License Plate Images Generation with Diffusion Models

Mariia Shpir^a, Nadiya Shvai^{a,b} and Amir Nakib^{b,c}

^aNational University of Kyiv-Mohyla Academy, Kyiv, Ukrane. ^bCyclope.ai, VINCI Autoroutes, Paris, France. ^cUniversity Paris Est Créteil, Laboratoire LISSI, Vitry sur Seine, France. ORCID (Nadiya Shvai): https://orcid.org/0000-0001-8194-6196, ORCID (Amir Nakib): https://orcid.org/0000-0001-9620-9324

Abstract. Despite the evident practical importance of license plate recognition (LPR), corresponding research is limited by the volume of publicly available datasets due to privacy regulations such as the General Data Protection Regulation (GDPR). To address this challenge, synthetic data generation has emerged as a promising approach. In this paper, we propose to synthesize realistic license plates (LPs) using diffusion models, inspired by recent advances in image and video generation. In our experiments a diffusion model was successfully trained on a Ukrainian LP dataset, and 1000 synthetic images were generated for detailed analysis. Through manual classification and annotation of the generated images, we performed a thorough study of the model output, such as success rate, character distributions, and type of failures. Our contributions include experimental validation of the efficacy of diffusion models for LP synthesis, along with insights into the characteristics of the generated data. Furthermore, we have prepared a synthetic dataset consisting of 10,000 LP images, publicly available at https://zenodo.org/doi/10.5281/zenodo. 13342102. Conducted experiments empirically confirm the usefulness of synthetic data for the LPR task. Despite the initial performance gap between the model trained with real and synthetic data, the expansion of the training data set with pseudolabeled synthetic data leads to an improvement in LPR accuracy by 3% compared to baseline.

1 Introduction

LPR plays important role in traffic management, automatic ticketing, security and surveillance. For the last decade, the state-of-theart results in LPR have been achieved with deep learning methods [12, 10, 15, 25, 31]. The unprecedented success of neural networks (NN) in this and other computer vision tasks is based on the usage of large datasets. However, privacy protection laws such as GDPR [19] in European Union impose severe limitations on data collection, storage, and usage. Consequently, only a limited number of open LPR datasets are available, many of them including fewer than a thousand images [26, 16, 14, 13, 20], particularly for those of European origin.

One viable solution to restricted data availability is synthetic data. Template-based LP image generation [23] is a robust and controllable approach that, nevertheless, requires meticulously handcrafted data augmentations to transform the template into a realistic image. Furthermore, it is highly dependable on the template itself and cannot be easily transferred to a similar problem. Generative adversarial networks (GANs) have allowed researchers to automate the templateto-real-image transformation [24, 27], and to achieve higher diversity and real-life likeliness [22]. However, stability of GAN training remains a challenge.

Inspired by recent advances of diffusion models for image and video generation, we propose to use them to LP synthesis. To experimentally validate this approach, we trained a Denoising Diffusion Probabilistic Model (DDPM) [5] on a Ukrainian LP dataset consisting of 78k images. Consequently, 1000 images were generated for a detailed analysis. They were manually classified as success and failure cases. Successfully generated LPs were annotated with the corresponding LP text. Based on these annotations, a character distribution analysis was performed. Furthermore, synthetic LPs were labeled for LPR task, and corresponding model was compared to the baseline trained with a real LP dataset.

Our contributions in this research can be summarized as follows:

- we experimentally verify the viability of realistic LP synthesis with diffusion models;
- we analyze the outputs of the generative model, particularly the success rate, the character distribution, and the failure cases;
- we prepare a large synthetic dataset of 10,000 LP images that we release [21] alongside this paper;
- we empirically confirm the usefulness of synthetic data for the LPR task.

The rest of the paper is organized as follows: Section 2 discusses the related work on LP generation and open LPR datasets. Section 3 provides details on the experiments methodology, results and their analysis. Section briefly describes the dataset of synthetic Ukrainian LPs publicly released as a result of this work. Finally, Section 5 concludes the article.

2 Related work

2.1 License plate generation

Generating synthetic LP images has become a practical solution to address the restrictions imposed by privacy protection laws on data collection and usage. Over the past few years, researchers have been investigating different methods to create realistic LP images for the purpose of training deep learning models in LPR tasks.

