhance their robustness against various distortions.

ALPR in Unconstrained Scenarios, published by Silva and Jung [12], has a well-documented architecture and design details. It was designed to address challenging image-capture scenarios, such as oblique camera views in which the vehicle's license plate is significantly tilted. It comprises four main components (Figure 1). The vehicle detection component uses the pre-trained YOLOv2 [36] network, modified to merge all vehicle-related classes into a single entity and discard other classes. YOLOv2 was chosen for its balance of speed and precision [12]. The license plate detection component uses Warped Planar Object Detection NETWORK (WPOD-NET) to locate the license plate and regress one affine transformation per detection, producing bounding box coordinates and affine transformation parameters. This network consists of a total of 21 convolutional layers (filter size 3×3), followed by ReLU activation functions, and four max pooling layers (size 2×2 with stride 2). The final detection head is divided into two branches. The first branch uses the Softmax activation function to predict the probability that the license plate is located in the current block of pixels. The second branch uses the identity activation function to regress the six affine parameters. The areas with object probability higher than the specified thresholds are predicted to be the license plates and the affine parameters are used for their rectification, transforming the license plate to resemble a front view. Finally, the OCR network recognises the characters on the rectified license plate. This network is based on a modified YOLO network [6]. Silva and Jung [12] trained the WPOD-NET network using 196 images from three different datasets. Various data augmentation techniques were applied to all training images. The OCR network was based on a modified YOLO network [6], and was trained with a large




Figure 1. Block diagram showing the components of the ALPR in Unconstrained Scenarios [12] model. The image shown is a cropped version of one of the images in the UFPR-ALPR dataset [18].

number of synthetic and augmented images to improve robustness to regional variations in license plate formats [12].

HyperLPR [34] is an open-source ALPR system designed to read Chinese license plates. Although it is not a part of a specific research publication, it has been used in various ALPR studies on adversarial attacks on ALPR systems [37, 38]. Information about its architecture was derived from scientific publications and analysis of the available implementation. The model is composed of three convolutional layers with increasing filter sizes ($3 \times 3 \times 32$, $3 \times 3 \times 64$, and $3 \times 3 \times 128$), followed by batch normalisation, ReLU activation, and 2×2 max pooling layers. The mentioned components of the network are used for detecting license plates, and the output is fed into a network with four gated recurrent units of 256 units each, a dropout layer, and a softmax output layer normalising an 84-unit probability distribution (corresponding to the number of possible license plate characters). There is no information on the dataset used for the training, and the training implementation is not available.

UltimateALPR-SDK, a SOTA commercial solution [35], is known to be the fastest ALPR implementation with high precision and accuracy [39]. During experimentation, it performed recognition significantly faster than the other two models. According to its documentation, the model is trained on license plates from more than 150 countries, primarily in Europe but also in the USA, Canada, Russia, Indonesia and others. There is no public information on its architecture and the code is protected in a way that makes it impossible to deduce it, making it a “black box”. The open-source version used in this study constrains the last character of the license plate. For example, if the plate shows AB123CD, the model outputs AB123C*.

2.2 Datasets

To ensure diversity of the origins and formats of the license plates, a comprehensive search for available datasets was conducted. Six datasets (Table I) from four different continents were obtained, containing license plates with different formats, font styles, and colours, acquired under various conditions (night, day, background scene complexity, etc.).

The Caltech Cars 1999 dataset [19] consists of parked cars images acquired from the rear. There is no information available on image capture. All images were acquired during daytime, from approximately the same distance, without camera tilt. The license plates vary in colour, format, and

Table I. Overview of the datasets used in our study.

Dataset name	Country	Number of images	Resolution	Compression
Caltech Cars 1999 [19]	USA	120*	896×592	Lossy
AOLP [20]	Taiwan	2,049	various	Lossy
Croatian (CRO) [21]	Croatia	502*	640×840	Lossy
PKU [40]	China	3,977	various	Lossy
RodoSol-ALPR [41]	Brazil	10,000	1280×720	Lossless
UFPR-ALPR [18]	Brazil	3,600	1920×1080	Lossless
Total number of images		20,248		

*Some of the images have been removed due to obstructed view of license plates.

number of characters, and consist of a combination of Arabic numerals and Latin characters.

The Application-Oriented License Plate Recognition (AOLP) dataset [20] contains images divided into three groups depending on the complexity of the intended application: 681 images for the access control (AC) group, 757 images for the law enforcement (LE) group, and 611 images for the road patrol (RP) group. The images were captured with different imaging devices and have different spatial resolutions and illumination conditions. The imaging devices were not disclosed. However, it is revealed that the AC and LE image groups were collected with fixed cameras at entry points or roadside, while the images of the RP group were captured with mobile devices. The license plates comprise black Latin characters and Arabic numerals on a white background, with each plate having a fixed length of six characters and numerals.

The Croatian (CRO) license plate dataset [21] contains images acquired with an Olympus Cmedia C-2040ZOOM digital camera. Most images of cars were taken from the rear, with no significant tilt. They were all captured at a very close distance, with few images taken at night.

The PKU dataset [40] consists of images captured under various conditions and is divided into five groups (G1–G5) depending on the acquisition conditions. G1 contains 810 images of cars on highways, captured during the day, while G2 consists of 700 images of cars and trucks on highways during daytime, with some glare from the sunlight present. In G3, there are 743 images of cars and trucks on highways captured during night. G4 contains 572 images of cars and trucks on city roads, both in the day and at night. Finally, G5 contains 1152 images of cars and trucks at intersections with crosswalks. There is no information on the capture

devices, but it is discernable that different devices were used for different groups (G1–G5). The license plates consist of white characters on a blue background and follow a fixed format: one Chinese character, one Latin character, followed by five characters, a combination of Latin characters and Arabic numerals.

The RodoSol-ALPR dataset [41] consists of 20,000 images captured by static cameras located at toll booths. Half of the dataset are images of cars, while the other half are images of motorcycles. The license plates of cars are formatted in a single line, while those of motorcycles are in two lines. There are variations in colour, including black, blue, or red characters on a white background, and white characters on a red background. Since the chosen models in this study do not handle two-line plates, images of motorcycles were excluded, resulting in 10,000 images used from this dataset. There is no information available on the capture devices. The acquisition conditions vary in terms of illumination, vehicle position in the scene, and camera distance.

The UFPR-ALPR dataset [18] was collected in the following way: three different cameras were mounted inside a vehicle (GoPro Hero4 Silver, Huawei P9 Lite and iPhone 7 Plus). The vehicle then followed 150 vehicles, 120 cars (passenger cars, buses, trucks, vans), and 30 motorcycles. For each tracked vehicle, 30 images were acquired, resulting in a total of 4500 images. The background was complex and there were variations in the illumination conditions (e.g. shadows). In this work, all 30 images are treated as separate images, as they vary significantly with respect to camera-to-license plate distance and illumination conditions. The license plates contain seven characters in fixed format: three Latin characters followed by four Arabic numerals. The dataset includes license plates with black characters on a white background and white characters on a red background. The car license plates are formatted in a single line, while the motorcycle license plates are formatted in two lines. The latter were excluded from this study, resulting in 3600 images used from this dataset. The images were acquired in daylight with clear weather conditions.

