Proceedings of CONF-SEML 2026 Symposium: Computational Analysis and Modeling in Complex Intelligent Syste (CONF-SEML 2026)
DOI: 10.54254/2755-2721/2026.GU34738

Vehicle Recognition in Complex Environments: A Fusion of Data Argumentation, Domain Adaptation and Image Enhancement

Miaoci Chen

Ipswich school, Ipswich, UK
teleixi9yue@163.com

Abstract. Recognizing vehicles under difficult environmental conditions is an enduring challenge to the effective implementation of intelligent transportation systems and autonomous driving technologies. In everyday life, many factors, including poor weather (rain), night-time illumination & darkness, and partial occlusion will greatly reduce the reliability of typical deep learning-based vehicle identifiers. This paper proposes a robust framework for vehicle recognition using 3 complementary strategies: data augmentation; domain adaptation using a Domain- Adversarial Neural Network (DANN); and applying an image enhancement module based on Zero-Reference Deep Curve Estimation (Zero-DCE). The UA-DETRAC dataset was used to develop a robust evaluation protocol based on 5 different scene types: sunny, rainy, night-time, low-light and occluded. At the validation stage, both ResNet50 and EfficientNet-B0 were used as baseline models. Experiments were carried out to compare each of the 3 strategies both separately and in combination. The results showed that using the image enhancement module produced the greatest improvement (8.3% & 7.9%) in coordination with the low-light & night-time scene types respectively. Domain adaptation has improved recognition performance (i.e., the accuracy of recognition) for almost every scene category (particularly those with occlusions), where the use of domain adaptation has increased recognition performance by 6.1%. In addition, combining all three approaches resulted in a very large overall performance increase from an average of 73.4% (baseline) to 87.6%, an improvement of 14.2%. The paper also looks at the time that each method takes to make a prediction and how they can be used together to improve vehicle recognition in applications.

Keywords: Vehicle recognition, complex environment, deep learning

1. Introduction

Recent developments in smart transportation systems and self-driving cars have significantly changed the way perceive the world around people [1-3]. One of the main tasks associated with understanding the environment is recognizing vehicles; therefore, it is important to identify what kind of vehicles are present from images, as it plays a critical role in monitoring traffic, navigating safely, and avoiding accidents. Recent advances in deep learning (especially using Convolutional

DOI: 10.54254/2755-2721/2026.GU34738

Neural Networks) have greatly improved the ability to accurately recognize vehicles [4-6]. With sufficient light and unobstructed views on major benchmarking datasets (e.g., UA-DETRAC, VeRi), modern algorithms have been shown to achieve better than 95% accuracy and reliability.

The functional distance between lab-based performance measurements and actual performance in the field is significant, with operational constraints making the ability of vehicle recognition systems to perform acceptably in real-life scenarios very difficult. At a minimum, all the adverse conditions that can arise in the environment where vehicles operate and interact produce some level of degradation in the vehicle recognition systems' performance. Conditions such as motion blur from the rain, the effect of raindrops on the surface of the vehicle, and the effect of reflections from wet surfaces, cause significant degradation to the performance of vehicle recognition systems. Similarly, at nighttime or low-light levels of illumination [7-9], the degradation of the vehicle recognition systems can be quite severe due to the signal-to-noise ratio (SNR), obscuring the vehicle's edges and the color of the vehicle. When a vehicle is occluded by something else (e.g., a tree, another vehicle, a building, etc.), visual cues needed to recognize the vehicle are either lost or fragmented, further degrading the performance of the vehicle recognition systems. While it is possible to demonstrate a vehicle recognition system achieves 90% accuracy in daytime, clean conditions, that accuracy can be reduced to less than 60% when operating in nighttime or rainy conditions. This level of performance degradation is completely unacceptable for autonomous vehicles, which must be able to provide reliable vehicle recognition in all environmental circumstances, including both weather and lighting.

