

FL-P4AV: A Federated Learning-Based Privacy-Preserving Personalized Path Planning Framework for Collaborative Autonomous Ground Vehicles

Saranya C

saranyac1@srmist.edu.in

SRM Institute of Science & Technology

Janaki G

SRM Institute of Science & Technology

Research Article

Keywords: Federated Learning, Privacy-Preserving Path Planning, Autonomous Ground Vehicles, Differential Privacy, Semantic A* Algorithm

Posted Date: October 28th, 2025

DOI: <https://doi.org/10.21203/rs.3.rs-7939959/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

FL-P4AV: A Federated Learning-Based Privacy-Preserving Personalized Path Planning Framework for Collaborative Autonomous Ground Vehicles

Saranya C¹, Janaki G²

^{1,2}Department of Electrical and Electronics Engineering, SRM Institute of Science & Technology, Kattankulathur, India
¹saranyac1.research@gmail.com

Received:

Revised:

Accepted:

Published:

Abstract - The growing deployment of autonomous ground vehicles in smart cities and logistics demands secure, efficient, and context-aware navigation systems. This paper proposes FL-P4AV, a Federated Learning-Enabled Privacy-Preserving Personalized Path Planning framework designed for collaborative autonomous vehicles. Unlike traditional centralized or static path planners, FL-P4AV allows each vehicle to train a lightweight local model that predicts navigation costs based on semantic features such as obstacle density and goal proximity. The models undergo refinement via federated learning, which maintains privacy by not sharing raw data and employing differential privacy mechanisms. The semantic weights obtained are incorporated into an enhanced A* algorithm to facilitate personalized and efficient route computation. Experimental evaluations in dynamic grid environments indicate that FL-P4AV results in shorter paths, fewer inflection points, reduced turning angles, and quicker planning times relative to baseline methods. Despite the existence of privacy-preserving noise, the system maintains steady convergence and adjusts dynamically to real-time environmental changes. FL-P4AV offers a scalable and secure framework for the coordination of decentralized autonomous vehicles in path planning, indicating substantial potential for real-world applications in smart transportation systems.

Keywords - Federated Learning, Privacy-Preserving Path Planning, Autonomous Ground Vehicles, Differential Privacy, Semantic A* Algorithm.

1. Introduction

Autonomous ground vehicles (AGVs) are progressively utilized in areas including logistics, transportation, military, and smart city frameworks. These systems depend significantly on effective path planning algorithms to traverse from source to destination while circumventing barriers in dynamic and uncertain situations [2]. Conventional methods like static A* or Dijkstra's algorithm, although computationally efficient, are incapable of adapting to environmental changes or integrating contextual knowledge like terrain, congestion, or mission-specific constraints. Figure 1 illustrates the comprehensive system flow for decentralized path planning employing several autonomous ground vehicles. Each vehicle monitors its surroundings and calculates its navigation route locally, while cooperation is facilitated using a privacy-conscious model synchronization mechanism [22]. The system facilitates dynamic adaptability, multi-agent cooperation, and data privacy among dispersed navigation agents.

Centralized path planning frameworks, dependent on comprehensive environmental knowledge and centralized computation, present several challenges. Initially, a single point of failure is established, leading to diminished system robustness. Furthermore, they exhibit susceptibility to latency and scalability challenges as the agent count rises [23]. Moreover, they necessitate the ongoing exchange of

local sensory data, which engenders substantial concerns regarding data privacy, particularly when the vehicles are owned by disparate businesses or utilized in sensitive contexts [24].




Fig. 1 Path Planning Framework for Collaborative AGVs

To mitigate these constraints, there is an increasing demand for decentralized, individualized, and privacy-preserving path planning systems. Personalization enables vehicles to make navigation judgments informed by their distinct environmental observations and limitations [20]. The technology must concurrently facilitate real-time adaption to dynamically altering impediments and grid topologies. Accomplishing this while preserving the confidentiality of

individual vehicle data necessitates an innovative amalgamation of federated learning (FL) and differentiated privacy (DP).

