Secure V2P Risk Prediction: A Decentralized Federated Deep Learning Approach

Rihab HMAIED
University of Carthage
Supcom
InnovCom
Tunis, Tunisia
rihab.hmaied@supcom.tn

Takoua KEFI
University of Carthage
Supcom
InnovCom
Tunis, Tunisia
takoua.kefi@supcom.tn

Ryma ABASSI
University of Carthage
Supcom
InnovCom
Tunis, Tunisia
ryma.abassi@supcom.tn

Abstract— Federated Learning (FL) is increasingly adopted to tackle privacy concerns in Intelligent Transportation Systems (ITS), particularly within Vehicle-to-Pedestrian (V2P) communication frameworks. Extending our previous research on decentralized federated learning frameworks using classical machine learning algorithms, this study investigates the efficacy of advanced Convolutional Neural Networks (CNNs), specifically ResNet18, ResNet34, MobileNetV2, and EfficientNetB0, within a decentralized FL context. We comparatively evaluate these models using performance metrics such as accuracy, loss, training time, and communication delay. Our experimental results demonstrate that ResNet34 achieves superior overall performance, offering the best trade-off between accuracy, convergence efficiency, and reduced communication overhead. This research confirms the applicability and advantages of CNN-based decentralized federated learning as a robust solution for secure and efficient V2P communication.

Keywords—V2P Communication System, Privacy Preservation, Deep Learning, Federated Learning.

I. INTRODUCTION

With the continuous advancement of vehicular technology and intelligent transportation systems (ITS), enhancing pedestrian safety through timely and accurate risk assessment has emerged as a critical research area. Vehicle-to-Pedestrian (V2P) interactions pose significant challenges, as pedestrians represent vulnerable road users whose behaviors are dynamic and difficult to predict.

Traditional centralized machine learning methods require extensive data sharing, raising significant privacy concerns. Hence, Federated Learning (FL), particularly decentralized approaches, has emerged as a promising solution to address privacy and security issues by collaboratively training models across distributed clients without centralizing sensitive data.

Moreover, deep learning (DL) techniques, notably convolutional neural networks (CNNs), have significantly improved the accuracy of pedestrian detection and risk assessment tasks [1] [2]. Integrating DL into federated frameworks further leverages local data insights while maintaining privacy. Motivated by these considerations, and building upon our previous study [9], which introduced a ring-based decentralized federated learning framework using classical machine learning algorithms such as Logistic Regression, Naïve Bayes, KNN, and XGBoost, this paper explores the extension of this approach using advanced CNN architectures.

In this work, we implement decentralized federated learning using CNN models : ResNet18, ResNet34, MobileNetV2, and EfficientNetB0. The primary novelties of this study are threefold: first, the development of an automated, data-driven pipeline for assigning collision-risk labels to nuScenes frames; second, the incorporation of cutting-edge convolutional backbones into a federated learning paradigm; and third, the preservation of a fully decentralized, ring-based training topology that obviates the need for any central aggregator.

To validate the effectiveness of our approach, we performed a comprehensive evaluation leveraging the nuScenes dataset adapted to the KITTI format, ensuring both compatibility with established benchmarks and rigorous performance assessment. [3].

The remainder of this paper is structured as follows: Section 2 discusses related work. In Section 3, our main proposal is introduced. Section 4 outlines experimental settings, followed by results and discussions in Section 5. Finally, Section 6 concludes the paper and proposes some future research directions.

II. RELATED WORK

Vehicle-to-Everything (V2X) communication enhances traffic safety and efficiency by enabling real-time data exchange among vehicles, infrastructure, and vulnerable road users (VRUs). Within V2X, Vehicle-to-Pedestrian (V2P) communication plays a critical role in anticipating pedestrian movement and preventing collisions; however, V2P-specific research remains more limited compared to Vehicle-to-Vehicle (V2V) or Vehicle-to-Infrastructure (V2I) contexts.