There are numerous options for improving model robustness. Researchers have proposed many techniques to do this, including Data Augmentation (the technique of using different transformations to create more training images from an original picture), and Domain Adaptation [10]. Domain adaptation creates a representation of features that is independent of any changes in the environment via a process called adversarial training, where the representation is produced through the interaction between a Feature Extractor and a Domain Classifier. Image enhancement techniques have also been used to improve the robustness of models. The image enhancement techniques are applied to images before they are provided to a classifier in order to create a clear image, such as removing moisture from low-light images, or improving the contrast of images. While each of these is a separate method of enhancing the data, there is limited research that has been done on the number of datasets used for each of these enhancement methods or to test how they will interact when combined together.

This paper discusses the difficulties associated with evaluating vehicles in hard to reach environments. Through an exhaustive study of the literature on this subject matter, this paper has made two significant contributions to the field of identifying vehicles in hard to reach environments:

1) This paper created a multi-scenario evaluation framework using the UA-DETRAC dataset. The UA-DETRAC dataset has been divided into five discrete (or unique) types of environments (sunny, rainy, night-time [hard lighting], low-light [day or night], and occlusion [one of the five]). By providing a performance evaluation of the system for each of the discrete environments rather than providing an average performance assessment for all five environments, it is possible to provide greater performance detail based upon the recent advances in vehicle identification technology.

2) This paper provides a comprehensive comparison between three methods of improving robustness in vehicle identification: (a) data augmentation; (b) domain adaptation techniques (DANN) to develop robustness; and (c) improving image quality (using enhancements – Zero-

DOI: 10.54254/2755-2721/2026.GU34738

DCE). This paper compared the relative performance of the three methods, along with performance breakdowns and inference times for each method for all five environments.

3)A proposal for a combined approach that integrates all three methods into one seamless end-to-end training system. The results from using this method show that, on average, the combination of all three approaches gives better results than the use of any of them singly or with their pairwise combinations. The average accuracy is 87.6% across complex scenarios, which is 14.2 percentage points higher than the baseline accuracy.

2. Methodology

This paper proposes a solution to the issue of recognizing vehicles in complex environments through image classification using deep learning. The process by which this was accomplished has been divided into four major sequential components (Stages): Data Pre-processing and Partitioning by Environmental Specifics, Selecting a Baseline Model, Designing and Integrating Three Robustness Enhancing Strategies and an Organized Methodology for Training and Evaluating.

The following paragraphs provide a detailed explanation for each of these stages.

2.1. Dataset and preprocessing

This paper utilizes the UA-DETRAC (University at Albany DEtection and TRACKing) dataset which is one of the largest/most diverse public benchmarks for vehicle recognition/tracking. To effectively benchmark the model, this paper has selected more than 100,000 frames from real-world traffic scenarios along with over 140,000 annotated bounding boxes. Another important aspect of these data are that they come with extensive metadata related to the environmental conditions at the time of the capture, hence allowing us to study object recognition under different weather/lighting conditions.

As part of the preparation process before model training, this paper performed a strict preprocessing pipeline. First, this paper examined all frames to identify those frames that had incomplete or ambiguous bounding box annotations. Next, this paper partitioned the data into five disjoint subsets for scenario-specific evaluation, using a dominant environmental characteristic of each frame's content as the partitioning criterion.

- "sunny" conditions are defined as having clear weather, lots of natural light and no rain or fog.
- "rainy" conditions are defined by the presence of streaks of rainfall, wet road surfaces and reduced visibility caused by falling rain.
- "night" conditions are defined as containing images taken after local sundown, where the predominant light sources are from artificial street lighting and/or vehicle headlights.
- "low-light" conditions are defined as being in twilight, highly overcast, or scenes with insufficient natural light but not necessarily total darkness. To increase diversity among the samples, this paper also applies gamma correction to reduce brightness of some sunny samples in synthetic samples.
- "occlusion" is defined as having greater than or equal to 30% occlusion of target vehicle by other objects such as trees, poles, and other vehicles. This paper augmentations to increase the set of occlusion samples, are randomly applying rectangular masks to training images for augmentation.