Federated Learning allows vehicles to collaboratively learn a shared navigation model by only exchanging model parameters instead of raw data, thus preserving local privacy [26]. Differential Privacy further ensures that even shared parameters do not inadvertently leak sensitive information, by injecting statistically bounded noise during model updates [25]. FL and DP collectively establish a robust basis for developing secure, collaborative, and contextually aware navigation systems. Notwithstanding these benefits, the integration of Federated Learning (FL) and Differential Privacy (DP) in the path planning of autonomous vehicles presents novel limits and design challenges.

- Local models must be streamlined to accommodate onboard computational constraints.
- Communication overhead should be reduced throughout federated rounds.
- Privacy assurances must be reconciled with model convergence precision.
- The environment is exceedingly dynamic, requiring real-time re-strategizing.

2. Related Work

The application of Federated Learning (FL) in Autonomous Ground Vehicles (AGVs) has increased significantly, owing to its ability to preserve privacy while facilitating collaborative intelligence. Traditional controlled learning methods often make data less private, which is why FL came about as a decentralized option. Numerous studies have investigated the integration of FL in vehicular environments, with particular emphasis on privacy-preserving navigation, trajectory prediction, and multi-agent coordination [1]–[5]. Numerous studies have incorporated differential privacy into federated learning frameworks to address privacy concerns in autonomous ground vehicles. Y Li et al. [3] and [8] Y. Tian et al. [8] offer privacy-preserving methods based on noise injection techniques to shield sensitive data from vehicle training sets. These works emphasize the trade-off between utility and privacy, necessitating a balance between model accuracy and robust privacy guarantees through calibrated noise. FL has been utilized for predictive traffic management [9], feedforward control in multi-agent systems [10], and intrusion detection in Transportation IoT networks [11].

In addition to direct federated learning applications, overall progress in vehicular learning systems has established fundamental methodologies. Ming et al. [14] and Tezerjani et al. [19] present taxonomy of path planning techniques in both organized and unstructured contexts, highlighting critical problems such as obstacle avoidance, dynamic limitations, and real-time responsiveness. Concurrently, [18]

Muzi Peng et al. [18] and Yanan Ma et al. [17] present FL-augmented models for trajectory prediction and vehicular crowdsensing, illustrating how decentralized data may improve autonomous decision-making without centralized aggregation. Human-robot cooperation in intelligent environments is also enhanced by privacy-preserving federated learning approaches. Research conducted by Sun and Rahmati [15] and Mahmud et al. [16] introduces FL-based reinforcement learning frameworks that effectively coordinate agents with common goals while safeguarding sensitive control signals. These findings are also corroborated by system-level surveys on the integration of blockchain [1], secure model aggregation [13], and AI-driven intrusion detection [6]. Furthermore, the integration of reinforcement learning and deep learning with federated learning has unveiled new avenues for autonomous guided vehicles. Advanced routing, temporal modeling, and individualized driving behaviors are currently integrated into federated frameworks that adhere to the communication, latency, and privacy limitations characteristic of real-world implementations.

In summary, the literature supports a strong foundation for integrating FL and differential privacy into AGV navigation systems. However, existing solutions often lack adaptive personalization and semantic contextualization in path planning. The FL-P4AV framework builds upon this prior work by introducing privacy-preserving semantic cost estimation for real-time, collaborative path planning tailored for AGVs operating in dynamic environments. This study explicitly defines the targeted issue as follows:

- Empower each autonomous ground vehicle to calculate a tailored, efficient, and secure navigation route from its origin to destination while safeguarding the confidentiality of its local data.
- Function in dynamic grid contexts by circumventing static and mobile impediments via real-time perception, without disclosing raw observations—only anonymised model changes.
- Enhance performance measures such as route length, turning angle, and planning duration, while maintaining collaborative learning stability across all decentralized agents.

This study presents the following key contributions:

- Introduces FL-P4AV, a federated learning system designed for privacy-preserving, individualized route planning inside autonomous vehicle networks.
- Integrates a lightweight, semantic-aware cost prediction model with a modified A* algorithm to improve route smoothness and contextual awareness.
- Utilizes differential privacy to protect model parameters during federated aggregation, maintaining data confidentiality while preserving accuracy.
- Exhibits the efficacy of the methodology via simulation-based assessment in dynamic contexts,

accompanied by a comprehensive study of route metrics, convergence behavior, and privacy-performance trade-offs.