There are certain limitations and restrictions that must be addressed in all presented datasets. First, it is not reasonable to expect that datasets with a small number of images, as well as those having similar and not-so-complex acquisition conditions (Caltech Cars 1999, CRO, AOLP), to be sufficient for training SOTA deep learning ALPR architectures. The other three datasets (PKU, RodoSol-ALPR, and UFPR-ALPR) could serve as valuable sources for training such architectures, especially UFPR-ALPR, which has significant variations in the acquisition conditions. Perhaps the most important limitation is the fact that the initial technical quality is not characterised. In other words, the exact capture conditions and distortions present in the images are unknown (i.e., noise level, lens blur, optical distortions, image processing and compression artifacts, etc.). In particular, for this study, characterising the noise originating from the imaging pipeline would allow the

quantification of the total image noise present in the images (original and simulated). To the best of our knowledge, no relevant dataset published so far includes characterised technical quality.

2.3 Distortions

Given the known generalisation capabilities of ALPR systems [25], it is of interest to investigate how different weather distortions influence the performance of ALPR systems. This can be approached in two ways: either by acquiring images in various weather conditions, or by simulating the distortions introduced by adverse weather conditions on existing license plate datasets. The first is less favourable, since it would involve capture of the same scene from a fixed point under different conditions. In the second approach, the recognition performance is first determined on the original images, then systematic and controlled simulated weather distortions are added to them. The weakness of this approach lies in the fact that simulated weather distortions deviate from real-world scenarios. For example, simulating rain by adding a rain mask to the image [42] does not account for the potential reduction in scene brightness due to clouds.

In this work, rain (Figure 2(b)), brightness (Fig. 2(c)), fog (Fig. 2(d)), frost (Fig. 2(e)), and snow (Fig. 2(f)) were simulated on images as weather distortions. The brightness distortion simulates camera overexposure, which reduces the performance of license plate detection algorithms due to reduced contrast, clipping, and loss of details in highlights [43]. The presence of fog, snow or rain represents significant challenges for modern ALPR systems [31, 44], as they can potentially occlude important parts of the scene. Frost on the camera lens is another type of distortion that affects the performance of SOTA object detection algorithms [24]. The brightness, fog, frost, and snow distortions were applied using the image corruption library introduced by Michaelis et al. [24] with five intensity levels defined in the library. Rain distortion was applied using Adobe After Effects [45] software, a well-known tool for developing deraining algorithms [42]. Three different rain configurations were created depending on the presence and direction of the wind (no wind W_0 , wind in one direction W_1 , and wind in the opposite direction W_2). The parameters used to create rain-distorted images were defined by trial-and-error: the number of raindrops (10,000), the size of the raindrops (level four), and the speed of the rainfall (5000). In cases where wind is present, the wind intensity was 2000 and the wind direction variations 20%.

The simulation of the imaging pipeline was chosen to systematically add read noise to the datasets (Fig. 2(g)), and read noise was modeled using Image Systems Evaluation Toolkit (ISET) [46]. Noise is generated during the process of analog-to-digital signal conversion [47]. The experimental configuration followed that of Plavac et al. [48], where 15 noisy images were generated for each original image, with read noise levels ranging from 0.2 to 14.2 mV in 1 mV increments.

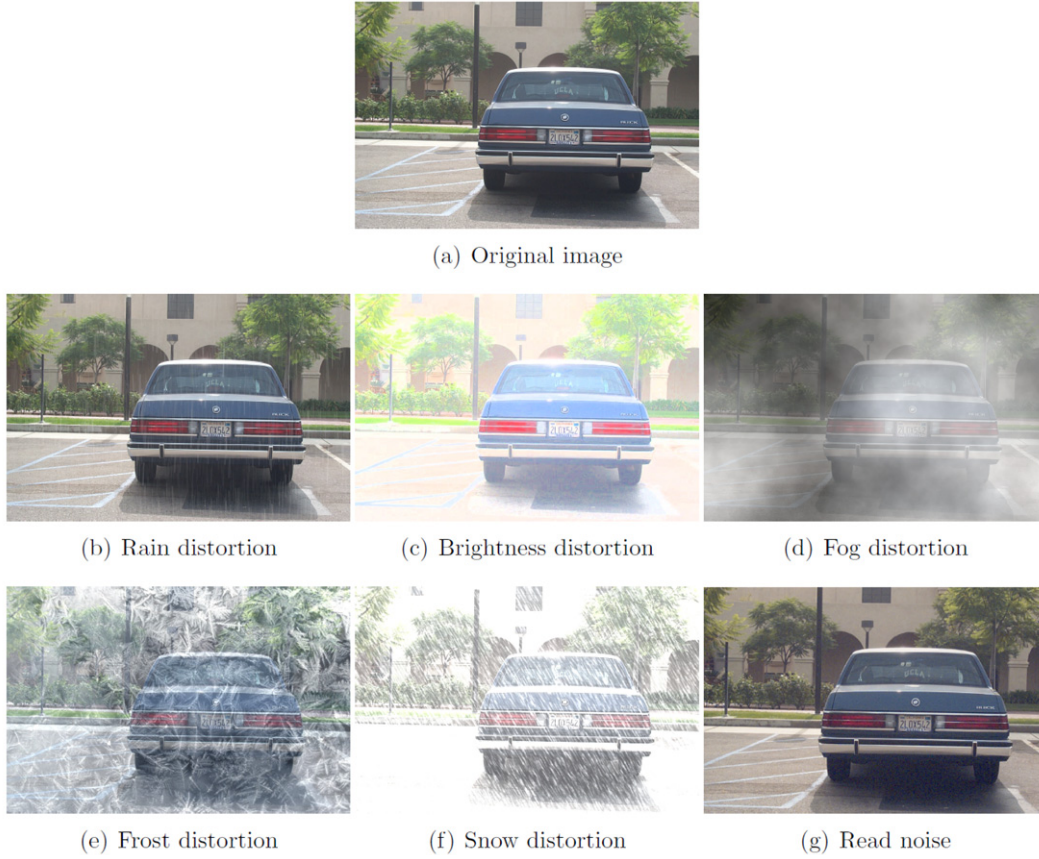


Figure 2. An example image from the Caltech Cars 1999 [19] dataset (a), along with sample images of maximum intensity level of rain simulation (b), brightness simulation (c), fog simulation (d), frost simulation (e), snow simulation (f), and the maximum intensity level of read noise distortion (g).

The distortion algorithms used are popular and widely used in the research community to investigate the robustness of DL-based models [49–51] and to develop deraining algorithms [42]. However, certain limitations in their realism and consistency should be noted. Across different types of distortions (Fig. 2), it is obvious that each distortion damaged the original image in a different way and to a different degree for the same intensity level. The algorithm used for fog distortion does not account for scene depth, despite fog intensity being a function of depth [52]. Additionally, it significantly reduces scene brightness, making it resemble smoke more than fog. The frost distortion algorithm introduces non-monotonic changes in image contrast as distortion levels increase. The contrast gradually decreases from levels one to three, but increases at levels four and five. Such inconsistencies can cause ambiguities in the performance evaluation, i.e. changes in ALPR system accuracy could be attributed to these inconsistencies rather than to the distortion level itself. The snow distortion does not fully convey a real snowfall, even though snow is falling, it is not seen on the ground that remains green (vegetation). Additionally, the occlusion of the license plate appears to be more pronounced at distortion level four than at level five. The rain distortion is simplified compared to what could be achieved with physically-based rain rendering algorithms. While the modeled read noise

provides grounds for a comprehensive evaluation of noise impact, the compound pre-existing noise level from capture affects the interpretation of results.

2.4 Performance Evaluation Metrics

Various metrics have been proposed in literature to evaluate the performance of ALPR systems [6, 12, 22]. In this work, the performance of ALPR is assessed using the following measures:

- **Total Accuracy (TA)** compares ground truth license plate characters with predicted license plate characters, and is calculated as:

$$TA(\%) = \frac{N_{\text{correct}}}{N} \times 100, \quad (1)$$

where N_{correct} is the number of complete matches, and N is total number of images.