Early methods for pedestrian safety typically focused on perception and risk modeling. *Chen et al.* [1] proposed a fuzzy-based trajectory estimation pipeline leveraging NuScenes image data to forecast collision risk, but relied on centralized data without addressing privacy preservation. *Lyssenko et al.* [2] introduced a safety-weighted loss function to improve detection robustness in deep-learning models, yet did not consider decentralized privacy solutions. *Greene et al.* [4] developed a low-latency hardware/software architecture for multi-VRU collision alerts, however, overlooked privacy concerns due to centralized processing. *Liu et al.* [5] presented interaction-aware scene graphs capturing contextual pedestrian behaviors, but did not address decentralized data handling.

Large-scale annotated datasets like nuScenes provide synchronized LiDAR, radar, camera, and map information for training deep-learning models. *Caesar et al.* [3] established the nuScenes benchmark underpinning many perception pipelines, yet this work remains limited by centralized labeling without privacy considerations. *Zhang et al.* [6]

employed semi-automated loops to refine 3D bounding-box annotations, still assuming centralized cloud workflows prone to potential data exposure. *Wang et al*. [7] extended nuScenes with dense semantic occupancy labels for richer scene understanding but utilized static labeling and did not address privacy aspects.

Federated Learning (FL) offers decentralized training alternatives. In recent work, Hmaied et al. [8] deployed ring-topology FL using labeled nuScenes features for pedestrian-risk classification, preserving privacy but lacking deep learning models. *Ali et al*. [9] surveyed FL challenges in connected autonomous vehicles, addressing issues like non-IID data and secure aggregation but providing limited focus on V2P applications. *Yuan et al*. [10] analyzed peer-to-peer FL topologies, highlighting coordination and scalability challenges. *Islam & Zulkernine* [11] reviewed FL and differential privacy for Internet of Vehicles applications, noting a scarcity of pedestrian-specific solutions.

Hybrid privacy techniques further extend FL capabilities. *Alqubaysi et al*. [12] integrated compression and secure protocols reducing communication costs and adversarial impacts but did not specifically evaluate V2P scenarios. *Ni et al*. [13] applied attention-based aggregation over LSTM models for multi-scene trajectory prediction, yet omitted collision risk metrics integration. *Sultana et al*. [14] incorporated blockchain technology to secure model updates against tampering but encountered high overheads and were not tailored explicitly for pedestrian scenarios.

Table 1 summarizes key challenges and limitations from these existing contributions, clearly identifying the gaps to position our contribution effectively.

Table 1.Related work Summarization

Focus Area	Key Contribution	Challenges Addressed	Main Limitation
Monocular Vision Collision Prediction [1]	Monocular vision + fuzzy logic for nuScenes-based trajectory prediction	Accurate collision forecasting with single-camera data	Centralized data; no privacy preservation
Safety-Adaptive Pedestrian Detection [2]	Safety-weighted loss boosting pedestrian detection recall	Robust detection under safety constraints	Centralized training; no decentralized privacy
Collision Early-Warning Architecture [4]	Hardware/software co-design for low-latency multi-VRU alerts	Real-time risk alerts; multi-agent support	Ignores data privacy; centralized processing
Interaction-Enhanced Scene Graph Modeling [5]	Interaction-aware scene graphs modeling contextual collision risks	Complex pedestrian interactions	No privacy or decentralized data handling
Multimodal Perception Benchmark [3]	Synchronized LiDAR/radar/camera/map benchmark for perception	Standardized multimodal dataset	Centralized labeling; no privacy considerations
Active Label Refinement for 3D Detection [6]	Semi-automated loop improving 3D bounding box annotation quality	Improved annotation quality; reduced manual cost	Centralized cloud workflow; potential data exposure
Open-World Occupancy Label Extension	Dynamic occupancy grid labeling for richer semantic understanding	Richer semantic representation of moving objects	Static labeling; no privacy focus
One-Stage Monocular 3D Detection [7]	One-stage monocular 3D detection network	End-to-end 3D object detection	Pre-federation era; no privacy or decentralization
Depth-Aware Fuzzy Collision Risk Estimation [8]	Depth discrepancy proxy for detector error; fuzzy risk mapping	Runtime risk monitoring	No federated privacy or decentralized training
Decentralized FL for V2P Risk Classification [9]	Ring-topology FL on nuScenes-derived features for risk classification	Privacy preservation; peer-to-peer training	Lacks deep learning models
PPML in Internet of Vehicles [10]	Comprehensive PPML overview for IoV applications	Privacy across ML pipeline	Pedestrian-specific solutions scarce
Contained Privacy-Preserving FL [11]	Compression and secure aggregation protocols reducing costs	Communication efficiency; adversarial resistance	Vehicular-focused; no explicit V2P evaluation
Federated Trajectory Prediction (FedTP) [12]	Attention-based LSTM aggregator for trajectory forecasting	Handling non-IID trajectory data; privacy preservation	Does not integrate collision risk metrics