In the early days, computer graphics techniques were employed to generate LPs by creating artificial images using pre-designed templates and fonts. These methods are still being applied due to their robustness and controllability. Notably, Silvano et al. [23] have proposed synthesizing Mersocur LP based on template and carefully designed image augmentations. These methods, although successful in producing significant amounts of data, frequently fell short in terms of diversity and realism required for training resilient LPR models. Moreover, they are difficult to extend due to their heavy dependence on a specific template.

Researchers have been investigating the potential of GANs in generating LP images that are more realistic. GAN-based approaches have demonstrated great potential in capturing the diverse range of factors present in real-world LPs, such as various fonts, lighting conditions, and backgrounds. In particular, Wang *et al.* [24] have proposed an improvement of the CycleGAN based model [30] that learns the mapping between the synthetic images generated by a script, and real images. Wu *et al.* [27] have adopted a similar approach and improve the image-to-image based Pix2Pix generative architecture [7]. Nevertheless, the challenges of generating highquality images and ensuring a stable training process have been persistent in these approaches.

In recent times, diffusion models have gained significant attention as a compelling alternative for image generation tasks. Inspired by their success in producing high-quality photos and videos, we propose using diffusion models for LP synthesis. Diffusion models have numerous advantages, such as the capability to accurately model complex data distributions and produce a wide range of high-quality samples. Through this study, we explore the possibility of using diffusion models to create realistic Ukrainian LP images and examine their potential in tackling the challenge of limited data in LPR research.

2.2 License plate datasets

Several LP datasets are openly available for research. They are commonly used for LP detection (LPD) or LPR tasks. We provide below their short description, and list them in Table 1 with the information on number of images, their origin region, and type of the annotation available. Caltech Cars 1999 dataset [26] contains 126 images of cars from the rear. These images were taken in the Caltech parking lots. Caltech Cars 2001 [16] is a car dataset containing 526 images of cars from the rear. They were taken on the freeways of southern California.

AOLP dataset was composed by Hsu *et al.* in [6] for LPD and LPR tasks. This dataset is divided into three distinctive subsets, namely access control (AC), law enforcement (LE), and road patrol (RP). AOLP-AC contains 681 images, AOLP-LE consists of 757 images, and AOLP-RP has 611 images. The origin of the data is Taiwan.

The Peking University LP dataset (PKU) contains in total 3,977 images split into five groups: G1, G2, G3, G4, G5 [29]. Bounding boxes annotations for LP are also available, thus making it well suitable for LPD task.

OpenALPR [14] dataset has three subsets according to the data origin: Europian Union (OpenALPR-EU), USA (OpenALPR-US) and Brazil (OpenALPR-BR). OpenALPR-EU contains 108 images, OpenALPR-US has 222 images, and OpenALPR-BR consists of 115 images. Annotations include LP bounding boxes, and LP texts, thus making OpenALPR dataset suitable for LPD and LPR tasks.

SSIG benchmark dataset [2] was collected by Gonçalves *et al.* It is composed of 2,000 Brazilian LPs. The corresponding 14,000 alphanumeric characters come with bounding box annotations.

CCPD proposed by Xu *et al.* in [28] is a large dataset consisting of over 300,000 Chinese LP images. It has rich annotations that contain

 Table 1. Number of images, their origin and annotation type for some commonly used LP datasets; *bbox* stands for *bounding box*.