- **License Plate Detection Error (LPDE)** is calculated as:

$$LPDE(\%) = \frac{N_{\text{empty}}}{N} \times 100, \quad (2)$$

where N_{empty} represents the number of images for which the model output is empty.

- **License Plate Recognition Error (LPRE)** is then:

$$LPRE(\%) = (100 - TA(\%)) - LPDE(\%) \quad (3)$$

The assumption that an empty output indicates a license plate detection error was derived during trials, but may not always be true. For instance, if the detected bounding box contains only part of the license plate, the recognition might be entirely correct for that partial detection (e.g., half of the characters). Using the current evaluation method, this scenario would be classified as a license plate recognition error. The common methodology in literature is to evaluate the performance of the license plate detection component using Intersection over Union (IoU) metric [22]. This was not feasible in this work since ground truth bounding boxes for license plate detection were available for only one out of the six datasets used.

The above-mentioned metrics were used to evaluate the performance of the three models on original and distorted images. Since the ALPR in Unconstrained Scenarios model is capable of detecting Latin characters and Arabic numerals, the PKU dataset was not part of this model evaluation. The HyperLPR model was developed specially to recognise Chinese license plates; therefore, it is used only on the PKU dataset. For the UltimateALPR-SDK model, modifications were made because its predictions are restricted, and the last character was ignored in the evaluation for all datasets used. The model also skips the first character in the PKU dataset (Chinese alphabet letter), which was accounted for during evaluation of performance. This left the performance evaluation of PKU dataset based on the remaining five characters.

3. RESULTS AND DISCUSSION

3.1 Performance of the Models with Original Images

Since the ALPR in Unconstrained Scenarios model only supports license plates containing Latin alphabet characters, the PKU dataset was not included in its performance evaluation. The performance of ALPR in Unconstrained Scenarios and UltimateALPR-SDK models on all datasets, except the PKU dataset, is shown in Table II. On average, the ALPR in Unconstrained Scenarios model showed good performance with an average TA of 78.33%. The highest TA was achieved with the RP group of AOLP dataset, while the lowest TA was observed in the AC group of images from the same dataset. This performance drop is attributed to the close proximity of the camera to the cars in the AC group, leading to car detection failures with YOLOv2-based vehicle detection module, which is reflected by LPDE of 29.52%.

Compared to other two models, the UltimateALPR-SDK model demonstrated very high performance with all datasets (Tables II and III). Except for RodoSol-ALPR and the LP group of AOLP dataset, the model achieved TA above 90% for all datasets used in the study. Similar to previous models, the license plate recognition component was the most error-prone component, while the detection component performed perfectly in the Caltech Cars 1999 dataset and two groups of the PKU dataset. It is important to note that this model had an advantage over the other two models because the last character in the model output was restricted.

Table II. Performance comparison of models on various datasets (excluding PKU dataset). The best performance for each metric is shown in bold.

Dataset	ALPR in Unconstrained Scenarios			UltimateALPR-SDK		
	TA (%)	LPDE (%)	LPRE (%)	TA (%)	LPDE (%)	LPRE (%)
Cars 1999	87.50	0.83	11.67	95.00	0	5.00
AOLP AC	57.27	29.52	13.22	97.06	0.15	2.79
AOLP LE	86.53	10.04	3.43	86.00	0.40	13.61
AOLP RP	89.85	1.80	8.35	92.96	0.82	6.22
CRO	83.86	5.38	10.76	93.82	0.80	5.38
RodoSol-ALPR	75.16	8.44	16.40	92.86	0.30	6.84
UFPR-ALPR	68.47	12.97	18.56	82.31	1.50	16.19
Average	78.38	9.85	11.77	91.43	0.57	8.00

Table III. Performance comparison of models on PKU dataset. The best performance for each metric is shown in bold.

Dataset	HyperLPR			UltimateALPR-SDK		
	TA (%)	LPDE (%)	LPRE (%)	TA (%)	LPDE (%)	LPRE (%)
PKU G1	98.77	0.49	0.74	94.57	0	5.43
PKU G2	98.71	0.14	1.15	94.86	0	5.14
PKU G3	98.92	0.40	0.67	93.94	0.13	5.92
PKU G4	96.68	0	3.32	95.28	0.17	4.55
PKU G5	96.09	0.52	3.39	91.02	0.13	8.85
Average	97.83	0.31	1.86	93.93	0.09	5.98

Therefore, the evaluation was based on the remaining characters. In the case of the PKU dataset, the first (Chinese) letter was also skipped.

HyperLPR was trained to recognize Chinese license plates. Therefore, the performance analysis focuses solely on the PKU dataset. Table III shows that the model performed exceptionally well in all five groups of PKU dataset, with an average TA of 97.83% which was expected, since the model was evaluated on the license plates in the same regional format as the ones used to train the model.

It can be seen that the license plate recognition component was usually more prone to errors compared to the license plate detection component in all the models used, with exceptions in the AC and LE groups of AOLP dataset for ALPR in Unconstrained Scenarios model. The YOLO-based vehicle detection component used in this model tends to fail with detection when dealing with images with great vehicle proximity.

3.2 Performance of the Models with Distorted Images

3.2.1 Weather Distortions

Results do not show significant differences in the performance metrics between the three models on angular configurations of the rain-distorted images. The comparison of TA between the original and rain-distorted images

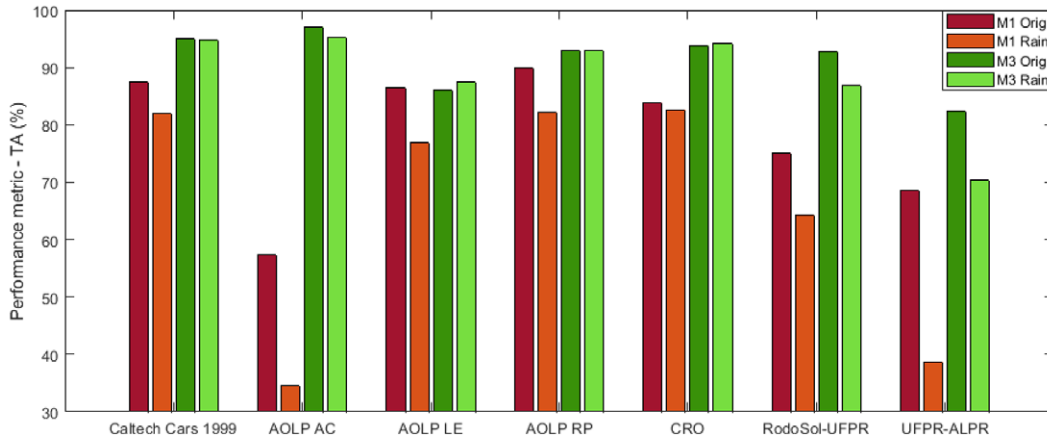


Figure 3. Total Accuracy (TA) for ALPR in Unconstrained Scenarios (M1) and UltimateALPR-SDK (M3) on rain-distorted images, averaged over the three angular configurations.




Figure 4. Total Accuracy (TA) for HyperLPR (M2) and UltimateALPR-SDK (M3) on rain distorted images of PKU dataset, averaged over the three angular configurations.

(averaged for the three angular configurations) for all datasets, except PKU dataset, for ALPR in Unconstrained Scenarios and UltimateALPR-SDK models is shown in Figure 3.

The UltimateALPR-SDK model showed much lower sensitivity to the introduced rain distortion compared to the ALPR in Unconstrained Scenarios model. In certain cases (LE group in the AOLP dataset, CRO dataset), introducing rain distortion led to a slight performance improvement for the UltimateALPR-SDK model.