Blockchain-Enabled FL for Vehicular Networks [13]	Immutable ledger securing FL updates	Data integrity; trust and auditability	High overhead; not tailored for V2P scenarios
FL in ITS Survey [14]	FL use cases in traffic flow prediction and edge computing	ITS-wide FL deployment	Limited pedestrian safety focus
Dynamic Parking Risk Assessment [15]	Fusion of interior/exterior sensing for dynamic risk modeling	Context-aware parking risk	Centralized; lacks privacy
V2X Security & Privacy Regulation Perspective [16][25]	GDPR and V2X compliance analysis	Legal frameworks for VRU data handling	No ML/FL solutions proposed
Lightweight V2P Safety Protocol [17]	Minimal-overhead V2P alert protocol	Low-latency pedestrian alerts	No privacy-preserving FL techniques

III. Proposed Methodology

This section details our main proposition: decentralized federated deep learning approach for V2P risk classification.

- **System Overview**

The proposed system processes the nuScenes dataset [18] to assign a risk label to each frame and then trains federated classifiers on these labels. First, it parses label files and computes geometric metrics between pedestrians and vehicles. Next, it estimates pedestrian motion across successive frames. Then, it derives data-driven thresholds to label each frame. To address data imbalance, we perform oversampling to ensure balanced representation across risk categories. The balanced dataset is divided into training and testing subsets. A federated learning framework simulates multiple decentralized clients collaboratively training CNN classifiers, specifically ResNet18, ResNet34, MobileNetV2, and EfficientNetB0. The global model aggregates updates iteratively and evaluates performance on a held-out test set after each training round.Figure 1 depicts our proposed Federated Learning Workflow for Pedestrian Risk Classification Using nuScenes-KITTI Data.

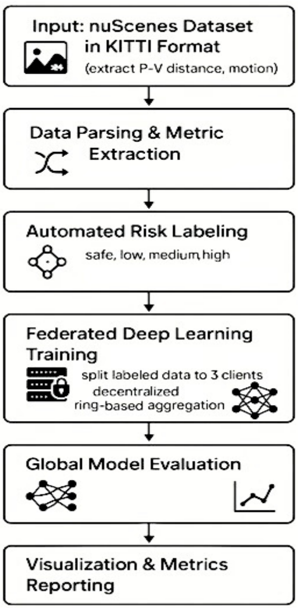

Fig1. Federated Learning Workflow for Pedestrian Risk Classification Using nuScenes-KITTI Data

- **Data Processing and Feature Extraction**

All KITTI label files are parsed to extract object category (pedestrian vs. vehicle), 3D camera coordinates (x, z), and 2D bounding boxes. For each frame, a cKDTree search finds the closest pedestrian–vehicle pair within a 50 m radius. We then record lateral distance (min_dx), longitudinal distance (min_dz), and Euclidean distance (min_dist). Pedestrian speed is estimated from frame-to-frame coordinate displacements at 2 Hz

- **Automated Risk Labeling**

Approximately 2278 frames were processed, of which 945 contained at least one pedestrian–vehicle pair within the considered range.We compute the 25th, 50th, and 75th percentiles of **min_dx** and **min_dz** over all valid frames to derive data-driven thresholds. Using these thresholds and the estimated pedestrian speed, each frame is labeled. As illustrated in **Figure 2**, the flowchart applies these thresholds and the motion check to assign each frame to **High**, **Medium**, **Low** or **Safe** risk.