For the purpose of preserving a wide variety of examples and preventing over-training on one specific dataset, each subset from above was created with 70/10/20 (70% train/10% validation/20% test) split ratio into three separate sets for train/validation/test. The validation dataset was used for

DOI: 10.54254/2755-2721/2026.GU34738

hyperparameter tuning and running early stopping methods, while the test dataset was set aside for final evaluation.

In addition to the image partition strategy, three traditional image improvement methods (preprocessing choices) for increasing input quality under poor conditions will also be applied: (i) histogram equalization as a way to stretch the contrast of the image; (ii) Retinex-based improvement to separate the illumination portion from the reflection portion of the image; and (iii) dehazing using a dark-channel method. Each method will be evaluated individually and combined; however, they will be superseded by the learnable Zero-DCE module in Section 2.3.

2.2. Baseline models

In this paper, two commonly used CNNs (Convolutional Neural Networks) models are chosen as the baseline for recognizing vehicles. Both of the chosen models have been widely used for classifying images and have complementary characteristics: efficiency in computation by using fewer resources and the ability to more accurately represent what is contained in the input image.

- ResNet50: As a representative of residual networks, ResNet50 implements skip connections (identity shortcuts), which allow gradients to flow through many of the layers in a direct fashion; thus helping to avoid vanishing gradients. With 50 layers, it has a good balance between depth and how costly it is to compute. ResNet50 has been widely used as a foundation for many recognition tasks including make/model recognition of vehicles, therefore ResNet50 serves as the major baseline for all of the experiments in this paper.
- EfficientNet-B0: EfficientNets uses a compound scaling approach, meaning they use uniform scaling across the network; the depth of the network, the width of the network and the resolution of the input images are all uniformly scaled (to create the, so called, "EfficientNets"). EfficientNet-B0 is the smallest version of the EfficientNet family and produces far fewer parameters than ResNet50 with little difference in accuracy. In this paper, by comparing ResNet50 and EfficientNet-B0, the authors hope to find out whether having a more complex model makes the proposed adaptation strategies more effective.

Both model implementations utilize weights that have been trained on ImageNet. The last fully connected layer has been removed and replaced with a new layer that will give the predicted classes of vehicles (cars, SUVs, trucks, etc.) in the UA-DETRAC dataset. Using pre-trained weights allows for a good basis for feature extraction when only limited amounts of data exist for some complex scenes.

2.3. Robustness-enhancing strategies

This paper presents 3 unique strategies used in improving vehicle recognition in complex environments. Each one addresses various sources of performance degradation.

Strategy 1: Data Augmentation: Data augmentation is a straightforward yet effective means of adding diversity to training examples without having to collect new data.

This report describes the stochastic image augmentation algorithm applied to each image during training, selecting at random the following augmentations at a probability of 0.5 from the following list of augmentations each time an image is presented:

- Random Brightness Adjustment: A uniformly sampled random value in the range [0.6, 1.4] is added as the brightness adjustment to simulate variations in lighting conditions.
- Adding Gaussian Noise: Zero-mean Gaussian noise with a standard deviation of 0.05 is added to replicate the effect of sensor noise in low-light situations.

DOI: 10.54254/2755-2721/2026.GU34738

- Applying Gaussian Blur: A Gaussian kernel (5x5) with a standard deviation of (σ) = 1.0 will be applied to the image; this simulates defocused images caused by the camera and/or rain.

- Random Occlusion: A rectangular mask will be placed in a random location on the image covering no more than thirty percent of the area of the image to simulate partial occlusion.

These augmentations will be applied online, so every transformed image will be created for each epoch. Through this process, the number of training images is effectively multiplied by a large factor and the model has thus been exposed to a more significant number of different types of visual variability and should therefore generalize better when tested on complex environments.