TA obtained for PKU dataset and the HyperLPR and UltimateALPR-SDK models is shown in Figure 4. These two models demonstrated good robustness to rain distortion. The latter model performed worse than the former model in all cases, but HyperLPR model showed a higher sensitivity to rain distortion. The UltimateALPR-SDK model did not suffer from significant performance drops across the five groups of PKU dataset.

In most cases, the performance decreased for distorted images, but there were few cases where the performance

increased after introducing rain-distorted images. It is of interest to investigate which component of the system is affected in both situations, license plate detection or license plate recognition. Table IV shows changes in license plate detection and license plate recognition performance for rain-distorted images using the ALPR in Unconstrained Scenarios model. It can be observed that introducing rain distortion generally increased LPDE across all datasets except Caltech Cars 1999 dataset. The decrease in LPRE does not indicate an improvement in the license plate recognition module, but rather that pre-existing errors are now more significantly absorbed by the license plate detection module.

HyperLPR model showed a stable and minimal performance decrease for all five groups of PKU dataset, with LPDE increasing more frequently than LPRE (Table A.1. For UltimateALPR-SDK model, the performance changes did not follow a clear pattern (Table A.2). In certain cases (RodoSol-ALPR, UFPR-ALPR datasets), the performance drop was mainly due to an increase in LPRE. When performance improved on rain-distorted images compared

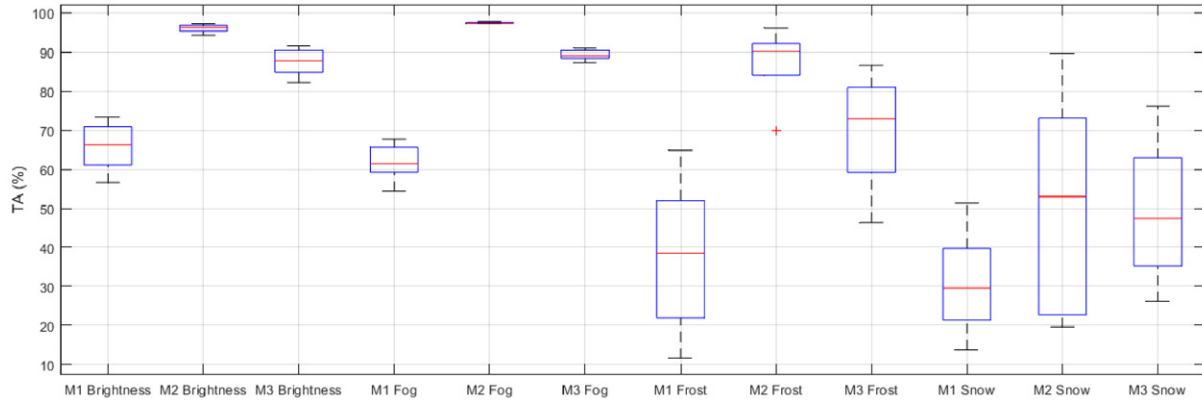


Figure 5. Total Accuracy (TA) distribution for three models across different weather distortion types (brightness, fog, frost, and snow). Performance of ALPR in Unconstrained Scenarios model is noted with M1, for HyperLPR model with M2, and for UltimateALPR-SDK model with M3. The distribution is extracted from the obtained TA across all datasets and distortion levels.

Table IV. Performance of the ALPR in Unconstrained Scenarios model on images with simulated rain distortion. The number in the parentheses represents the change from the performance obtained on the corresponding original images, with red and green colours indicating worse and better performance, respectively.

Dataset	LPDE (%)	LPRE (%)
Caltech Cars 1999	0.28 (−0.56)	17.78 (+6.11)
AOLP AC	56.14 (+26.62)	9.35 (−3.87)
AOLP LE	21.00 (+10.96)	2.05 (−1.38)
AOLP RP	10.75 (+8.95)	6.98 (−1.37)
CRO	6.04 (+0.66)	11.29 (+0.53)
RodoSol-ALPR	20.59 (+12.14)	15.10 (−1.29)
UFPR-ALPR	43.22 (+30.25)	18.05 (−0.51)

to original images (AOLP LE, CRO, PKU G3 datasets), the license plate detection component performed slightly worse or did not change compared to original images, while the license plate recognition component performed better on rain-distorted images. It is challenging to find valid explanations for this rather unexpected behavior. Rain distortion of images of this dataset may have enhanced the edges of the license plate characters, improving their visibility against the background. Another possible explanation is that the introduced rain distortion might have obscured certain artifacts or background elements that interfered with recognition, resulting in cleaner inputs for the recognition module. The assumptions for such behavior should be taken with great caution. Further investigation and image analysis are needed to fully understand this effect and determine the underlying causes of unexpected performance variations.

Other weather distortions applied are part of the “image corruptions library” [24], and are applied at five intensity levels. The distribution of TA for different models and types of weather distortion is shown in Figure 5. The ALPR in Unconstrained Scenarios model performed best under brightness distortion, followed by fog, frost,

and snow distortions. The increased variability in frost and snow distortions indicates greater inconsistency in performance across datasets. HyperLPR showed consistently robust performance under brightness and fog distortions with minimal variability, across all groups of PKU dataset. This model also performed fairly well under frost distortion, with one outlier and slightly higher variability. The worst performance was observed for snow distortion, with the highest variability. The UltimateALPR-SDK model showed high TA with relatively low variability under brightness and fog distortions. Performance dropped and variability increased significantly under frost and snow distortions. In summary, HyperLPR and UltimateALPR-SDK models showed excellent performance in handling brightness and fog distortions, while ALPR in Unconstrained Scenarios model showed a relatively good performance. The ALPR in Unconstrained Scenarios model struggled significantly with frost distortion, showing the highest variability, while the other two models still performed well. Snow distortion was challenging for all three models and showed high variability.

For ALPR in Unconstrained Scenarios, introducing brightness distortion decreased TA for all datasets, with the performance declining as the distortion intensity increased. The most severe performance drop was observed for RodoSol-ALPR dataset (from TA 65.50% at level 1 to 14.65% at level 5). Similar trend was observed in UltimateALPR-SDK model, but the magnitude of the performance drop varied significantly between datasets. The most severe performance drop was again observed for RodoSol-ALPR dataset (from TA 83.32% at brightness distortion level 1 to TA 28.62% at level 5). HyperLPR model showed the highest robustness to brightness distortion across the employed models, with a slight performance drop observed for all groups of PKU dataset (1–3% reduction in TA from level 1 to level 5), except for group G4 where the drop in TA was around 9%.

The ALPR in Unconstrained Scenarios model showed the highest sensitivity to fog distortion among the models used, with TA consistently decreasing as the distortion level




Figure 6. Total Accuracy (TA) for different intensity levels of certain simulated weather distortions across datasets. (a) HyperLPR, snow distortion. (b) ALPR in Unconstrained Scenarios, frost distortion.

increased across all datasets. The performance of the other two models declined minimally as the fog distortion level increased.

For frost-distorted images, the performance decreased as the distortion level increased across all datasets and models used in the study. The obtained TA for each dataset with different intensity levels of frost distortion for ALPR in Unconstrained Scenarios is shown in Figure 6(b). The lowest performance was observed for the AC group of AOLP dataset (0.14% TA at level 5). The highest drop was for the RP group (from 87.73% at level 1 to 1.96% at level 5). For UltimateALPR-SDK model, the highest performance drop was for AC group of AOLP dataset (from 96.04% to 35.39%). HyperLPR showed to be slightly more robust to frost distortion, with a severe performance drop for group G4 of PKU dataset (from 93.18% to 58.04%).