3. Materials and Methods

The proposed system, FL-P4AV (Federated Learning-based Privacy-Preserving Personalized route Planning for Autonomous Vehicles), gives a modular and intelligent approach to tackle the issues of safe, decentralized route planning in collaborative autonomous navigation. Unlike traditional centralized systems, which rely on raw data exchange and centralized computing, FL-P4AV allows each vehicle to learn navigation strategies autonomously by processing local sensor data and participating in federated learning rounds without releasing private information. This decentralized architecture promotes scalability, facilitates real-time responsiveness, and improves vehicle-level personalization according to specific environmental conditions. The system consists of interconnected modules: environment setup, lightweight learning, federated aggregation, differential privacy enforcement, semantic path planning, dynamic obstacle handling, and metric-based evaluation. The FL-P4AV framework facilitates secure and collaborative path planning for autonomous ground vehicles while maintaining data privacy [27]. This framework establishes a strong basis for scalable, privacy-conscious, and contextually aware autonomous mobility in dynamic settings.

3.1. System Architecture

FL-P4AV's architecture is structured into three fundamental layers: The Autonomous Vehicle Layer, the Federated Intelligence Layer, and the Semantic Path Planning Layer. At the vehicular level, each autonomous vehicle gathers localized environmental data, including obstacle distribution and goal distance, using onboard sensors and systems. This data is utilized to train a streamlined local neural network that encapsulates semantic patterns pertinent to navigation. To ensure privacy, the vehicle incorporates differential privacy (DP) noise into its model parameters prior to dissemination. These privacy-preserving updates are sent to the Federated Aggregation Layer, where a central engine securely consolidates all changes into a unified global model. The global model is thereafter sent to all participating cars, augmenting their learning while safeguarding individual raw datasets. The revised local models are employed in the Semantic Path Planning Layer, where each vehicle produces a semantic cost map utilized by an advanced A* algorithm to calculate optimum and individualized navigation routes. A Dynamic Obstacle Handling Module facilitates re-planning in response to real-time environmental changes. Path quality is ultimately assessed through key metrics, including path length, planning time, turning angles, and complexity. The outputs are transmitted to a Visualization Engine for performance analysis.




Fig. 2 Proposed FL-P4AV Architecture

Figure 2 illustrates that the architecture provides a distinct modular progression from sensor-level data collecting to federated learning and intelligent decision-making. The versatility and privacy-centric design of the FL-P4AV architecture provide it a viable and novel option for safe multi-agent navigation systems.

3.2 Module Descriptions

The FL-P4AV system has seven interrelated elements. Each module is designed to execute a specific role within the comprehensive route planning process, while safeguarding privacy, facilitating real-time flexibility, and enhancing path efficiency.

3.2.1 Environment Grid Initialization

The Environment Grid Initialization module constructs the spatial representation of the navigation space used throughout the FL-P4AV framework. The environment is modelled as a discrete 2D matrix $G \in \mathbb{Z}^{N \times N}$, where each grid cell is assigned a semantic label indicating its role in the navigation task. This includes free space, static obstacles, the vehicle's source position, and the destination point. Grid sizes of 20×20 , 30×30 , and 50×50 were employed to simulate varying levels of complexity. To increase realism and experimental validity, structural layouts and obstacle configurations were also inspired by publicly available datasets—namely the Waymo Open Dataset [31] and the KITTI Stereo Dataset [32]. These datasets provided real-world spatial patterns and object placements, which were adapted into grid-based abstractions suitable for training and evaluation within the FL-P4AV simulation environment. The semantic encoding of these cells is defined in Equation (1):

$$G(i, j) = \begin{cases} 0, & \text{if } (i, j) \text{ is a free cell} \\ 1, & \text{if } (i, j) \text{ contains an obstacle} \\ 2, & \text{if } (i, j) = (x_s, y_s) \text{ Source} \\ 3, & \text{if } (i, j) = (x_d, y_d) \text{ Destination} \end{cases}$$

To simulate a realistic navigation challenge, a total of O obstacles are distributed randomly across the grid. These obstacle coordinates are sampled from a uniform distribution over valid grid indices, as shown in Equation (2):

$$(x_o^{(k)}, y_o^{(k)}) \sim \mathcal{U}(0, N-1), \quad \text{for } k = 1, 2, \dots, O$$