Strategy 2: Domain Adaptation via DANN: A change in scenery (sunny/rainy) creates a difference in the type of object you can see. For instance, something that was previously an acceptable object in one type of scenery may not be so in another. One way to reduce the effect of this domain gap is to develop a domain adversarial neural network (DANN) approach. The concept behind DANN is that a domain classifier (a neural network trained to differentiate between images from the source domain vs. the target domain) will be able to learn how to better classify images based on the features that it has been trained to learn about. The second part of DANN is that the feature extractor will produce features that are not based on the domain classifier's predicted labels (domain label). This way, the feature extractor's output is invariant to the domain label (or representation).

This paper described how to implement the DANN architecture. In the case of the DANN architecture, this paper is using the base model (ResNet50 and EfficientNet-B0) as a matching feature extractor. From the extracted features, this paper will attach 2 parallel branches to the extracted features: (i) a label classifier for predicting the level/direction of the vehicle, (ii) a domain classifier for outputting binary labels defining the vehicle's location (or source). The total loss calculated at the end of the training process (DANN model) computes the losses of both branches combined and takes into account the presence of a gradient reversal layer for enabling adversarial training. The domain labels are known to be associated with a particular scene (i.e., all images coming from scene categories will be assigned a particular domain). The goal of this adversarial training process is to force the encoder model (feature extractor) to learn features that correlate with both sources, neutralizing any inherent differences between the two domains (e.g., an image captured outside versus inside).

Strategy 3: Image Enhancement Module (Zero-DCE): This paper propose to use a learnable enhancement module in place of fixed enhancement techniques as part of their recognition pipeline. The authors specifically use Zero-Reference Deep Curve Estimation (zero-DCE), which is a lightweight, deep learning model that provides pixel-level adjustment curves to enhance the quality (i.e., brightness, contrast, colour) of the image. Zero-DCE is trained in a zero-reference setting, using non-reference loss functions that promote moderate brightness increases, contrast stretching and colour constancy.

Zero-DCE has several advantages compared to existing methods. First, it works entirely in the image space and does not require paired low-to-normal-light training images (as opposed to other techniques). Second, it performs efficiently, taking only 2-3 ms longer on a modern GPU. Finally, it can be jointly fine-tuned with the vehicle recognition model, enabling the enhancement module to adapt to the specific requirements of the downstream task. Here, Zero-DCE is used as a pre-processing step before the backbone network. The entire pipeline is trained end-to-end using only the vehicle classification data.

DOI: 10.54254/2755-2721/2026.GU34738

2.4. Training protocol and evaluation metrics

The same standardized training protocol was used for all experiments in this paper. Adam optimizer was used with an initial learning rate of 0.001. Cosine annealing learning rate schedule decreases learning rate from 0.001 to 1e-5 over 50 epochs. Batch size is equal to 32. Cross entropy loss function used to classify the vehicles. Domain classification loss weight of 0.3 used for DANN experiments after performing an initial grid search. Early stopping based on validation accuracy was implemented (patience=10 epochs). All models were run on one NVIDIA Tesla V100 GPU (16GB) using all available GPUs.

The evaluation metrics used in this paper were:

• Accuracy- Fraction of vehicle instances correctly classified divided by total number of samples tested.

• Recall- Proportion of correctly identified true positive samples for each vehicle category followed by a macro average across all vehicle categories.• F1 Score- Mean of precision and recall as a single number that gives equal weight to performance between classifiers, especially for unevenly distributed categories.

• Inference Time (in milliseconds)- Mean time taken to process one image after performing 1000 iterations of the same test so that the time can be averaged out and any warmup period on the machine is included.Four different experimental conditions were used (baseline, one strategy, two strategies, and three strategy) and for each of these conditions, there were three independent replications of the experiment and so the average and standard deviation of the training performance were computed for all of the experiments. The statistical significance of the results was assessed using the paired t-test as warranted.

3. Experimental results

3.1. Experimental design

The report includes a collection of comparison tests to determine how well the recommended techniques perform. A reference is created using normal images without the addition of augmentation or Domain Adaptation or Image Enhancement to serve as a baseline to test how well the various techniques perform.