Snow distortion also greatly affected the performance of the models in all datasets. Performance for the HyperLPR model is shown in Fig. 6(a). As mentioned earlier, the occlusion of license plates is more pronounced at distortion level 4 than at level 5. This is confirmed by TA metric: performance decreases as distortion level increases for all datasets, but more so at level 4 than level 5. The least affected was group 3 of PKU dataset, and the highest performance drop was for group 1 of PKU dataset. Similar performance trends were observed for the other two models. The CRO dataset was the least affected by the introduced distortions, while the lowest score was observed for the RodoSol-ALPR dataset.

The average LPDE and LPRE of UltimateALPR-SDK model for each distortion type at five intensity levels are shown in Table V. The license plate detection and recognition components experienced higher error rates as the distortion levels increased. For all distortion types considered (brightness, fog, frost, and snow), the license plate recognition component was more affected by the distortion. At high levels of snow distortion, LPDE approached relatively close to LPRE. Observations for ALPR in Unconstrained

Table V. Averaged LPDE and LPRE across datasets of the UltimateALPR-SDK model for different weather distortion types: brightness, fog, frost, and snow.

	Level 1	Level 2	Level 3	Level 4	Level 5
LPDE (Brightness)	0.75	1.33	2.21	3.71	5.38
LPRE (Brightness)	7.71	8.49	9.90	10.59	12.47
LPDE (Fog)	1.14	1.79	2.39	2.50	3.46
LPRE (Fog)	7.92	7.88	8.55	8.65	9.17
LPDE (Frost)	2.75	5.46	8.16	13.12	18.54
LPRE (Frost)	10.63	15.48	18.94	23.35	35.18
LPDE (Snow)	4.30	13.01	15.38	31.59	27.87
LPRE (Snow)	19.53	28.54	37.21	42.28	33.93

Scenarios model are quite different (Table A.3). With brightness distortion, both LPDE and LPRE increased with increasing levels, the license plate detection component was more prone to errors. For fog distortion, LPRE increased until level 3, after which it decreased. The license plate detection component again experienced increased errors with higher distortion levels, being the more error-prone component. For frost and snow distortions, the license plate detection errors increased with the distortion level and the detection component was the dominant error source, while LPRE slightly decreased with higher distortion levels. This does not imply better performance of the license plate recognition component, but rather that LPDE absorbed errors from the recognition component. HyperLPR behaved similarly to UltimateALPR-SDK for brightness, fog, and frost distortions (Table A.4). In these cases, LPDE and LPRE increased with distortion levels, with recognition errors being dominant. Interestingly, snow distortion reversed this trend after level 3, making detection the more error-prone component of the HyperLPR model.




Figure 7. The HyperLPR model: Total Accuracy (TA) for read noise distortion over PKU dataset. Distortion level zero indicates performance on original images. The remaining distortion levels correspond to read noise with intensity from 0.2 mV to 14.2 mV with 1 mV increments.

3.2.2 Read Noise Distortion

Introducing read noise led to a decrease in TA of HyperLPR model in all groups of PKU dataset (Figure 7). As distortion levels increased, performance decreased in all groups. The highest performance drop was observed for groups G3 and G1 of PKU dataset (from 97% TA on the original images to 79% at maximum read noise level). G4 showed the lowest performance drop with respect to read noise distortion. LPDE increased with increasing read noise levels for all groups in PKU dataset (Figure A.1). The highest increase was observed for G1 and G3, corresponding with the observed TA reduction. LPRE experienced a slight increase over the levels of read noise introduced, with the effect most prominent in G5 of PKU dataset (Figure A.2).

For ALPR in Unconstrained Scenarios model, TA decreased with increasing read noise levels (Figure A.3). This effect is visible across all datasets except for Caltech Cars 1999 dataset, which exhibited a rather noisy behavior. In the case of this model, LPDE increased as distortion levels increased for most of the datasets (Figure A.4). However, the Caltech Cars 1999 and CRO datasets did not follow this trend, with LPDE barely experiencing any changes with increasing read noise levels. LPRE showed changes of much smaller intensity and increased variability, making it difficult to identify a

consistent pattern across multiple datasets (Figure A.5). The only clear increase in LPRE with increasing read noise levels was observed in UFPR-ALPR dataset. From these results, it can be concluded that read noise mainly influenced the performance of ALPR in Unconstrained Scenarios by reducing the performance of the license plate detection component.

Interpreting TA changes for UltimateALPR-SDK is more complex (Figure A.6). Clear performance reductions were observed for UFPR-ALPR, RodoSol-ALPR, and G1 and G3 groups of PKU dataset. However, CRO dataset and LE group of the AOLP dataset showed an increase in TA with the read noise levels. Other datasets showed rather noisy behavior, oscillating around a certain TA interval. A clear increase in LPDE was observed for UFPR-ALPR, RodoSol-ALPR, and G1 and G3 groups of PKU dataset. The license plate recognition component demonstrated a variable behavior across datasets. UFPR-ALPR dataset, G5 group of PKU dataset, and RP group of AOLP dataset showed an increase in LPRE with increasing read noise levels. Other datasets showed a fluctuating LPRE around a certain intensity or even a decrease in LPRE (Caltech Cars 1999, LE group of AOLP dataset).

In designing the experiment, our aim was to replicate real-world conditions by simulating read noise levels. Although we were unable to find reliable sources for typical read noise levels in literature, anecdotal evidence suggest upper limits of 3–5 mV. The upper limits in our experiment significantly exceeded these limits, but the simulation results (Fig. 7) indicate that relatively low levels of read noise (0.2 mV, corresponding to distortion level 1) significantly affect recognition accuracy.

4. CONCLUSION AND FUTURE WORK

This study evaluated the impact of various simulated weather distortions and camera read noise on the performance of deep learning-based Automatic License Plate Recognition (ALPR) systems. Three models, each representing different architectural approaches, were analyzed across six datasets to evaluate their robustness and identify areas for improvement. Introduction of weather distortions (i.e. rain, brightness, fog, frost, and snow) significantly impacted the performance of the ALPR systems used in this study. Snow and frost had the most severe effects, often reducing system accuracy to near zero. The modeled camera read noise also led to a performance decrease, with even minimal read noise levels (0.2 mV) causing noticeable drops in recognition accuracy. Performance changes were not consistent across different datasets. While some datasets followed general trends (e.g., performance reduction under frost distortion), others exhibited unique behaviors (e.g., performance improvement under rain distortion for certain datasets). The sensitivity to distortions varied significantly among the models. HyperLPR model showed the greatest robustness to weather distortions, followed by UltimateALPR-SDK model, and then ALPR in Unconstrained Scenarios. For read noise, the UltimateALPR-SDK model demonstrated the highest robustness, whereas ALPR in Unconstrained Scenarios was again identified as the most sensitive model. Both license plate detection and recognition components were affected by distortions, but the extent varied depending on the type of distortion. For instance, rain distortion and read noise primarily affected the detection component, while brightness, fog, frost, and snow distortions mainly affected the recognition component.

There are several directions for future work. First, further image analysis of the distorted images is planned, with the aim of explaining the inherent image attribute changes (e.g., contrast, brightness, colour, edge density, blur, noise, etc.) when distortion level increases. This can elucidate the effect of distortions on the detection and recognition performance of ALPR systems. Second, future research should explore additional distortion types that can occur in the imaging pipeline, and their combined effects on performance of ALPR systems. Including real-world distorted license plates in the study would be beneficial to confirm the observations made on the simulated distortions. Once the effects of various distortions are understood, the next step would be to investigate and develop different

strategies to mitigate these negative effects. This includes replacing, modifying, or fine-tuning the parts of the systems that were found to be the most sensitive to distortions. Accordingly, enhancements can be incorporated into the model pipeline, forming a new hybrid model. The findings underscore the importance of addressing image distortions in the design and deployment of ALPR systems. Robustness to environmental and imaging pipeline-induced distortions is crucial for maintaining high accuracy and reliability in real-world applications.