Dataset name	# of images	Origin	Annotations
Caltech 1999 [26]	126	USA	None
Caltech 2001 [16]	526	USA	None
AOLP-AC [6]	681	Taiwan	None
AOLP-LE [6]	757	Taiwan	None
AOLP-RP [6]	611	Taiwan	None
PKU-G1 [29]	810	China	LP bboxes
PKU-G2 [29]	700	China	LP bboxes
PKU-G3 [29]	743	China	LP bboxes
PKU-G4 [29]	572	China	LP bboxes
PKU-G5 [29]	1,152	China	LP bboxes
	100	F U	LP bboxes,
OpenALPR-EU [14]	108	EU	LP text
			LP bboxes,
OpenALPR-US [14]	222	USA	LP text
	115	D '1	LP bboxes,
OpenALPR-BR [14]	115	Brazil	LP text
SSIG [2]	2,000	Brazil	LP characters bboxes
			LP bboxes,
CCPD [28]	>300,000	China	LP text,
			LP four vertices
MediaLab LPR	101	a	None originally,
database [13]	181	Greece	LP bboxes available in [17]
			LP bboxes,
KarPlate Korea [3]	3,893	Korea	LP characters bboxes,
			LP text
University of Zagreb [20]	509	Croatia	None
CARLENGI	0.175	.	LP bboxes
GAP-LP [9]	9,175	Tunisia	LP characters bboxes
	1 500	D	LP bboxes,
UFPK-ALPK [11]	4,500	Brazil	LP characters bboxes,
			LP text,
RodoSol-ALPR [12]	20,000	Brazil	LP layout,
			LP four vertices
IR-LPR-LPD [18]	20,967	Iran	LP bbox
IR-LPR-LPR [18]	27,745	Iran	LP characters bboxes

information on LP bounding box coordinates, LP text, LP four vertices locations, LP area *wrt* the whole image, tilt degree, brightness and blurriness of the LP region.

MediaLab LPR database [13] consists of 181 Greek LP images. The dataset is split into multiple categories (day color images, day grayscale images, day images with blur, day images with shadows, day images with close view, and night capture). In addition, some difficult cases are presented, such as images with multiple vehicles, dirt and shadow, shadows in the LP, truck LP captured during the day, and truck LP captured during the night. Originally, no annotations were available for this data set; however, currently the corresponding LP bounding boxes can be found in the external source [17].

University of Zagreb LP dataset was collected and released in course of License Plate Detection, Recognition and Automated Storage students' project in 2003 [20]. The image database has been prepared by using OLYMPUS CAMEDIA C-2040ZOOM digital camera. It contains over 500 images of the rear views of different vehicle types (cars, trucks, buses), taken under various lighting conditions (sunny, cloudy, rainy, twilight, night light).

GAP-LP is a Tunisian LP database presented by Kessentini *et al.* in [9]. The dataset consists of 9,175 images that were acquired with different quality cameras under different resolutions, view angles and daylight lighting conditions. The images are available at a dedicated website [8] together with LP bounding boxes and LP characters bounding boxes annotations. Consequently, GAP-LP is straightforward to be used for LPD and LPR tasks.

UFPR-ALPR is a Brazilian LP dataset presented by Laroca *et al.* in [11]. It contains 4,500 images with LP bounding boxes and LP characters bounding boxes annotations, making it well suitable for LPD and LPR tasks. The images were obtained from three different cameras, 1,500 per each one. LPs in the UFPR-ALPR dataset belong to cars and motorcycles.

The RodoSol-ALPR dataset contains 20,000 images from static cameras along the 67.5-kilometer ES-060 highway in Espírito Santo, Brazil, operated by Rodovia do Sol (RodoSol). Detailed in paper by Laroca *et al.* [12], it features various vehicle types captured under diverse conditions, including day and night, different weather, and varying distances from the camera. Every image has the following information available in a text file: the vehicle type (car or motorcy-cle), the LP layout (Brazilian or Mercosul), LP text, LP four vertices coordinates (x,y).

IR-LPR is an Iranian LP dataset collected by Rahmani *et al.* [18]. It includes 20,967 car images with the LP bounding box annotations. The total number of LP images with LP characters bounding boxes is 27,745 images.

Although some of the datasets detailed above are quite large, *e.g.* CCPD containing over 300,000 images, or RodoSol-ALPR containing 20,000 images, European datasets are small and few: MediaLab LPR database of Greek LPs has only 181 images, OpenALRP-EU contains 108 images, and the University of Zagreb dataset has 509 images. Such scarcity is explained by the strict legislation on data privacy. To meet this challenge, we propose to generate synthetic LP images with DDPMs. We hope that the availability of a large number of realistic LP images will be beneficial for ANPR research.

3 Experiments, results and discussion

In this section, we detail the experimental setup and perform the analysis of the generated images.

3.1 Experimental Setup

3.1.1 Dataset

Standardization of Ukrainian Vehicle License Plate Codes. In 1995, Ukraine standardized its vehicle LP codes to include only 12 Ukrainian Cyrillic letters with Latin alphabet visual equivalents: A, B, E, I, K, M, H, O, P, C, T, X. Additionally, a two-digit regional code was added to the LP.