Seven experimental scenarios have been defined as follows:

A). Data Augmentation Only

B). Domain Adaptation Only (DANN)

C). Image Enhancement Only (Zero-DCE)

D). Data Augmentation + Domain Adaptation

E). Data Augmentation + Image Enhancement

F). Domain Adaptation + Image Enhancement

G). Data Augmentation + Domain Adaptation + Image Enhancement (Total Fusion)

Results were analyzed for all five scenes (Sunny, Rainy, Nighttime, Low-Light, Occlusion). All scenarios were tested using ResNet50 and EfficientNet-B0 constructions; however only results using ResNet50 will be reported in this document, while a summary of results using EfficientNet-B0 is provided discreetly.

DOI: 10.54254/2755-2721/2026.GU34738

3.2. Quantitative results

For Table 1 displays the accuracy rate %, for ResNet50, within 5 different environmental scenarios. Each value within this table is the average of 3 independent trials with their respective standard deviations being < 1.2% (not included for clarity).

Table 1 - Accuracy Rate %'s of ResNet50 Recognition of 5 different types of scenes using various strategies to improve upon baseline (previous) performances.

Table 1. The performance of models

Scene	Baseline	A(Aug)	B(DANN)	C(Zero-DCE)	D(A+B)	E(A+C)	F(B+C)	G(A+B+C)
Sunny	80.2	83.1	84.5	82.3	86.2	85.4	86.0	88.5
Rainy	73.5	77.8	79.2	76.4	81.5	80.3	81.0	84.2
Nighttime	58.7	65.2	66.8	66.6	71.3	72.1	71.8	75.6
low-light	61.3	67.5	68.9	69.6	73.4	74.2	73.9	77.8
conclusion	69.4	74.1	75.5	72.8	78.3	77.6	78.0	81.9
average	73.4	77.5	78.6	77.5	82.1	81.9	82.1	87.6

There are three important observations to be made on Table 1. Firstly, there is a significant variation in overall performance across the three different types of scenes used for testing the algorithm. The baseline performance of the algorithm for the sunny scenes was found to be 80.2% and for the night time scenes was only found to be 58.7%. There was therefore a difference in baseline performance of 21.5%. This shows how difficult complex environments can be in terms of evaluating algorithms. Secondly, the three techniques implemented have all improved upon the baseline technique in each of the three types of scenes, this indicates that all three of the techniques have a positive influence on the overall performance of the algorithm. Finally, the amount of improvement found in performance varies between the different types of scenes. The greatest improvements were found when using condition C (Zero-DCE) for low-light (+8.3%) and night time (+7.9%) scenes, as most would expect as this technique was designed specifically to address the problem of very little light. The second greatest overall level of improvement was when using condition B (DANN), where the average gain was 5.2% for each of the three types of scenes tested; furthermore, the highest amount of gain when using DANN was found for the occlusion scene (+6.1%). The third greatest level of improvement was seen when using condition A (data augmentation), with the average amount of gain being 4.7% for the occlusion scene and 4.3% for the rain scene.

In general, when employing more than one strategy to approach a problem, the results tend to be greater than when using one strategy alone. In this case, all three pairs of combinations (D, E, and F) had pairwise average accuracies > 81.9%, with at least a 3.3% improvement from individual results to combined results. D (A+B) and F (B+C) produced the same pairwise average accuracy of 82.1%, while E (A+C) resulted in a slightly lower average accuracy of 81.9%. These results point towards good synergy among the three strategies: two of them will mutually benefit from one another when attempting to achieve results using a particular image enhancement technique.

The overall average accuracy achieved by combining the three strategies (using the 3 approaches to obtain maximum results across the various image types) was 87.6%, representing an increase of 14.2% over baseline performance levels. Furthermore, the analysis shows that the three strategies are complementary and that the combined result is greater than the separate implementation of each

DOI: 10.54254/2755-2721/2026.GU34738

of the strategies. The night scene achieved the highest percentage of improvement (16.9%), while the sunny scene achieved the minimum (8.3%).