ACKNOWLEDGMENT

Seyed Ali Amirshahi was supported by the project “VQ4MediCS: Video Quality Assessment and Enhancement for Pre-Hospital Medical Services” (grant number 329034) from the Research Council of Norway. Marius Pedersen and Sophie Triantaphillidou were supported by the Research Council of Norway through the “Quality and Content” project (grant number 324663).

APPENDIX A.

Table A.1. Performance of HyperLPR model on images with simulated rain distortion, averaged over three wind configurations. The numbers in the parentheses represent the changes from the performance obtained on the corresponding original images, with red and green colours indicating worse and better performance, respectively.

Dataset	LPDE (%)	LPRE (%)
PKU G1	1.77 (+1.28)	0.95 (+0.21)
PKU G2	1.10 (+0.96)	1.22 (+0.08)
PKU G3	1.35 (+0.94)	1.26 (+0.58)
PKU G4	1.05 (+1.05)	2.91 (−0.41)
PKU G5	1.09 (+0.57)	4.69 (+1.29)

Table A.2. Performance of UltimateALPR-SDK model on images with simulated rain distortion. The numbers in the parentheses represent the changes from the performance obtained on the corresponding original images, with red and green colours indicating worse and better performance, respectively.

Dataset	LPDE (%)	LPRE (%)
Caltech Cars 1999	0 (0)	5.28 (+0.28)
AOLP AC	1.42 (+1.27)	3.38 (+0.59)
AOLP LE	0.86 (+0.46)	11.56 (−2.05)
AOLP RP	0.76 (−0.06)	6.28 (+0.06)
CRO	1.26 (+0.46)	4.45 (−0.92)
PKU G1	0 (0)	5.93 (+0.50)
PKU G2	0.14 (+0.14)	5.10 (−0.04)
PKU G3	0.13 (0)	5.38 (−0.54)
PKU G4	0.58 (+0.41)	4.78 (+0.23)
PKU G5	0.39 (+0.26)	9.24 (+0.39)
RodoSol-ALPR	2.69 (+2.39)	10.45 (+3.61)
UFPR-ALPR	5.48 (+3.98)	24.17 (+7.98)

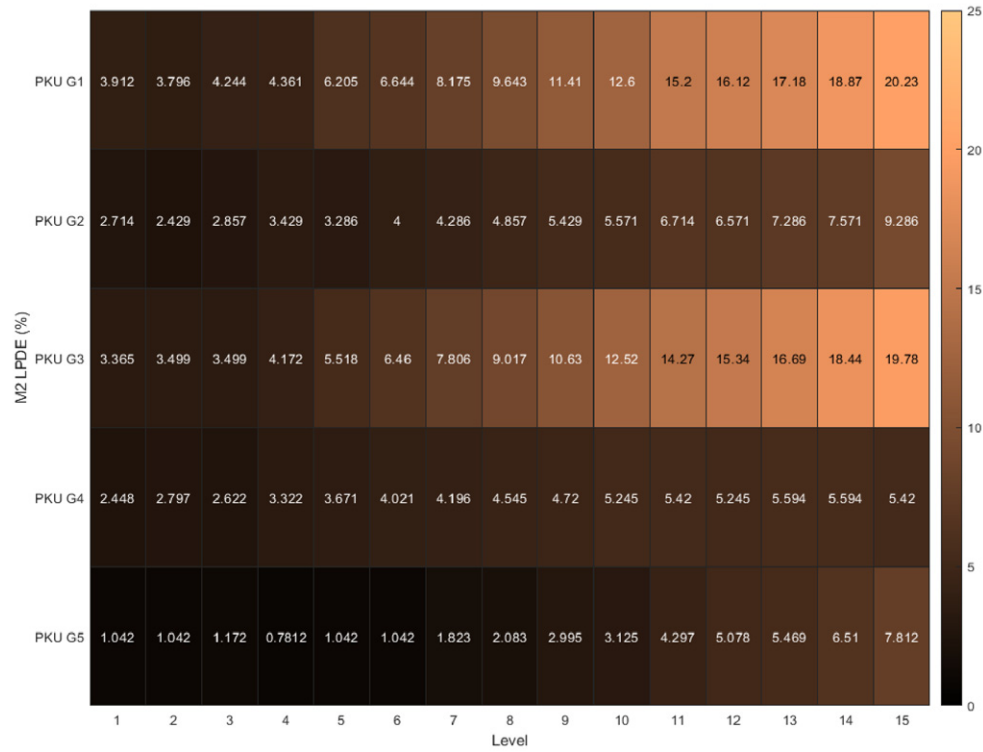


Figure A.1. The HyperLPR model: License Plate Detection Error (LPDE) for read noise distortion on PKU dataset. The distortion levels correspond to read noise with intensity from 0.2 mV to 14.2 mV with 1 mV increments.




Figure A.2. The HyperLPR model: License Plate Recognition Error (LPRE) for read noise distortion on PKU dataset. The distortion levels correspond to read noise with intensity from 0.2 mV to 14.2 mV with 1 mV increments.




Figure A.3. The ALPR in Unconstrained Scenarios model: Total Accuracy (TA) for read noise distortion over all datasets, except PKU dataset. Distortion level zero indicates performance on original images. The remaining distortion levels correspond to read noise with intensity from 0.2 mV to 14.2 mV with 1 mV increments.



Figure A.4. The ALPR in Unconstrained Scenarios model: LPDE for read noise distortion over all datasets, except PKU dataset. The distortion levels correspond to read noise with intensity from 0.2 mV to 14.2 mV with 1 mV increments.

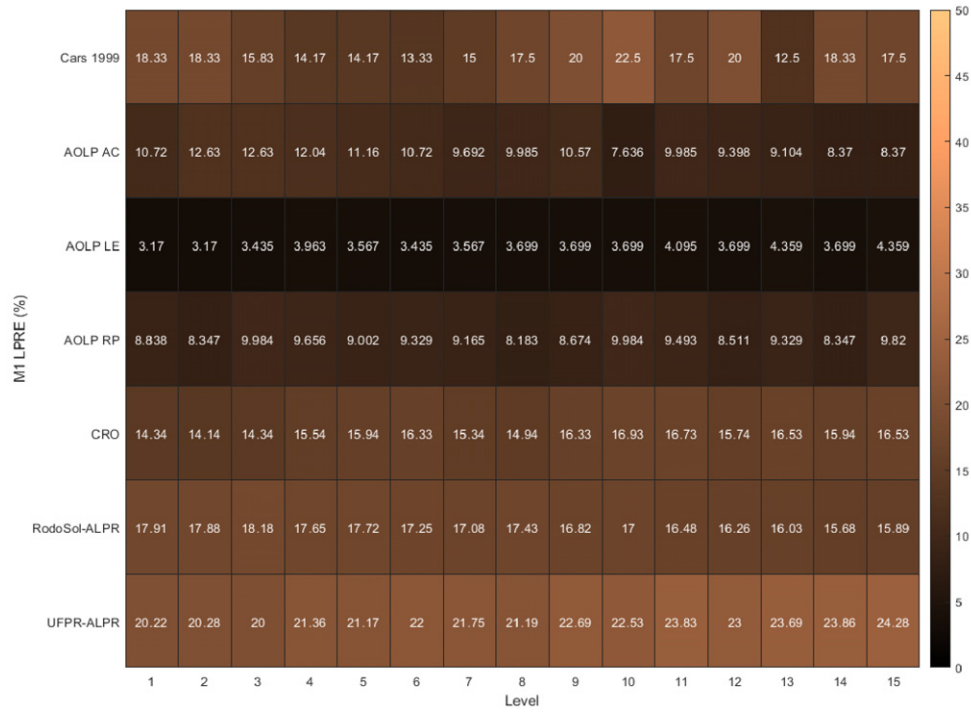


Figure A.5. The ALPR in Unconstrained Scenarios model: LPRE for read noise distortion over all datasets, except PKU dataset. The distortion levels correspond to read noise with intensity from 0.2 mV to 14.2 mV with 1 mV increments.