To ensure at least one valid path can exist between the source and destination, the number of obstacles is constrained as per Equation (3):

$$O \leq N^2 - 2$$

The source and destination coordinates are chosen from contrasting quadrants to enhance spatial traversal and increase path complexity. The assignments are represented in the grid according to Equation (4):

$$G(x_s, y_s) = 2, \quad G(x_d, y_d) = 3$$

In dynamic scenarios, the grid changes over distinct time intervals. At each timestep t , obstacles can transition to neighboring cells according to a specified probability p_{move} . The grid's dynamic update is regulated by the stochastic function presented in Equation (5):

$$G_t = f_{\text{dynamic}}(G_{t-1}, p_{\text{move}})$$

The dynamic behavior is essential for validating the real-time reactivity of the FL-P4AV path planner in varying environmental conditions. This module offers a structured and adaptable framework for the testing and evaluation of decentralized, privacy-preserving navigation strategies.

3.2.2 Lightweight Local Model

The Lightweight Local Model module allows each autonomous vehicle to acquire navigation-relevant context from its surrounding environment while ensuring that no raw data is shared. This module aims to predict the traversal cost of a grid cell by analyzing local features, including cell position, obstacle density, and estimated distance to the goal. Each vehicle employs a linear regression model, characterized by its minimal computational requirements, which facilitates real-time inference on embedded hardware. The model receives a feature vector $\vec{x} \in R^4$ that encapsulates the semantic characteristics of a grid cell, and it produces a predicted cost $\hat{c} \in R$. The predictive model is represented in a linear format, as shown in Equation (6):

$$\hat{c} = \vec{w} \cdot \vec{x} + b$$

Here, $\vec{w} \in R^4$ is the weight vector and $b \in R$ is the bias term. The model is trained using local training samples $(\vec{x}^{(i)}, c^{(i)})$ is the ground truth cost of a cell estimated from vehicle observations. The training objective is to minimize the mean squared error (MSE) between the predicted and actual costs, defined in Equation (7):

$$L = \frac{1}{n} \sum_{i=1}^n (\widehat{c}^{(i)} - c^{(i)})^2$$

The model parameters $\vec{w}$ and b are updated using gradient descent. For a learning rate η , the update rules for a single batch are described in Equations (8) and (9):

$$\vec{w} \leftarrow \vec{w} - \eta \cdot \nabla_{\vec{w}} L$$

$$b \leftarrow b - \eta \cdot \nabla_b L$$

Feature extraction is crucial for model performance. The feature vector $\vec{x}$ includes: x_n, y_n : normalized coordinates of the cell, D_o : local obstacle density within a fixed-radius window, D_t : normalized Euclidean distance to the destination. Obstacle density D_o is computed by counting the number of obstacle cells within a square window centered at the cell (x, y) , normalized by the total number of cells in that window, as given in Equation (10):

$$D_o = \frac{1}{(2r+1)^2} \sum_{i=-r}^r \sum_{j=-r}^r \mathbb{1}_{\{G(x+i, y+j)=1\}}$$

The final trained model is used by each vehicle to generate a semantic weight map of the grid, representing predicted costs for path planning. These weights directly influence the heuristic in the A* algorithm, enabling personalized navigation adapted to local terrain and contextual risks.

3.2.3 Federated Aggregation

The Federated Aggregation module facilitates joint model training among various independent vehicles, ensuring that data remains local and private. Rather than sending unprocessed sensor data, each vehicle autonomously develops its local model based on its individual observations and subsequently shares solely the model parameters—namely, the weight vector $\vec{w}_i$ and bias b_i . These updates are then securely aggregated by a central server to form a global model. This approach ensures scalability, reduces communication overhead, and prevents sensitive data exposure during the learning process. Let N represent the total number of participating vehicles. Each vehicle $i \in \{1, 2, \dots, N\}$ sends its updated parameters $(\vec{w}_i, b_i)$ to the

federated server. The global weight vector $\vec{w}_g$ and global bias b_g are computed using simple arithmetic averaging, as shown in Equations (11) and (12):