The findings behind the model EfficientNet-B0 consistently matched those of the model ResNet50; however, the overall performance scores were lower than ResNet50 by approximately 1 to 1.8 percentage points across the experimental conditions tested. For example, the average accuracy for EfficientNet-B0 in condition G was 86.1% compared to 87.6% for ResNet50. Thus, it can be stated that smaller, efficient backbone models do benefit from fusional approaches, albeit at a smaller degree than larger, more powerful ResNet50 backbone models.

3.3. Inference time analysis

The performance of autonomous vehicles and transportation systems depends heavily upon their ability to function in real-time. This paper showed the average end result time per image from the ResNet50 model comparing 1,000 forward passes through the Tesla V100 GPU when processing images in three different conditions of enhanced image processing: Baseline (no enhancement) = 12.3ms , Condition C (Zero-DCE only) = 15.1 ms (+2.8ms), Condition G (Full Fusion) =15.3ms (+0.2ms over Zero-DCE).

From the analysis, the Zero-DCE enhancement was found to add 2.8ms additional latency, resulting in both enhancing images and enabling realtime systems with frame rates of 30 –60 frames per second (33.3–16.7 ms between frames) as specified for many other realtime applications. Therefore, using all available data augmentation techniques and using domain adaptation on every sample would be expected to give the full fusion model an opportunity to be usable for a variety of real-world applications.

3.4. Further analysis and discussion

This paper aims to give a better understanding of the rationale behind using all three of these similar approaches together compared to using two of these approaches together resulting in improved efficacy overall when combined. One aspect of achieving superior performance overall when combining all three methods is through the introduction of multiple different training distributions caused by data augmentation. By providing the model with a larger number of training examples using data augmentation, the model will become more resilient to the effects caused by differences in light level, as well as the presence of noise, occlusion and other factors that may affect the input images to the model. Even though there will be a larger number of training examples created using data augmentation to improve the model's performance, this does not remove differences between the statistical distributions of images taken during daylight hours compared to those taken during nighttime hours. Domain adaptation techniques such as DANN have been shown to generate domain-invariant features from training image sets taken during morning and evening hours, thereby helping to reduce variability in a model's performance. While DANN can aid in representation learning, it does present a number of challenges when processing low quality images (examples: completely black images due to the absence of light, or images with very poor visibility due to obstructions from rain or fog).

Zero-DCE's image enhancement tool enhances the quality of low-quality input images and will assist the domain invariant feature generation and the domain classifier in successfully classifying low-quality input images. Each of these three approaches will have a different scale with which to represent their influence on the same issue considered in 3 Levels: Image Enhancement pertains to "Input Level Issues"; Domain Invariance Feature Generation pertains to "Domain Invariant Feature

DOI: 10.54254/2755-2721/2026.GU34738

Level Issues"; and Data Augmentation pertains to "Sample Level Issues". Each of these three approaches, when utilized together, creates an acceptable and viable solution to the three separate levels of degraded input images and may have a lesser impact when only utilizing one of the three levels individually.

Section 2 also describes examples of experiments when failures in testing the Complete Fusion Method have occurred. For example, while testing the Complete Fusion Method described in Section 3.2 on vehicles, the method identified images of vehicles as vehicles when there was a high degree of ambiguity due to the presence of extremely low levels of illumination (i.e., images were so dark that it was impossible to see the vehicle shape), or because the vehicle was concealed by something; as was discussed in Section 3.4. It will be very difficult for the Complete Fusion Method to accurately determine whether what it sees as a possible vehicle is actually the shapes of vehicles, especially if the vehicle shapes cannot be confidently determined or if the vehicle was completely concealed by some large object, because it will not be able to use anything other than the thermal energy levels in the image. Images such as those above can have an enhancement capability through the Zero DCE (Zero Data Characteristic Enhancement), but will provide no useful information to identify or to find a vehicle from the image if there is no vehicle present. Data from other sensor types such as radar or LiDar and temporal data from the video images need to be collected to address these extreme situations.