In 2004, numerical regional codes were replaced with new letter codes. In 2013, further modifications were made, replacing the initial letters in existing codes as follows: A to K (except Crimea, where AK became MA), B to H, and C to I. Table 2 presents the list of LP prefixes introduced in 2004, along with their corresponding 2013 codes and associated regions.

Since 2015, a new design featuring a blue band on the left side of the plate with the letters "UA" in white below the Ukrainian national flag was introduced. In 2021, Ukrainian regions were assigned two additional letter codes, although these have not yet been widely adopted. Furthermore, since 2020, vehicles powered by electric motors without internal combustion engines are designated with the characters Y or Z in their LP numbers.

Currently, all LPs issued in Ukraine, including Soviet-era plates, remain valid and can still be legally used on the country's roads. For this study, we will focus on one-line optimized plates of Type 1 (Regular vehicles) and Subtype 1-1-1 (Electric motor-powered vehicles), with codes registered between 2004 and 2021 [1]. These specific types of plates are the most commonly observed and used in Ukraine.

Dataset Details. Our study uses a private dataset of Ukrainian LP images, each captured in different lighting conditions and angles. In total, the dataset has 78,855 images (75,654 of Type 1, and 3,201 of Subtype 1-1-1), offering a wide range of real-world scenarios. See Figure 1 for more examples from the dataset.

Table 2. Correspondance of LP prefixes to regions of Ukraine.

2004 Prefix	2013 Prefix	Region
AA	KA	Kyiv city
AB	KB	Vinnytsia
AC	KC	Volyn
AE	KE	Dnipropetrovsk
AH	KH	Donetsk
AI	KI	Kyiv
AK	MA	Crimea
AM	KM	Zhytomyr
AO	KO	Zakarpattia
AP	KP	Zaporizhzhia
AT	KT	Ivano-Frankivsk
AX	KX	Kharkiv
BA	HA	Kirovohrad
BB	HB	Luhansk
BC	HC	Lviv
BE	HE	Mykolaiv
BH	HH	Ödesa
BI	HI	Poltava
BK	HK	Rivne
BM	HM	Sumy
BO	HO	Ternopil
BT	HT	Kherson
BX	HX	Khmelnytskyi
CA	IA	Cherkasy
CB	IB	Chernihiv
CE	IE	Chernivtsi
CH	IH	Sevastopol city



Figure 1. License Plate Formats and Dataset Samples. The top section shows the format types of Ukrainian LPs used in this study, and the bottom section provides corresponding sample images from the dataset.

3.1.2 Generative model training

For our experiments, we used the Denoising Diffusion Probabilistic Model (DDPM), a class of generative models that iteratively reduce noise from a sample over a set number of steps. This process progressively transforms a sample from a noise distribution to a data distribution. We followed the training configuration proposed by the authors in the original DDPM paper, with specific adjustments tailored to our task.

Images were resized to 64x64 pixels without center cropping or

flipping. For training, we used a batch size of 64. The chosen model consists of five feature map resolutions, ranging from 64x64 to 4x4, with dimensions [128, 128, 256, 256, 512]. The model includes two convolutional residual blocks per resolution level and self-attention blocks at the 16x16 and 8x8 resolutions between the convolutional layers.

The learning rate was initially set to 10^{-4} , and subsequently changed using a cosine scheduler and a 5000-step warmup phase. Higher learning rates led to unstable training convergence. We used the AdamW optimizer with standard hyperparameters and applied EMA with a decay rate of 0.9999 to stabilize training. A dropout rate of 0.1 was applied to reduce overfitting.

The diffusion model was trained for a total of 100 epochs on a Tesla T4 GPU with 15GB of VRAM, which took approximately 30 hours to complete. For inference, the generated images were initially produced at a resolution of 64×64 pixels and later resized to 193×72 pixels to match the average aspect ratio of the training dataset.

Figure 2 provides visualization of generated LP images over various training steps. We can clearly observe the progression: in the initial steps, the model learns the overall structure of the plates, and in the later steps, it adapts the text to match the target distribution.

