Traffic Sign Recognition in Rainy Conditions Based on Federated Learning

Yilin Chen

Computer Science, Wuhan University of Science and Technology, 430065 Wuhan, China

Abstract. The challenge of traffic sign recognition in rainy conditions poses significant difficulties for autonomous driving systems, primarily due to obscured visibility and altered sign characteristics. To tackle this issue, this paper simulated rainy environments to improve the recognition accuracy of traffic signs in real world. This paper utilized OpenCV to preprocess images by adding a rain effect, thereby enhancing the dataset's realism. Subsequently, this study implemented a LeNet model within a Federated Learning framework, which enables decentralized training while preserving data privacy. The approach involved leveraging the Belgium Traffic Sign Classification Benchmark dataset, achieving an impressive accuracy of approximately 93% in recognizing traffic signs despite the simulated rainy conditions. The federated learning model effectively aggregated knowledge from multiple clients, resulting in a more resilient and efficient recognition system. The proposed method is demonstrated by experimental results to enhance performance in challenging weather conditions while also maintaining data privacy in machine learning applications. Overall, this paper underscores the potential of integrating federated learning with CNNs to improve traffic sign recognition capabilities.

1 Introduction

In recent years, autonomous driving has rapidly advanced, promising significant improvements in transportation by enhancing road safety, reducing traffic congestion, and minimizing human errors. To navigate and make decisions, autonomous vehicles use a combination of sensors, machine learning algorithms, and computer vision [1]. Recognizing and interpreting traffic signs is a crucial aspect of autonomous driving, which is vital in informing vehicles about speed limits, road conditions, and potential hazards. Traffic signs serve as visual cues that help ensure vehicles comply with traffic regulations and maintain safe driving practices. Traffic Sign Recognition (TSR) technology allows autonomous vehicles to detect and categorize road signs accurately in real-time, enabling them to respond appropriately to changing road conditions. TSR technology makes it possible for autonomous vehicles to recognize and categorize road signs in real-time, which enables them to respond appropriately to changes in road conditions. Given the wide variety of traffic signs and

Corresponding author: cyl202113407434@wust.edu.cn

challenges such as occlusion, varying weather conditions, and lighting, developing reliable TSR systems is essential to ensure the smooth operation of autonomous driving systems.

The development of autonomous driving technology has been significant in recent years. For instance, in [2], graph optimization techniques and deep learning models were utilized to improve path planning. The approach involved constructing urban traffic network models using graph optimization algorithms and incorporating deep learning to predict traffic patterns. This enabled dynamic route updates, optimized path selection, and reduced travel time. he application of these technologies has enhanced the path planning capabilities of autonomous driving systems in complex traffic environments. Traffic sign recognition plays a crucial role in autonomous driving because it directly affects vehicle safety and decision-making capabilities. In [3], Zhang et al. introduced an improved method for traffic sign recognition by designing a teacher network and a student network. First, the teacher model increases network depth by combining different feature channels. Then, knowledge distillation is employed to enhance the accuracy of the student model. Finally, the student model undergoes pruning to reduce the number of parameters and computational cost. This approach maintains high recognition accuracy while minimizing computational resource requirements. Moreover, in [4], the authors developed an efficient traffic sign recognition system based on Convolutional Neural Networks (CNNs). The system, accelerated by GPU training, achieved an impressive recognition rate of over 98% and was further improved to 99.15% with additional training. They designed a flexible CNN implementation capable of training large networks in a short time and incorporated a committee approach combining CNNs with Multi-Layer Perceptrons (MLPs) to further enhance recognition performance. However, these approaches primarily focus on ideal scenarios where traffic signs are clear and unobstructed. They lack deeper exploration of extreme cases, such as when weather conditions like rain, fog, or smog impair the visibility of signs. Addressing such real-world challenges is critical, as autonomous vehicles often encounter these conditions during operation. Moreover, the models developed in these studies do not incorporate privacy-preserving techniques, which are essential in modern applications. For instance, methods like federated learning, which allows for decentralized data training without compromising privacy, could be leveraged to enhance both security and performance.

In this regard, this study first used OpenCV to preprocess the dataset by simulating various challenging conditions, such as adverse weather and visibility issues. Following this, a CNN model was trained using federated learning The final model achieved a 93% accuracy as a result.

2 Method

2.1 Dataset preparation

The Belgium Traffic Sign Classification Benchmark dataset used in this study contains 2, 534 images in total, categorized into 62 different traffic sign classes [5]. All images are in RGB format and saved as PPM files, but their sizes are not uniform, varying in width and height depending on the specific image. Fig. 1 provides some sample images.

Fig. 1. Collected dataset's sample images [5].

During the data preprocessing phase, several steps were taken to ensure that the dataset was suitable and effective for training purposes. All images were resized to a uniform size of 28×28 pixels in order to maintain consistency across inputs to the model. The model's generalization ability was enhanced by using data augmentation techniques such as random horizontal flipping and random rotation. To further improve data quality, color jitter was applied to adjust the brightness, contrast, saturation, and hue of the images.