Section 3 discusses how generalizable this work is across different datasets. While UA-DETRAC was used for testing this system, the methods used in this work are not unique to UA-DETRAC. For all methods tested in this work, data augmentation, DANN and Zero-DCE, are well known and have a history of being successfully migrated from one domain to another. Therefore, this work anticipates that the same benefits obtained in this research will similarly apply to other vehicle datasets (e.g., BDD100K, Cityscapes) or even object recognition applications.

4. Conclusion

The primary aim of this research is to determine how well vehicles can be identified from challenging circumstances such as nighttime, low-light, and restricted visibility. The data collected from this research demonstrates that ordinary deep learning techniques do not effectively detect vehicles in these scenarios as demonstrated by ResNet50's inability to classify vehicles during the night with an accuracy of 58.7% and with low-light conditions having an accuracy of only 61.3% when compared to daytime. As a result of these findings, there is a need for alternative methods to be developed which are practical for use in the real world.

Three different strategies were examined to improve the accuracy of the ResNet50 baseline model through the use of enhancement techniques: data augmentation, domain adaptation (DANN), and image enhancement (Zero-DCE). Each enhancement technique improved the results of the ResNet50 baseline model; data augmentation improved the accuracy overall of the model, whereas DANN specifically improved the performance of occluded images and Zero-DCE significantly improved the performance of low-light detection.

The combination of the three enhancement techniques (data augmentation, domain adaptation, and Zer0-DCE) provided the best overall results when used on the ResNet50 baseline model; in average, an accuracy of 87.6% was attained, which represents a significant increase from the baseline accuracy of ResNet50. Furthermore, the increase in the processing time for images, after applying the enhancement techniques, would be minimal and could still allow for real-time processing.

DOI: 10.54254/2755-2721/2026.GU34738

This study is limited to the examination of one dataset and image based tasks only. Future extensions of the work include examining other datasets and extending this technique to video-based recognition.

References

[1] Oladimeji, D., Gupta, K., Kose, N. A., Gundogan, K., Ge, L., & Liang, F. (2023). Smart transportation: an overview of technologies and applications. Sensors, 23(8), 3880.

[2] Zantalis, F., Koulouras, G., Karabetsos, S., & Kandris, D. (2019). A review of machine learning and IoT in smart transportation. Future Internet, 11(4), 94.

[3] Azgomi, H. F., & Jamshidi, M. (2018). A brief survey on smart community and smart transportation. In 2018 IEEE 30th international conference on tools with artificial intelligence (ICTAI) (pp. 932-939). IEEE.

[4] Li, Z., Liu, F., Yang, W., Peng, S., & Zhou, J. (2021). A survey of convolutional neural networks: analysis, applications, and prospects. IEEE transactions on neural networks and learning systems, 33(12), 6999-7019.

[5] O'shea, K., & Nash, R. (2015). An introduction to convolutional neural networks. arXiv preprint arXiv: 1511.08458.

[6] Gu, J., Wang, Z., Kuen, J., Ma, L., Shahroudy, A., Shuai, B., ... & Chen, T. (2018). Recent advances in convolutional neural networks. Pattern recognition, 77, 354-377.

[7] Yu, Y., Chen, W., Chen, F., Jia, W., & Lu, Q. (2024). Night-time vehicle model recognition based on domain adaptation. Multimedia tools and applications, 83(4), 9577-9596.

[8] Bell, A., Mantecón, T., Díaz, C., del-Blanco, C. R., Jaureguizar, F., & Garcia, N. (2021). A novel system for nighttime vehicle detection based on foveal classifiers with real-time performance. IEEE Transactions on Intelligent Transportation Systems, 23(6), 5421-5433.

[9] Wang, C. C., Huang, S. S., & Fu, L. C. (2005). Driver assistance system for lane detection and vehicle recognition with night vision. In 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (pp. 3530-3535). IEEE.

[10] Pan, S. J., Tsang, I. W., Kwok, J. T., & Yang, Q. (2010). Domain adaptation via transfer component analysis. IEEE transactions on neural networks, 22(2), 199-210.