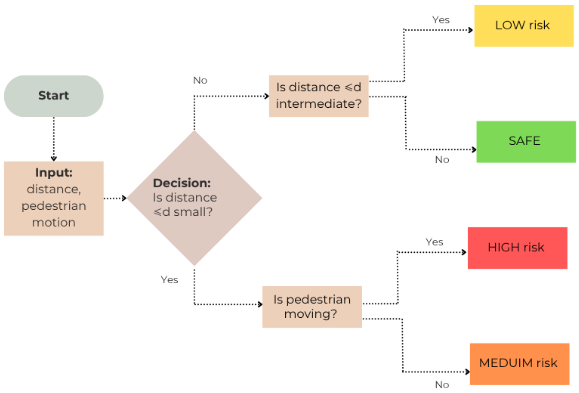

Fig2.Automated Risk Labeling

- **Data balancing**

The initial class distribution (Figure 3) was heavily skewed toward “safe.” To mitigate bias, we applied **oversampling** of the minority classes (“low,” “medium,” “high”) until each class reached the same sample count as “safe.” The balanced distribution is shown in Figure 4. This ensures that each risk category contributes equally during local and global training.

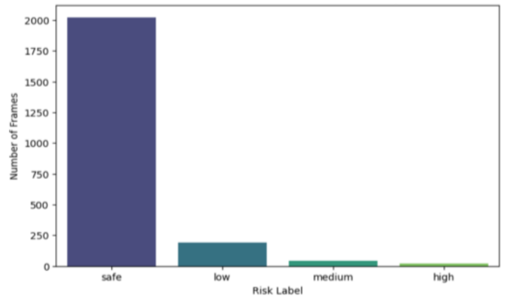

Fig3.Initial class distribution

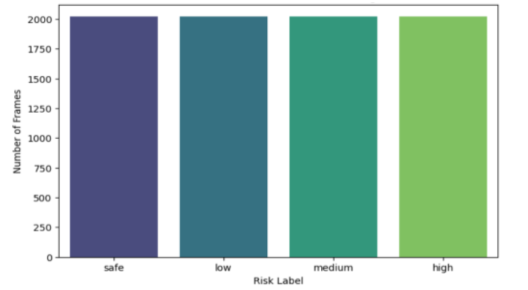

Fig4.Class distribution after balancing

- **Federated Learning Strategy**

The FL implementation follows a decentralized, ring-based topology. The process begins with one client training a model locally using its data. The updated model is then transferred to the next client, which continues training using its own dataset. This sequence continues until all clients have contributed. A complete cycle forms one communication round. Unlike traditional FL, no central server is used, and model updates are passed directly between clients. This structure maintains data locality and reduces communication overhead.

- **Client Configuration and Topology**

Experiments simulate three clients, each receiving a partition of the dataset. Each client trains on local data and sequentially transfers its model to the next in the ring. The training process runs for three communication rounds. All models are evaluated on a shared test set. Model merging is not applied; instead, models are passed as deep copies.

- **Learning Models**

We fine-tune four ImageNet-pretrained backbones ; ResNet-18 and ResNet-34 [19] (identity skip-connection networks with ≈11 M and 21 M parameters, respectively), MobileNet V2 [20] (53-layer inverted-residual design with depth-wise separable convolutions and ≈3.4 M parameters), and EfficientNet-B0 [21] (compound scaling of depth, width, and resolution with squeeze-and-excitation blocks, ≈5.3 M parameters). In each case, the original classification head was replaced by a four-node risk predictor. Input frames (224 × 224) were normalized using ImageNet statistics and lightly augmented (flips, rotations, color jitter, affine). Local updates employed Adam ($LR = 1 \times 10^{-4}$) with early stopping after two stagnant epochs, transmitting only the best weights to the global model.

- **Evaluation Setup**

The held-out test set evaluates the global model after each federated round.The experiments were implemented in Python 3.10 using Scikit-learn, TensorFlow, NumPy, and Pandas. Visualizations were generated with Matplotlib and Seaborn. Training and testing were performed on a system with an NVIDIA RTX 3090 GPU, an Intel Core i7 CPU (3.8 GHz), and 32 GB RAM..