To simulate rainy conditions, a custom rain effect was introduced using the RainTransform class. The intensity parameter was set to 1, 000 to control the strength of the rain; length was set to 20 to define the length of the raindrops; angle was set to -20 to determine the angle of the rain; width was set to 3 to control the width of the raindrops; and beta was set to 0.8 to adjust the color saturation of the rain. These settings enable the RainTransform class to effectively replicate the visual effects of rain. This custom rain effect provides a simulation of rainy conditions in the training dataset, aiding the model in handling similar environmental challenges in real-world applications. Fig. 2 provides some sample images which have been preprocessed.

Fig. 2. Sample images that have been preprocessed (Photo/Picture credit : Original).

2.2 Federated learning-based CNN model

2.2.1 Preliminaries of federated learning

Federated Learning is a distributed machine learning approach where multiple clients collaboratively train a shared model while keeping their data decentralized [6-8]. The core mechanism involves training local models on each client and periodically aggregating their learned parameters to update a global model. This approach ensures data privacy and reduces the need for centralized data storage. In Federated Learning, clients independently compute updates on their local datasets, which are then communicated to a central server. The server performs aggregation, typically averaging the weights from each client to form a new global model. This iterative process continues until the model converges or meets predefined performance criteria.

2.2.2 Convolutional neural networks-based prediction

In the implementation carried out in this study, a CNN model within a Federated Learning framework was employed to address the challenge of traffic sign recognition using the BelgiumTS dataset. LeNet is the CNN architecture used, which has two convolutional layers and three fully connected layers [9]. This architecture is particularly suited for image classification tasks due to its effective feature extraction and classification capabilities. The implementation of Federated Learning involves multiple clients, each responsible for training a local model on a subset of the data. The setup includes five clients, each training on a portion of the dataset. After local training, the model weights are aggregated using Federated Averaging [10], where the average of the weights from all clients is computed to update the global model. The global model is then tested to evaluate its performance, reflecting the combined knowledge from all participating clients.

2.3 Implementation details

In the implementation of the Federated Learning-based CNN model, several crucial elements were meticulously configured to optimize performance. To ensure a balanced approach between convergence speed and stability, the learning rate was adjusted to 0.1. The Cross-Entropy Loss function was employed, as it is well-suited for classification tasks by quantifying the discrepancy between predicted probabilities and actual labels. Training was conducted over 10 local epochs for each client, providing ample iterations for the model to learn from the local data before aggregating the results. Model performance was assessed based on accuracy on the test set, which reflects the proportion of correctly classified images.

3 Results and discussion

Table 1. The comparison between regular situation and rainy situation.

Rounds	Regular Situation		Rainy Situation	
	Test Loss	Test Accuracy	Test Loss	Test Accuracy
1	1.8727	67.65%	2.9379	23.86%
2	0.5766	91.35%	1.1636	78.13%
3	0.2545	93.13%	0.6051	85.23%
4	0.1720	93.93%	0.4584	87.50%
5	0.1036	95.08%	0.3502	90.00%
10	0.0390	95.67%	0.1898	91.98%
15	0.0225	95.95%	0.1171	93.65%

In this paper, the performance of the model is compared under typical conditions without rain simulation to that under simulated rainy conditions. As shown in Table 1, the accuracy reaches 90% by the second round in the normal scenario, while it takes until the fifth round to achieve the same accuracy in the rainy scenario. The accuracy stabilizes around 95% under normal conditions, whereas it settles at approximately 93% under rainy conditions. Furthermore, the loss for the rainy scenario is consistently higher than that of the normal conditions in each round.

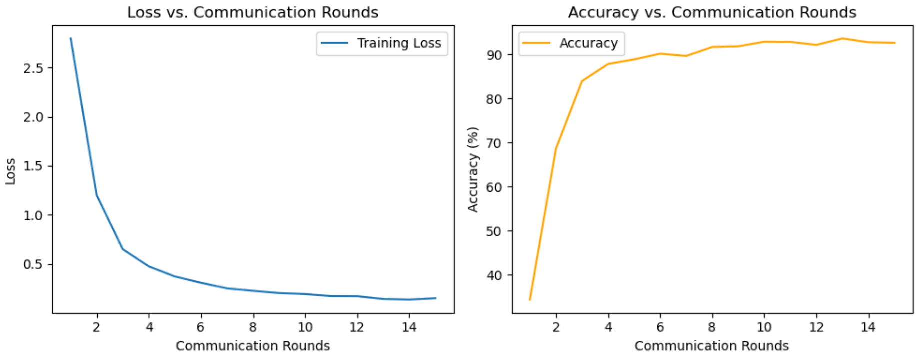

Fig. 3. The rainy situation's loss and accuracy (Photo/Picture credit : Original).