3.2 LP generation results

Successful Image Generation Criteria. For further analysis, we established criteria for *successful image generation*. The criteria were based on the images being fully readable and meeting the following requirements:

• The LP pattern must adhere to the format "AA0000AA," which is the standard LP format registered between 2004 and 2021.

• The prefix must correspond to a valid region code as per the Ukrainian regions listing (refer to Table 2).

• The suffix is limited to the characters A, B, C, E, I, K, M, H, O, P, T, X, Y, and Z.

Images that do not meet these criteria were classified as *failed image generation*.

Analysis of Image Generation Success and Failure Distribution. For the detailed analysis, we generated a total of 1,000 synthetic LP images using a trained diffusion model. Upon evaluation, 864 of these images were classified as successful (824 of Type 1 and 40 of Subtype 1-1-1), and 136 images were found to be unreadable or contained incorrect elements (see Figure 3 for successful examples and Figure 4 for failed examples).

Additionally, Figure 5 provides a graphical representation of the distribution rates for Type 1, Subtype 1-1-1, and failure rates for both synthetic and real images in the LP generative model.

Quantitative Generation Quality Metrics. To further assess the performance of LP generative model, we calculated the Fréchet Inception Distance (FID) [4] scores for a generated synthetic dataset. The FID score is a widely used metric for evaluating the similarity between generated and real images. The FID score for a dataset of 15,864 synthetic images is 22.47.

3.3 Character Distribution Analysis

Symbol Distribution Analysis. We analyzed the distribution of prefix, digit, and suffix symbols in synthetic and real LPs (see Figure 6). The distributions show that synthetic data closely replicates real data, with only minor differences. For instance, the prefix symbol

'A' is highly frequent in both datasets, appearing in over 30% of the images. The digit symbols are also similarly distributed, with '0' being the most common at around 12% for both synthetic and real plates. However, some inconsistencies were observed.

Specifically, the digit '8' and the letter 'B' show differences in their frequencies. In synthetic LPs, '8' appears slightly more frequently than in real plates, and 'B' is more common as a suffix in generated LPs. This discrepancy could be due to the complexity in accurately drawing these symbols. The symbol 'I,' represented as a simple line, also appears more frequently in synthetic plates. This suggests that the model may prefer simpler shapes, which are easier to generate.

Interestingly, the variance in symbol distributions for prefix symbols is minimal compared to digits and suffix symbols. This is likely because there are fewer prefix combinations in the training data, leading to less variation in their synthetic representation.

Regional Distribution Analysis We also analyzed the regional distribution of LPs from different periods and types (synthetic and real) as shown in Figure 7. Despite the complexity of regional data, the synthetic models show a high degree of similarity to the real data. However, some inconsistencies are present, which could be due to the model's difficulty in replicating the exact distribution of complex regional data. These variations highlight areas where the model could be refined for better accuracy.

Overall, despite these minor differences, the synthetic data generation process produces distributions that are largely similar to realworld data.

3.4 LPR with synthetic data

Methodology for Training Character Detection Models. To evaluate the usefulness of generated LP images for LPR task, we adopted the OCR-by-character-detection method prevalent in this domain [18, 11, 12].

We initiated the process by manually labeling 864 successfully generated synthetic LP images with bounding boxes. For a fair comparison, we also labeled 864 randomly sampled images from the training dataset. We then trained the YOLOv9-c model, chosen for character detection. The model trained on the real subset of images served as a baseline for accuracy comparison.

The YOLOv9-c model was trained using the Adam optimizer with a learning rate of 0.001 and a batch size of 16. We applied various image enhancement techniques, such as HSV adjustment, rotation, translation, scaling, shearing, and perspective transformation. To address the challenge of unbalanced character distribution, the training process incorporated image weights, allowing for image sampling proportional to the number of targets per image and inversely proportional to class frequency. The best model from this phase was selected based on the highest validation loss observed over 30 epochs. A dataset share of 10% was reserved for validation.

For the synthetic data, we trained several instances of the model based on a gradual training dataset expansion with pseudolabeling. Initially, 864 synthetic images were manually annotated and used for training. Subsequently, an additional 5,000 synthetic images were pseudolabeled using the model trained on the initial 864 synthetic images, expanding the synthetic dataset to 5,864 images. In the final phase, 10,000 more synthetic images were pseudolabeled using the model trained on the 5,864-image dataset, bringing the total to 15,864 synthetic images.