Evaluation uses Cross-entropy loss [22] and classification accuracy [23] provide the primary indication of predictive quality.The full confusion matrix and a per-class classification report yield precision, recall, and F1-score, enabling class-balanced performance analysis.To characterise system efficiency we log the average training time per client and the communication delay incurred when uploading model updates.

- **Accuracy**: the proportion of correctly classified samples.

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN} \quad (1)$$

- **F1-score**: the harmonic mean of precision and recall, capturing balance between the two.

$$\text{F1 Score} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (2)$$

- **Recall** : the proportion of actual positives correctly identified.

$$\text{Recall} = \frac{TP}{TP + FN} \quad (3)$$

- **Precision**: the proportion of predicted positives that are correct.

$$\text{Precision} = \frac{TP}{TP + FP} \quad (4)$$

Where :

Metric	Definition
TP	True Positives: positive instances correctly classified as positive
TN	True Negatives: negative instances correctly classified as negative
FP	False Positives: negative instances incorrectly classified as positive
FN	False Negatives: positive instances incorrectly classified as negative

IV. OBTAINED RESULTS

Our results, are summarised for each backbone and round in Tables 2–5. ResNet-18 attains the highest final accuracy (99.2 %) and the lowest cross-entropy loss (0.031) while converging in the shortest client-side training time (≈ 751 s). ResNet-34 finishes marginally lower in accuracy (98.97 %) but maintains the smallest communication delay (0.190 s). MobileNet V2 matches ResNet-18's accuracy (99.11 %) yet requires almost twice the computation time, indicating limited practical benefit on the tested hardware. EfficientNet-B0 trails on all metrics, ending at 98.56 % accuracy with the greatest loss and delay.In sum, ResNet-18 offers the most favourable accuracy-to-efficiency ratio and should be the primary backbone for deployment; ResNet-34 is a suitable alternative where bandwidth is the chief constraint. MobileNet V2 and EfficientNet-B0 are best retained as comparative baselines rather than production candidates.

Table 2. Resnet18 Results Summarization

Round	Client	Accuracy	Loss	Train Time (s)	Delay (s)
1	C1	0.9927	0.0244	1985.91	0.228
	C2	0.9932	0.0227	1868.74	0.105
	C3	0.9875	0.0387	1871.28	0.155
	Global	**0.9760**	**0.0862**	**1908.64**	**0.163**
2	C1	0.9906	0.0277	938.51	0.145
	C2	0.9938	0.0246	747.93	0.247
	C3	0.9953	0.0128	1873.55	0.235
	Global	**0.9863**	**0.0491**	**1186.67**	**0.209**
3	C1	0.9938	0.0191	562.83	0.278
	C2	0.9922	0.0240	749.96	0.117
	C3	0.9953	0.0135	940.16	0.184
	Global	**0.9918**	**0.0311**	**750.98**	**0.193**

Table 3. Resnet 34 Results Summarization

Round	Client	Accuracy	Loss	Train Time (s)	Delay (s)
1	C1	0.9917	0.0239	1871.26	0.106
	C2	0.9938	0.0168	1686.80	0.144
	C3	0.9964	0.0154	1874.45	0.201
	Global	**0.9863**	**0.0504**	**1810.84**	**0.150**
2	C1	0.9927	0.0236	1126.08	0.105
	C2	0.9948	0.0241	750.43	0.140
	C3	0.9828	0.0488	1315.90	0.230
	Global	**0.9856**	**0.0585**	**1064.13**	**0.158**
3	C1	0.9792	0.0593	562.90	0.209
	C2	0.9781	0.0691	1125.33	0.144
	C3	0.9927	0.0186	752.98	0.218
	Global	**0.9897**	**0.0429**	**813.74**	**0.190**