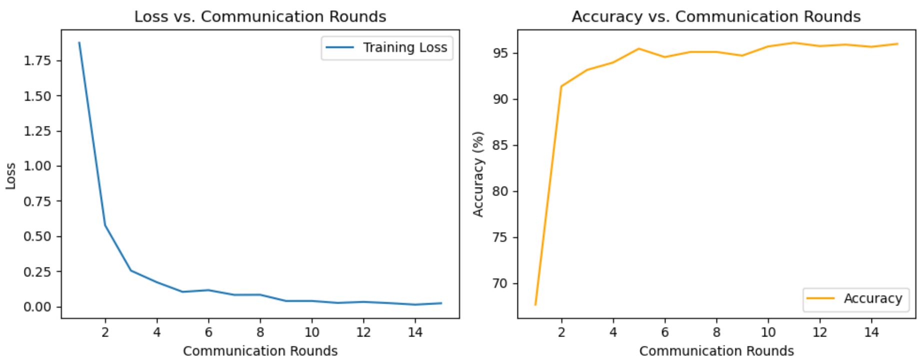

Fig. 4. The regular situation's loss and accuracy (Photo/Picture credit : Original).

As shown in Fig. 3 and Fig. 4, the trends in loss and accuracy are quite similar for both normal and rainy conditions. However, the loss is slightly higher, and the accuracy is lower in the rainy scenario compared to the normal scenario.

The introduction of noise may obscure some useful features, making the learning process more challenging for the model. Despite this reduction in accuracy, the model still performs relatively well. Additionally, the federated learning approach used in this study enhances privacy, which is a significant advantage. However, while this paper considers rainy conditions, autonomous driving scenarios can be even more complex, necessitating further exploration of factors such as exposure and nighttime conditions in future work.

The introduction of noise from the simulated rain may obscure certain useful features, making the learning process more challenging for the model. This interference can hinder the model' s ability to capture important patterns, resulting in a reduction in accuracy. Despite this decline, the model still demonstrates relatively excellent performance. The federated learning approach used in this study improves privacy by enabling decentralized training without sharing raw data, but it may also impact accuracy. The distribution of data among clients can lead to variations in model updates, potentially causing fluctuations in performance. This trade-off between privacy and accuracy is an important consideration in future work. The LeNet model used in this study may have its feature extraction capabilities compromised under conditions with added noise. The model's architecture, while effective, may not be robust enough to fully adapt to the challenges posed by rainy scenarios. This limitation suggests that more complex models could provide better feature extraction and overall performance. Also, this study relies on a simple CNN framework, which may not adequately capture the specifics of the task.

4 Conclusion

This study explored traffic sign recognition in rainy conditions, focusing on how simulated rain affects model performance. A federated learning approach using a LeNet model was proposed, employing the FedAvg strategy for averaging model weights across clients. The model's accuracy in recognizing traffic signs under simulated rainy conditions was 93%, as revealed by the experiments. Looking ahead, advanced models could be explored in future studies to enhance training outcomes. Moreover, while this paper focuses on rainy conditions, autonomous driving scenarios encompass a wider array of complexities, including variations in exposure and nighttime conditions.

References

1. E. Yurtsever, J. Lambert, A. Carballo, K. Takeda, A Survey of Autonomous Driving: Common Practices and Emerging Technologies. IEEE Access, vol. 8, pp. 58443-58469 (2020).
2. C. Liu, S. Lee, S. Varnhagen, H. E. Tseng, Path planning for autonomous vehicles using model predictive control, 2017 IEEE Intelligent Vehicles Symposium (IV), Los Angeles, CA, USA, pp. 174-179 (2017).
3. J. Zhang, W. Wang, C. Lu et al., Lightweight deep network for traffic sign classification. Ann. Telecommun., 75, 369–379 (2020).
4. D. Cireşan, U. Meier, J. Masci, J. Schmidhuber, A committee of neural networks for traffic sign classification, The 2011 International Joint Conference on Neural Networks, San Jose, CA, USA, pp. 1918-1921 (2011).
5. R. Timofte, BelgiumTS Dataset, https://btsd.ethz.ch/shareddata/ (2010).
6. L. Li, Y. Fan, M. Tse, K. Y. Lin, A review of applications in federated learning. Computers & Industrial Engineering, 149, 106854 (2020).
7. P. M. Mammen, Federated learning: Opportunities and challenges. arXiv preprint arXiv:2101.05428 (2021).
8. P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, ... S. Zhao, Advances and open problems in federated learning. Foundations and Trends in Machine Learning, 14(1–2), 1-210 (2021).
9. Y. LeCun, L. Bottou, Y. Bengio, P. Haffner, Gradient-based learning applied to document recognition. Proceedings of the IEEE, 86(11), 2278-2324 (1998).
10. T. Sun, D. Li, B. Wang, Decentralized federated averaging. IEEE Transactions on Pattern Analysis and Machine Intelligence, 45(4), 4289-4301 (2022).