To ensure high-quality image samples during pseudolabeling, we set the confidence threshold to 0.8. Each pseudolabeled image was



Figure 2. Visualization of the generated license plate images at different training stages.



Figure 3. Examples of successful LP image generation.



Figure 4. Examples of failed LP image generation.

selected based on the successful criteria described in Section 3.2. Each synthetic dataset was used to train separate models under the same setup as the baseline. It is important to note that pseudolabeling models are sensitive to overfitting, as overfitting can lead to the model ignoring failed parts of letters and focusing only on the main symbol features and their placement.

Validation and Testing of LPR Models. Validation involved 200 real images to determine the optimal confidence threshold, maximizing LPR accuracy across the models. We considered the binary accuracy measure: 1 for a correctly read LP text and 0 for an incorrectly read LP text, regardless of the number of errors in the LP text or their type. The results of the validation experiments are shown in Table 3.

The test phase used 1,000 real LP images to evaluate the model performance. The models were applied with the optimal confidence thresholds found during the validation phase. The test results are given in Table 4. The baseline model, trained with 864 manually labeled real images, resulted in 94.1% LPR accuracy. Its direct com-



Figure 5. Distribution of image categories for synthetic (left) and real (right) LP images.

petitor, the model trained with 864 manually labeled synthetic images, showed a lower performance of 90.9%. However, as the number of synthetic images increased, the LPR accuracy improved significantly. We observe that with equal manual labeling effort (of 864 images), generating and pseudolabeling more images significantly improves the LPR accuracy. The accuracy reaches 96.3% with 5,864 synthetic images and further increases to 97.5% with 15,864 synthetic images, both of which outperform the baseline accuracy of 94.1% achieved by the model trained on real images.

These results demonstrate the high potential of synthetic data for enhancing LPR accuracy in scenarios with limited data availability, with a clear benefit from increasing the volume of synthetic training data.

4 Dataset Description

To address the challenge of data scarcity in LPR tasks and to fill the gap in available Ukrainian LP datasets, we are releasing a dataset of 10,000 synthetic Ukrainian vehicle LP images [21]. This dataset, which was instrumental in the second step of our dataset expansion process and contributed to training the highest-performing model in our experiments, is intended to serve as a robust resource for the training and validation of LPR models. By providing diverse conditions, including variations in lighting and angles, this dataset supports the exploration of model performance in real-world scenarios, ultimately aiding the development of more accurate and resilient LPR systems.

Dataset Composition The dataset contains 10,000 images, each with a resolution of 193×72 pixels. These images represent two main types of Ukrainian LPs: regular vehicles and electric motor-powered vehicles. The dataset covers a wide range of scenarios,



Figure 6. Comparison of symbol distributions in license plates. Left: Prefix symbols. Middle: Digit symbols. Right: Suffix symbols.



Figure 7. Comparison of region prefix distribution between synthetic and real datasets.

including variations in lighting conditions, viewing angles, and regional codes, ensuring coverage of the standard LP formats used in Ukraine between 2004 and 2021. The sample images shown in Figure 8) highlight different scenarios and conditions, providing a better sense of the dataset quality and variety.

Note: As the data samples are synthetically generated, there may be slight inaccuracies in the representation of the intended distance between the letters and the exact color of the LPs. While these aspects have been approximated to closely resemble real-world conditions, they might not perfectly match the specifications of actual LPs.

Data Annotation Each image in the dataset is annotated in the YOLO Darknet format, which includes precise bounding box coordinates for each character on the LP. The annotations follow the standard LP format "AA0000AA," where:

• AA represents the regional code.

- 0000 represents the numerical sequence.
- AA represents the suffix, corresponding to specific Ukrainian Cyrillic letters with Latin equivalents.

5 Conclusions

In this paper, we have proposed a novel approach for synthesizing realistic LPs using diffusion models, addressing the challenge of limited data availability in LPR tasks. Through experimental validation, we have demonstrated the efficacy of diffusion models in generating synthetic LP images that closely resemble real-world data and their usefulness for the LPR task. Our analysis of success rates, character distributions, and failure cases provides valuable insights into the capabilities and limitations of the proposed approach.