$$\vec{w}_g = \frac{1}{N} \sum_{i=1}^N \vec{w}_i$$

$$b_g = \frac{1}{N} \sum_{i=1}^N b_i$$




Fig. 3 Federated Aggregation Process in FL-P4AV Framework

Figure 3 depicts the federated aggregation process inside the FL-P4AV system. Each autonomous car transmits differentially private model updates to the central aggregator. The data is averaged to create a global model, which is then disseminated to all cars, facilitating collaborative learning while preserving data privacy. Upon aggregation, the global model $(\vec{w}_g, b_g)$ is redistributed to all vehicles. Each vehicle integrates the global knowledge into its local model through a weighted update to balance local specialization and global generalization. This is achieved using an interpolation factor $\alpha \in [0,1]$, as formulated in Equations (13) and (14):

$$\begin{aligned} \vec{w}_i &\leftarrow \alpha \cdot \vec{w}_i + (1 - \alpha) \cdot \vec{w}_g \\ b_i &\leftarrow \alpha \cdot b_i + (1 - \alpha) \cdot b_g \end{aligned}$$

This mechanism ensures that each vehicle retains its locally learned biases while benefiting from the collaborative intelligence of the entire fleet. The federated learning process is performed over multiple communication rounds, with the model accuracy and convergence improving incrementally.

3.2.4 Privacy-Preserving Noise Injection

To maintain the confidentiality of local data from each vehicle during collaborative learning, the Privacy-Preserving Noise Injection module employs differential privacy (DP)

methods prior to the transmission of model parameters to the federated server. The primary aim of this module is to introduce calibrated random noise into the model's weights and biases, hence preventing the reverse-engineering of the final shared parameters to disclose critical training data.




Fig. 4. Privacy-Preserving Noise Injection Workflow

Figure 4 depicts the process by which each vehicle introduces calibrated Gaussian noise to its local model parameters prior to transmitting them to the federated aggregator. The procedure guarantees differential privacy by altering both weights and biases, thereby safeguarding sensitive training data while facilitating collaborative learning. Differential privacy ensures that the inclusion or exclusion of any one data point in a vehicle's training dataset does not materially affect the resultant model. This is accomplished by perturbing each parameter using zero-mean Gaussian noise. Let $\vec{w}_i \in R^d$ and $b_i \in R$ denote the weight vector and bias of the local model from vehicle i , respectively. The parameters $\vec{w}_i'$ and b_i' , subjected to noise injection, are calculated using Equations (15) and (16):

$$\vec{w}_i' = \vec{w}_i + \mathcal{N}(0, \sigma^2 I)$$

$$b_i' = b_i + \mathcal{N}(0, \sigma^2)$$

Here, $\mathcal{N}(0, \sigma^2)$ denotes Gaussian noise with zero mean and variance σ^2 , and I is the identity matrix. The variance σ^2 is controlled by the privacy budget ϵ , which determines the trade-off between privacy protection and model utility. A smaller ϵ leads to greater privacy but may reduce learning accuracy, whereas a larger ϵ provides better model utility at the cost of privacy. The privacy guarantee of the mechanism can be formally expressed as:

$$\Pr[\mathcal{M}(D_1) \in S] \leq e^\epsilon \cdot \Pr[\mathcal{M}(D_2) \in S] + \delta$$

where $\mathcal{M}$ is the randomized mechanism (in this case, the noise-injected model), D_1 and D_2 are neighboring datasets differing in one record, and δ is a small failure probability. This privacy-preserving mechanism ensures that even if adversaries gain access to the shared parameters, they cannot infer individual vehicle-specific data. The noise-injected parameters $\vec{w}_i'$ and b_i' are then forwarded to the Federated Aggregation module as described earlier.

3.2.5 Semantic-Aware A* Path Planner

The Semantic-Aware A* Path Planner module enhances the traditional A* search method by including semantic cost estimates derived from the local model. Conventional A* employs a heuristic derived from Euclidean or Manhattan distance to approximate the cost from a certain node to the target. In practical settings, purely geometric heuristics may inadequately represent navigation complexity arising from obstacles, terrain features, or contextual conditions. This module incorporates a semantic cost map, produced by the vehicle's local model, into the A* algorithm to enhance route planning with greater realism and personalization. Define the current node as n , the starting node as s , and the goal node as g . The classical A* evaluation function is altered in FL-P4AV to incorporate the semantic cost prediction $c(n)$, leading to the composite cost function as defined in Equation (18):