Figure A.6. The UltimateALPR-SDK model: TA for read noise distortion over all datasets. Distortion level zero indicates performance on original images. The remaining distortion levels correspond to read noise with intensity from 0.2 mV to 14.2 mV with 1 mV increments.

Table A.3. Averaged LPRE and LPDE across datasets of ALPR in Unconstrained Scenarios model for different weather distortion types: brightness, fog, frost, and snow.

	Level 1	Level 2	Level 3	Level 4	Level 5
LPDE (Brightness)	14.86	16.7	18.98	22.51	27.26
LPRE (Brightness)	11.77	13.11	14.73	14.84	16.15
LPDE (Fog)	20.71	22.49	24.9	26.27	34
LPRE (Fog)	11.59	12.59	13.57	12.84	11.63
LPDE (Frost)	22.67	40.04	50.05	62.69	82.13
LPRE (Frost)	12.51	12.22	11.45	11.86	6.27
LPDE (Snow)	33.54	49.76	54.87	71.68	63.91
LPRE (Snow)	15.09	14.41	15.59	14.55	12.2

Table A.4. Averaged LPRE and LPDE across datasets of HyperLPR model for different weather distortion types: brightness, fog, frost, and snow.

	Level 1	Level 2	Level 3	Level 4	Level 5
LPDE (Brightness)	0.43	0.69	0.76	1.00	1.65
LPRE (Brightness)	2.23	2.55	2.89	3.25	4.03
LPDE (Fog)	0.27	0.26	0.45	0.55	0.55
LPRE (Fog)	1.82	2.14	2.09	1.97	2.09
LPDE (Frost)	1.16	2.59	3.52	3.51	7.37
LPRE (Frost)	2.67	6.53	7.75	6.31	22.76
LPDE (Snow)	2.95	11.97	17.33	41.89	46.79
LPRE (Snow)	7.34	20.53	29.69	38.48	29.52

REFERENCES

- J. Shashirangana, H. Padmasiri, D. Meedeniya, and C. Perera, "Automated license plate recognition: A survey on methods and techniques," *IEEE Access* **9**, 11203–11225 (2021).
- R. K. Prajapati, Y. Bhardwaj, R. K. Jain, and K. K. Hiran, "A review paper on automatic number plate recognition using machine learning: An in-depth analysis of machine learning techniques in automatic number plate recognition: Opportunities and limitations," *Int'l. Conf. on Computational Intelligence, Communication Technology and Networking (CICTN)* (IEEE, Piscataway, NJ, 2023).
- S. Du, M. Ibrahim, M. Shehata, and W. Badawy, "Automatic license plate recognition (ALPR): A state-of-the-art review," *IEEE Trans. Circuits Syst. Video Technol.* **23**, 311–325 (2013).
- H. Li, P. Wang, and C. Shen, "Towards end-to-end car license plate detection and recognition with deep neural networks," *IEEE Trans. Intell. Transp. Systems* **20**, 1126–1136 (2017).
- H. Li, P. Wang, M. You, and C. Shen, "Reading car license plates using deep neural networks," *Image Vis. Comput.* **72**, 14–23 (2018).
- S. Montazzolli and C. Jung, "Real-time Brazilian license plate detection and recognition using deep convolutional neural networks," *Conf. on Graphics, Patterns and Images (SIBGRAPI)* (IEEE, Piscataway, NJ, 2017), pp. 55–62.
- Z. Xu, W. Yang, A. Meng, N. Lu, H. Huang, C. Ying, and L. Huang, "Towards end-to-end license plate detection and recognition: A large dataset and baseline," *European Conf. on Computer Vision (ECCV) ECCV 2018* (Springer, Cham, 2018), Vol. 11217.
- J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," *IEEE/CVF Computer Vision and Pattern Recognition Conf. (CVPR)* (IEEE, Piscataway, NJ, 2016), pp. 779–788.
- W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and Alexander C. Berg, "SSD: Single shot multibox detector," *European Conf. on Computer Vision* (Springer, Cham, 2016).
- S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Analysis and Machine Int.* **39**, 1137–1149 (2016).
- M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, "Spatial transformer networks," *Conf. on Neural Information Processing Systems (NeurIPS)* (Curran Associates, Red Hook, NY, 2015).
- S. M. Silva and C. R. Jung, "License plate detection and recognition in unconstrained scenarios," *European Conf. on Computer Vision (ECCV)* (Springer, Cham, 2018), Accessed: (2024) Code available on: <https://github.com/sergiomsliva/alpr-unconstrained>.
- N. Awalganar, P. Bratakke, and R. Chaugule, "Automatic license plate recognition system using SSD," *Int'l. Symposium of Asian Control Association on Intelligent Robotics and Industrial Automation (IRIA)* (IEEE, Piscataway, NJ, 2021), pp. 394–399.
- T. Saidani and Y. E. Touati, "A vehicle plate recognition system based on deep learning algorithms," *Multimedia Tools Appl.* **80** (2021).
- A. Bakshi and S. S. Udmale, "ALPR: A method for identifying license plates using sequential information," *Computer Analysis of Images and Patterns (CAIP)* (Springer, Cham, 2023).
- A. Farhat, O. Hommos, A. Al-Zawqari, A. Al-Qahtani, F. Bensaali, A. Amira, and X. Zhai, "Optical character recognition on heterogeneous SoC for HD automatic number plate recognition system," *J. Image and Video Process.* **2018** (2018).
- B. V. Kakani, D. Gandhi, and Sagar J., "Improved OCR base automatic vehicle number plate recognition using features trained neural network," *Int'l. Conf. on Computing, Communication and Networking Technologies (ICCCNT)* (IEEE, Piscataway, NJ, 2017).
- R. Laroca, E. Severo, L. A. Zanlorensi, L. S. Oliveira, G. R. Goncalves, W. R. Schwartz, and D. Menotti, "A robust real-time automatic license plate recognition based on the YOLO detector," *Int'l. Joint Conf. on Neural Networks (IJCNN)* (IEEE, Piscataway, NJ, 2018), (Accessed: 02.06. 2024) Dataset available on: <https://github.com/raysonlaroca/ufpr-alpr-dataset>.
- M. Weber, P. Perona, and C. Cars, (1999) (1.0) [Data set], 1999. CaltechDATA, <https://doi.org/10.22002/D1.20084> (Accessed: 2024) Available: <https://data.caltech.edu/records/fmbpr-ezq86>.
- G.-S. Hsu, J.-C. Chen, and Y.-Z. Chung, "Application-oriented license plate recognition," *IEEE Trans. Veh. Technol.* **62**, 552–561 (2012) Link to access the dataset: <https://github.com/AvLab-CV/AOLP>.
- University of Zagreb. Project "License plates", 2003 . (Accessed: 2024). [Online]. Available: <http://www.zemris.fer.hr/projects/LicensePlates/english/>.
- H. Li, P. Wang, and C. Shen, "Toward end-to-end car license plate detection and recognition with deep neural networks," *IEEE Trans. Intell. Transp. Syst.* **23**, 1126–1136 (2019).
- R. Geirhos, C. R. M. Temme, J. Rauber, H. H. Schutt, M. Bethge, and F. A. Wichmann, "Generalisation in humans and deep neural networks," *Conf. on Neural Information Processing Systems (NeurIPS)* (Curran Associates, Red Hook, NY, 2018).
- C. Michaelis, B. Mitzkus, R. Geirhos, E. Rusak, O. Bringmann, A. S. Ecker, M. Bethge, and W. Brendel, "Bench-marking robustness in object detection: Autonomous driving when winter is coming," *Conf. on Neural Information Processing Systems (NeurIPS)* (Curran Associates, Red Hook, NY, 2019).
- D. Dengxin and L. V. Gool, "Dark model adaptation: Semantic image segmentation from daytime to nighttime," *Int'l. Conf. on Intelligent Transportation Systems (ITSC)* (IEEE, Piscataway, NJ, 2018), pp. 3819–3824.
- J. Yim and K.-A. Sohn, "Enhancing the performance of convolutional neural networks on quality degraded datasets," *Int'l. Conf. on Digital Image Computing: Techniques and Applications (DICTA)* (IEEE, Piscataway, NJ, 2017).
- L. Kong, H. Wen, L. Guo, Q. Wang, and Y. Han, "Improvement of linear filter in image denoising," *Proc. SPIE* **9808**, 98083F (2015).
- S. Gai and Z. Bao, "New image denoising algorithm via improved deep convolutional neural network with perceptive loss," *Expert Syst. Appl.* **138**, 112815 (2019).
- E. Kim, J. Kim, H. Lee, and S. Kim, "Adaptive data augmentation to achieve noise robustness and overcome data deficiency for deep learning," *Appl. Sci.* **11** (2021).
- M. Momeny, A. M. Latif, M. A. Sarraam, R. Sheikhpour, and Y. D. Zhang, "A noise robust convolutional neural network for image classification," *Results Eng.* **10**, 100225 (2021).