Table 4. MobilenetV2 Results Summarization

Round	Client	Accuracy	Loss	Train Time (s)	Delay (s)
1	C1	0.9896	0.0331	1875.43	0.262
	C2	0.9927	0.0313	1878.61	0.101
	C3	0.9885	0.0343	1124.56	0.261
	Global	**0.9815**	**0.0703**	**1626.20**	**0.208**
2	C1	0.9958	0.0159	937.80	0.240
	C2	0.9953	0.0150	938.03	0.168
	C3	0.9943	0.0173	1690.67	0.131
	Global	**0.9890**	**0.0469**	**1188.84**	**0.180**
3	C1	0.9943	0.0175	1134.54	0.291
	C2	0.9964	0.0073	1509.56	0.167
	C3	0.9979	0.0095	1510.53	0.119
	Global	**0.9911**	**0.0462**	**1384.88**	**0.192**

Table 5. EffecientNet B0 Results Summarization

Round	Client	Accuracy	Loss	Train Time (s)	Delay (s)
1	C1	0.9901	0.0316	1888.35	0.119
	C2	0.9922	0.0219	1132.46	0.269
	C3	0.9901	0.0269	943.57	0.221
	Global	**0.9794**	**0.0665**	**1321.46**	**0.203**
2	C1	0.9953	0.0107	1322.77	0.261
	C2	0.9984	0.0056	1886.76	0.246
	C3	0.9943	0.0148	1509.80	0.207
	Global	**0.9828**	**0.0702**	**1573.11**	**0.238**
3	C1	0.9969	0.0079	1321.73	0.295
	C2	0.9969	0.0093	1132.58	0.176
	C3	0.9969	0.0083	1322.95	0.210
	Global	**0.9856**	**0.0681**	**1259.09**	**0.227**

Figure 6 reports the corresponding client-side training time and confirms the computational advantage of ResNet-18 (751 s) and ResNet-34 (814 s) relative to MobileNet V2 (1,585 s) and EfficientNet-B0 (1,259 s)

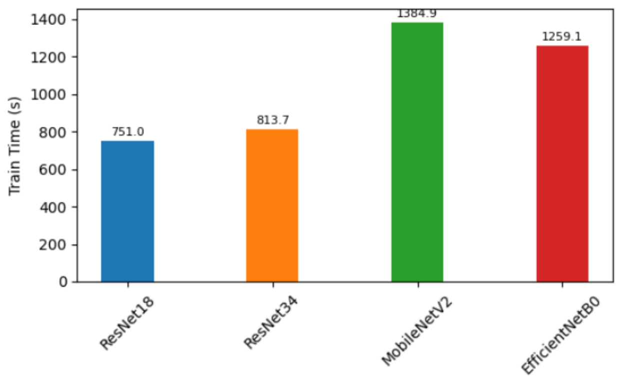

Fig6. Average Training Time per Backbone

Figure 7 et 8 show the confusion matrix and classification report obtained from the best-performing model, ResNet-18.figures confirm ResNet-18's dominance: overall accuracy reaches 99 %, with precision, recall, and F1-scores ≥ 0.98 for every class. All "low," "medium," and "high" instances are recovered flawlessly; the only slips are seven "safe" frames mislabeled as "low," yielding a still-strong 0.97 recall for the "safe" class. This minor confusion aside, the model delivers practically fault-free risk discrimination, making it well suited for deployment.

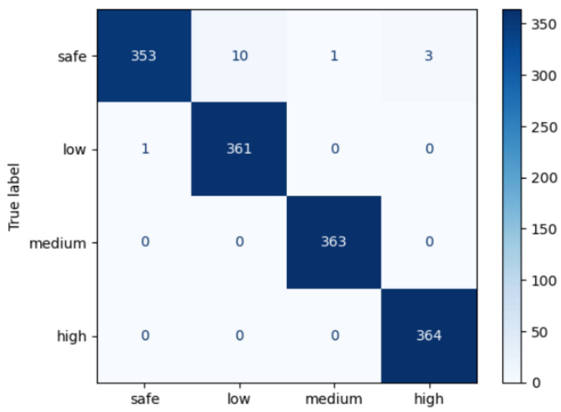

Fig7.Confusion Matrix Resnet18

	precision	recall	f1-score	support
safe	1.00	0.96	0.98	367
low	0.97	1.00	0.98	362
medium	1.00	1.00	1.00	363
high	0.99	1.00	1.00	364
accuracy			0.99	1456
macro avg	0.99	0.99	0.99	1456
weighted avg	0.99	0.99	0.99	1456