Furthermore, the creation of a synthetic dataset comprising 10,000 LP images adds to the existing resources available for LPR research, offering researchers and practitioners access to a diverse set of data

 Table 3.
 Validation results showing LPR accuracy across different confidence thresholds for models trained on the datasets. The bolded accuracy value represents the chosen optimal confidence threshold for each model, which was determined based on maximizing LPR accuracy.

Trainii	ng Dataset	Accuracy at Confidence Threshold								
Size	Туре	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
864	Real	80.0%	89.5%	93.5%	93.0%	91.0%	85.5%	79.0%	39.0%	0.0%
864	Synthetic	72.5%	86.5%	92.0%	94.0%	94.5%	89.5%	80.5%	53.8%	1.5%
5864	Synthetic	84.0%	91.5%	94.5%	96.5%	96.5%	95.5%	90.0%	71.5%	0.0%
15864	Synthetic	95.5%	95.5%	97.5%	98.5%	98.0%	97.0%	96.5%	91.5%	0.0%



Figure 8. Sample Images from the Synthetic Dataset. These images showcase the variety in lighting conditions, viewing angles, and regional codes present in the dataset.

 Table 4.
 Performance comparison of LPR models on a test set of 1,000

 real LP images. The bolded accuracy indicates the highest performance achieved among the models.

Training Dataset		Test Accuracy		
Size	Туре	itsi Attulaty		
864	Real	94.1%		
864	Synthetic	90.9%		
5864	Synthetic	96.3%		
15864	Synthetic	97.5%		

for training and evaluation purposes. By releasing this dataset alongside the paper, we aim to contribute to the advancement of LPR technology and foster further research in this area.

Looking ahead, future research directions could explore enhancements to the proposed synthesis approach, such as incorporating additional sources of variability to generate more diverse LP images. Additionally, investigating the generalization capabilities of models trained on synthetic data to real-world scenarios would be a valuable avenue for further exploration.

Overall, our work highlights the potential of synthetic data generation using diffusion models in augmenting data availability for LPR tasks and lays the groundwork for future advancements in this field.

References

- [1] Requirements for vehicle license plates. URL https://zakon.rada.gov. ua/laws/show/z0627-23. Accessed: 2024-07-20.
- [2] G. R. Gonçalves, S. P. G. da Silva, D. Menotti, and W. R. Schwartz. Benchmark for license plate character segmentation. *Journal of Electronic Imaging*, 25(5):053034–053034, 2016.

- [3] C. Henry, S. Y. Ahn, and S.-W. Lee. Multinational license plate recognition using generalized character sequence detection. *IEEE Access*, 8: 35185–35199, 2020.
- [4] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium, 2018. URL https://arxiv.org/abs/1706.08500.
- [5] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models, 2020.
- [6] G.-S. Hsu, J.-C. Chen, and Y.-Z. Chung. Application-oriented license plate recognition. *IEEE transactions on vehicular technology*, 62(2): 552–561, 2012.
- [7] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros. Image-to-image translation with conditional adversarial networks. In *Proc. of the IEEE conference* on computer vision and pattern recognition, pages 1125–1134, 2017.
- [8] Y. Kessentini, M. D. Besbes, S. Ammar, and A. Chabbouh. Gap-lp website. https://sites.google.com/site/matdbparking, 2019. Accessed: 26 March 2024.
- [9] Y. Kessentini, M. D. Besbes, S. Ammar, and A. Chabbouh. A twostage deep neural network for multi-norm license plate detection and recognition. *Expert systems with applications*, 136:159–170, 2019.
- [10] I. R. Khan, S. T. A. Ali, A. Siddiq, M. M. Khan, M. U. Ilyas, S. Alshomrani, and S. Rahardja. Automatic license plate recognition in real-world traffic videos captured in unconstrained environment by a mobile camera. *Electronics*, 11(9):1408, 2022.
- [11] R. Laroca, E. Severo, L. A. Zanlorensi, L. S. Oliveira, G. R. Gonçalves, W. R. Schwartz, and D. Menotti. A robust real-time automatic license plate recognition based on the yolo detector. In 2018 international joint conference on neural networks (ijcnn), pages 1–10. IEEE, 2018.
- [12] R. Laroca, E. V. Cardoso, D. R. Lucio, V. Estevam, and D. Menotti. On the cross-dataset generalization in license plate recognition. In *International Conference on Computer Vision Theory and Applications* (VISAPP), pages 166–178, Feb 2022. ISBN 978-989-758-555-5. doi: 10.5220/0010846800003124.
- [13] MediaLab. Medialab lpr database. https://http://www.medialab.ntua.gr/ research/LPRdatabase.html, 2007. Accessed: 22 March 2024.
- [14] OpenALPR. Openalpr dataset. https://github.com/openalpr/ benchmarks/tree/master/endtoend, 2016. Accessed: 22 March 2024.
- [15] H. Padmasiri, J. Shashirangana, D. Meedeniya, O. Rana, and C. Perera.