- ³¹ A. Rio-Alvarez, J. de Andres-Suarez, M. Gonzales-Rodriguez, D. Fernandez-Lanvin, and B. Lopez Perez, "Effects of challenging weather and illumination on learning-based license plate detection in noncontrolled environments," *Sci. Program.* **6897345**, 16 (2019).
- ³² J. Spanhel, J. Sochor, R. Juranek, A. Herout, L. Marsik, and P. Zemcik, "Holistic recognition of low quality license plates by CNN using track annotated data," *IEEE Int'l. Conf. on Advanced Video and Signal Based Surveillance (AVSS)* (IEEE, Piscataway, NJ, 2017).
- ³³ F. Xu, C. Chen, Z. Shang, Y. Peng, and X. Li, "A CRNN-based Method for chinese ship license plate recognition," *IET Image Process.* **18** (2023).
- ³⁴ J. Yan, J. Yu, and X. Xiao, "HyperLPR3 - high performance license plate recognition framework," (2023) Accessed: (2024). [Online]. Available <https://github.com/szad670401/HyperLPR>.
- ³⁵ Doubango Telecom. UltimateALPR-SDK, (2020). Accessed: (2024). [Online]. Available: <https://github.com/DoubangoTelecom/ultimateALPR-SDK>.
- ³⁶ J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)* (IEEE, Piscataway, NJ, 2017).
- ³⁷ Z. Gu, Y. Su, C. Liu, Y. Lyu, Y. Jian, H. Li, Z. Cao, and L. Wang, "Adversarial attacks on license plate recognition systems," *Comput. Mater. Continua* **65**, 1437–1452 (2020).
- ³⁸ M. Zha, G. Meng, C. Lin, Z. Zhou, and K. Chen, "RoLMA: A practical adversarial attack against deep learning-based LPR systems," *Information Security and Cryptology. Inscrypt 2019*, Lecture Notes in Computer Science (Springer, Cham, 2020), Vol. 12020.
- ³⁹ Doubango Telecom. UltimateALPR-SDK Documentation, (2020). Accessed: (2024). [Online]. Available: <http://www.doubango.org/SDKs/anpr/docs/index.html>.
- ⁴⁰ Y. Yuan, W. Zou, Y. Zhao, X. Wang, X. Hu, and N. Komodakis, "A robust and efficient approach to license plate detection," *IEEE Trans. Image Process.* **26** (2017) Dataset available on: https://github.com/ofeeler/LPR/tree/master/pku_vehicle_dataset.
- ⁴¹ R. Laroca, E. V. Cardoso, D. R. Lucio, V. Estevam, and D. Menotti, "On the cross-dataset generalization in license plate recognition," *Int'l. Conf. on Computer Vision Theory and Applications (VIS- APP)* (2022), (Accessed: 2024) Dataset available on: <https://github.com/raysonlaroca/rodosol-alpr-dataset>.
- ⁴² J. Chen, C.-H. Tan, J. Hou, L.-P. Chau, and H. Li, "Robust video content alignment and compensation for rain removal in a CNN framework," *IEEE/CVF Computer Vision and Pattern Recognition Conference (CVPR)* (IEEE, Piscataway, NJ, 2018), pp. 6286–6295.
- ⁴³ O. Bulan, V. Kozitsky, P. Ramesh, and M. Shreve, "Segmentation- and annotation-free license plate recognition with deep localization and failure identification," *IEEE Trans. Intell. Transp. Syst.* **18**, 2351–2363 (2017).
- ⁴⁴ A. Bakshi, S. Gulhane, T. Sawant, V. Sambhe, and S. S. Udmale, "ALPR - An intelligent approach towards detection and recognition of license plates in uncontrolled environments," *Int'l. Conf. Distributed Computing and Intelligent Technology (ICDCIT)* (Springer, Cham, 2023), pp. 253–269.
- ⁴⁵ Adobe After Effects Software, (2024). (Accessed: 2024) [Online]. Available www.adobe.com/AfterEffects.
- ⁴⁶ J. E. Farrell, F. Xiao, P. B. Catrysse, and B. A. Wandell, "A simulation tool for evaluating digital camera image quality," *Proc. SPIE* **5294** (2003).
- ⁴⁷ K. Wei, Y. Fu, Y. Zheng, and J. Yang, "Physics-based noise modeling for extreme low-light photography," *IEEE Trans. Pattern Anal. Mach. Intell.* **44**, 8520–8537 (2023).
- ⁴⁸ N. Plavac, S. A. Amirshahi, M. Pedersen, and S. Triantaphillidou, "The influence of read noise on an automatic license plate recognition system," *London Imaging Meeting* **5**, 12 (2024).
- ⁴⁹ C. Kamann and C. Rother, "Benchmarking the robustness of semantic segmentation models," *IEEE/CVF Computer Vision and Pattern Recognition Conf. (CVPR)* (IEEE, Piscataway, NJ, 2020), pp. 8825–8835.
- ⁵⁰ T. Rothmeier, D. Wachtel, T. von dem Bussche-Hunnefeld, and W. Huber, "I had a bad day: Challenges of object detection in bad visibility conditions," *IEEE Intelligent Vehicles Symposium* (IEEE, Piscataway, NJ, 2023), pp. 1–6.
- ⁵¹ K. Kireev, M. Andriuschchenko, and N. Flammarion, "On the effectiveness of adversarial training against common corruptions," *Conf. on Uncertainty in Artificial Intelligence* (Microtome Publishing, Brookline, MA, 2021).
- ⁵² A. K. Tripathi and S. Mukhopadhyay, "Removal of fog from images: A review," *IETE Tech. Rev.* **29** (2012).