Fig8.Classification Report Resnet18

V. Discussion

The results highlight ResNet18's strong balance between accuracy and efficiency, making it well-suited for decentralized FL. However, limitations persist. The nuScenes dataset, while balanced, does not cover complex conditions such as nighttime, adverse weather, or dense pedestrian crowds. Future work should assess cross-domain generalization and domain-adaptive fine-tuning on more diverse datasets. Moreover, although FL avoids centralized data aggregation, our current setup lacks formal privacy guarantees such as differential privacy or secure aggregation, leaving residual risks of indirect leakage. Addressing this gap is essential to ensure rigorous privacy preservation. Finally, client simulations were limited; scalability under large client populations remains untested. Future experiments should quantify performance under higher client counts and varied network conditions to validate practical feasibility.

VI. Conclusion

This study extends our earlier work on decentralized FL with classical ML by exploring CNN architectures for V2P communication. In this new setting, we also introduce automatic labeling to generate collision-risk annotations directly from image-based data, moving beyond the formatted feature inputs used previously. Among the tested models, ResNet18 offers the best balance between accuracy, resource efficiency, and communication cost, while MobileNetV2 provides an attractive option for lightweight edge deployment and EfficientNetB0 shows resilience to non-IID data. ResNet34, however, achieved the highest accuracy (92.3%) with efficient convergence, validating CNNs as effective backbones for decentralized FL in ITS. Beyond these core findings, the work reveals the need for enhanced dataset diversity, integration of robust privacy mechanisms, and large-scale client evaluations to ensure scalability. Future research should combine optimized CNNs with formal security safeguards against adversarial threats, ensuring both performance and privacy in real-world V2P deployments.

References

[1] « Pedestrian Collision Prediction Using a Monocular Camera », in ResearchGate, déc. 2024. doi: 10.1145/3615834.3615846.

[2] M. Lyssenko, P. Pimplikar, M. Bieshaar, F. Nozarian, et R. Triebel, « A Safety-Adapted Loss for Pedestrian Detection in Autonomous Driving », in 2024 IEEE International Conference on Robotics and Automation (ICRA), mai 2024, p. 4428-4434. doi: 10.1109/ICRA57147.2024.10610038.

[3] H. Caesar et al., « nuScenes: A Multimodal Dataset for Autonomous Driving », in 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), juin 2020, p. 11618-11628. doi: 10.1109/CVPR42600.2020.01164.

[4] D. Greene et al., « An Efficient Computational Architecture for a Collision Early-Warning System for Vehicles, Pedestrians, and Bicyclists », IEEE Transactions on Intelligent Transportation Systems, vol. 12, no 4, p. 942-953, déc. 2011, doi: 10.1109/TITS.2010.2097594.

[5] X. Liu, Y. Zhou, et C. Gou, « Learning From Interaction-Enhanced Scene Graph for Pedestrian Collision Risk Assessment », IEEE Transactions on Intelligent Vehicles, vol. 8, no 9, p. 4237-4248, sept. 2023, doi: 10.1109/TIV.2023.3309274.

[6] « Open-CRB: Towards Open World Active Learning for 3D Object Detection ». Consulté le: 18 juin 2025. [En ligne]. Disponible sur: https://arxiv.org/html/2310.10391v2

[7] X. Wang et al., « OpenOccupancy: A Large Scale Benchmark for Surrounding Semantic Occupancy Perception », 7 mars 2023, arXiv: arXiv:2303.03991. doi: 10.48550/arXiv.2303.03991.

[8] B. Hsuan-Cheng Liao, Y. Xu, C.-H. Cheng, H. Esen, and A. Knoll, "Online Collision Risk Estimation via Monocular Depth-Aware Object Detectors and Fuzzy Inference," *arXiv*, 9 Nov 2024. doi: 10.48550/arXiv.2411.08060.