Automated license plate recognition for resource-constrained environments. *Sensors*, 22(4):1434, 2022.

- [16] B. Philip, P. Updike, and P. Perona. Caltech cars 2001, Apr 2022.
- [17] L. Plates. Media-lab dataset. https://universe.roboflow.com/ license-plates-vcmxh/media-lab, sep 2022. URL https://universe. roboflow.com/license-plates-vcmxh/media-lab. visited on 2024-03-23.
- [18] M. Rahmani, M. Sabaghian, S. M. Moghadami, M. M. Talaie, M. Naghibi, and M. A. Keyvanrad. Ir-Ipr: A large scale iranian license plate recognition dataset. In 2022 12th International Conference on Computer and Knowledge Engineering (ICCKE), pages 053–058. IEEE, 2022.
- [19] P. Regulation. Regulation (eu) 2016/679 of the european parliament and of the council. *REGULATION (EU)*, 679:2016, 2016.
- [20] S. Ribarić, Z. Kalafatić, T. Hrkać, V. Srebrić, D. Sokol, D. Zrno, K. Kraupner, and G. Adrinek. License plate detection, recognition and automated storage. https://www.zemris.fer.hr/projects/LicensePlates/ english/, 2003. Accessed: 26 March 2024.
- [21] M. Shpir, N. Shvai, and A. Nakib. Synthetic ukrainian license plate dataset. https://zenodo.org/doi/10.5281/zenodo.13342102, 2024. Accessed: 22 August 2024.
- [22] N. Shvai, A. Hasnat, and A. Nakib. Multiple auxiliary classifiers gan for controllable image generation: Application to license plate recognition. *IET Intelligent Transport Systems*, 17(1):243–254, 2023.
- [23] G. Silvano, V. Ribeiro, V. Greati, A. Bezerra, I. Silva, P. T. Endo, and T. Lynn. Synthetic image generation for training deep learning-based automated license plate recognition systems on the brazilian mercosur standard. *Design Automation for Embedded Systems*, 25(2):113–133, 2021.
- [24] X. Wang, Z. Man, M. You, and C. Shen. Adversarial generation of training examples: applications to moving vehicle license plate recognition. arXiv preprint arXiv:1707.03124, 2017.
- [25] Y. Wang, Z.-P. Bian, Y. Zhou, and L.-P. Chau. Rethinking and designing a high-performing automatic license plate recognition approach. *IEEE transactions on intelligent transportation systems*, 23(7):8868– 8880, 2021.
- [26] M. Weber and P. Perona. Caltech cars 1999, Apr 2022.
- [27] S. Wu, W. Zhai, and Y. Cao. Pixtextgan: structure aware text image synthesis for license plate recognition. *IET Image Processing*, 13(14): 2744–2752, 2019.
- [28] 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. In *Proceedings of the European conference on computer vision (ECCV)*, pages 255–271, 2018.
- [29] Y. Yuan, W. Zou, Y. Zhao, X. Wang, X. Hu, and N. Komodakis. A robust and efficient approach to license plate detection. *IEEE Transactions on Image Processing*, 26(3):1102–1114, 2016.
- [30] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proc. of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
- [31] Y. Zou, Y. Zhang, J. Yan, X. Jiang, T. Huang, H. Fan, and Z. Cui. License plate detection and recognition based on yolov3 and ilprnet. *Signal, Image and Video Processing*, 16(2):473–480, 2022.