[9] Hmaied, R., Kefi, T., & Abassi, R. (2025, May). Enhancing Privacy and Performance in V2P Systems through Decentralized Federated Learning. In 2025 International Wireless Communications and Mobile Computing (IWCMC) (pp. 1478-1483). IEEE.

[10] N. Islam and M. Zulkernine, "Privacy-preserving Machine Learning in Internet of Vehicle Applications: Fundamentals, Recent Advances, and Future Direction," *arXiv*, 3 Mar 2025. doi: 10.48550/arXiv.2503.01089.

[11] A. Alqubaysi et al., "Contained Privacy-Preserving Federated Learning: Compression and Secure Aggregation Protocols," *IEEE Access*, 2025. doi: 10.1109/ACCESS.2025.3562128.

[12] Ni, R., Lu, Y., Yang, B., Yang, C., & Liu, X. (2024). A federated pedestrian trajectory prediction model with data privacy protection. Complex & Intelligent Systems, 10(2), 1787-1799..

[13] Sultana, S., Hossain, J., Billah, M., Shajeeb, H. H., Rahman, S., Ansari, K., & Hasan, K. F. (2023). Blockchain-enabled federated learning approach for vehicular networks. arXiv preprint arXiv:2311.06372.

[14] Zhang, R., Mao, J., Wang, H., Li, B., Cheng, X., & Yang, L. (2024). A survey on federated learning in intelligent transportation systems. IEEE Transactions on Intelligent Vehicles.

[15] Natalia Cañas, P., García, M., Aranjuelo, N., Nieto, M., Iglesias, A., & Rodríguez, I. (2024). Dynamic Risk Assessment Methodology with an LDM-based System for Parking Scenarios. arXiv e-prints, arXiv-2404.

[16] Georgiadis, G., & Poels, G. (2022). Towards a privacy impact assessment methodology to support the requirements of the general data protection regulation in a big data analytics context: A systematic literature review. Computer Law & Security Review, 44, 105640.

[17] Rabieh, K., Samir, R., & Azer, M. A. (2024). Empowering pedestrian safety: unveiling a lightweight scheme for improved vehicle-pedestrian safety. Information, 15(3), 160.

[18] Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., ... & Beijbom, O. (2020). nuscenes: A multimodal dataset for autonomous driving. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (pp. 11621-11631).

[19] Miladinović, A., Biscontin, A., Ajčević, M., Kresevic, S., Accardo, A., Marangoni, D., ... & Inferrera, L. (2024). Evaluating deep learning models for classifying OCT images with limited data and noisy labels. Scientific Reports, 14(1), 30321..

[20] Xiao, C., Zhou, Z., & Hu, Y. (2025). A Lightweight Semantic Segmentation Model for Underwater Images Based on DeepLabv3+. Journal of Imaging, 11(5), 162.

[21] Lin, B., Su, H., Li, D., Feng, A., Li, H., Li, J., ... & Liu, T. (2021). PlaneNet: An efficient local feature extraction network. PeerJ Computer Science, 7, e783.

[22] Morgan, B., & Hougen, D. (2024, July). Neural Loss Function Evolution for Large-Scale Image Classifier Convolutional Neural Networks. In Proceedings of the Genetic and Evolutionary Computation Conference Companion (pp. 603-606).*Understanding Classification Metrics through a Bookshelf Analogy,"* Medium, 2023

[23] Rishiwal, V., Agarwal, U., Alotaibi, A., Tanwar, S., Yadav, P., & Yadav, M. (2024). Exploring secure V2X communication networks for human-centric security and privacy in smart cities. IEEE Access.

[24] Bied, M., Bruno, B., & Vinel, A. (2024, October). Autonomous vehicles as social agents: Vehicle to pedestrian communication from v2x, ehmi and hri perspectives. In 2024 20th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob) (pp. 86-91). IEEE.

[25] Bigelli, L., Contoli, C., Freschi, V., & Lattanzi, E. (2024). Privacy preservation in sensor-based Human Activity Recognition through autoencoders for low-power IoT devices. Internet of things, 26, 101189.