$$f(n) = g(n) + h(n) + \lambda \cdot c(n)$$

In this formulation: $g(n)$: accumulated cost from the start node s to the current node n , $h(n)$: heuristic estimate of the cost from node n to goal g (e.g., Euclidean distance), $c(n)$: semantic cost predicted by the lightweight local model, λ : scaling coefficient used to control the influence of the semantic cost term. The semantic cost $c(n)$ reflects the complexity or risk of navigating through node n , based on inputs such as obstacle density and terrain proximity. The heuristic $h(n)$ typically uses the Euclidean distance:

$$h(n) = \sqrt{(x_n - x_g)^2 + (y_n - y_g)^2}$$

During pathfinding, the node with the lowest $f(n)$ value is selected for expansion, thereby balancing geometric efficiency with semantic safety and personal relevance. This hybrid formulation makes the planner context-aware, allowing it to adapt to various terrain complexities and vehicle-specific navigation preferences. Once the goal node is reached, the algorithm reconstructs the path $P = \{s, \dots, g\}$ by backtracking through the parent nodes. This path is then passed to the Path Quality Analyzer module for further evaluation of its efficiency and smoothness [29].

3.2.6 Dynamic Obstacle Handling

The Dynamic Obstacle Handling module simulates time-varying environmental changes by introducing movement in obstacle positions during the simulation timeline. In real-world navigation, autonomous vehicles often encounter unpredictable obstacles such as pedestrians, other vehicles, or mobile objects. To ensure robust path planning in such scenarios, the FL-P4AV framework integrates a stochastic model for dynamic obstacle movement and enables real-time path re-evaluation. At each timestep t , the environment grid G_t is derived from the previous state G_{t-1} through a

probabilistic update mechanism. Let (x, y) be the location of a static obstacle at $t - 1$. With a probability $p_{\text{move}} \in [0, 1]$, the obstacle may shift to a neighboring cell, selected uniformly from the valid Moore neighborhood (i.e., 8-connectivity). The grid update is defined in Equation (20):

$$G_t(x', y') = \begin{cases} 1, & \text{if obstacle moved from } (x, y) \text{ to } (x', y') \\ 0, & \text{if } (x, y) \text{ was previously occupied} \\ G_{t-1}(x', y'), & \text{otherwise} \end{cases}$$

The transition function can be compactly written as:

$$G_t = f_{\text{move}}(G_{t-1}, p_{\text{move}})$$

The function f_{move} uses uniform sampling among available adjacent free cells to determine the new position of the moving obstacle, ensuring obstacle collisions and overlaps are avoided. When any obstacle enters the predicted path P of a vehicle, a re-planning trigger is initiated. The vehicle then invokes the Semantic-Aware A* Path Planner again using the updated grid G_t . This mechanism allows for adaptive path correction in response to environmental dynamics, thereby increasing the reliability and safety of the navigation system. To reduce computational overhead, only a limited number of obstacles are allowed to move during each timestep, and vehicles operate asynchronously to maintain independent reactions to dynamic changes.

3.2.7 Path Quality Analyzer

The Path Quality Analyzer module is responsible for evaluating the performance and efficiency of the path generated by the Semantic-Aware A* planner. This module quantifies the path's effectiveness using multiple metrics that capture both geometric and dynamic characteristics of the planned route. The evaluation assesses the optimality of the planned path concerning length, smoothness, computational cost, and navigational complexity. The planned path is denoted as a sequence of k nodes, $P = \{n_1, n_2, \dots, n_k\}$ where n_1 is the start node and n_k is the destination. The path length L , defined as the number of steps taken to reach the destination, is computed using:

$$L = k - 1$$

The total Euclidean distance D between all consecutive nodes on the path is given by:

$$D = \sum_{i=1}^{k-1} \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2}$$

To evaluate the complexity of the path in terms of direction changes, the number of inflection points I is

calculated by counting the number of times the direction vector changes between three consecutive points:

$$I = \sum_{i=2}^{k-1} \delta(\vec{v}_{i-1}, \vec{v}_i), \quad \delta = \begin{cases} 1, & \text{if } v_{i-1} \neq \vec{v}_i \\ 0, & \text{otherwise} \end{cases}$$

where $\vec{v}_i = n_{i+1} - n_i$ is the direction vector at step i . The total turning angle Θ of the path is evaluated by summing the angular deviations at each inflection point:

$$\Theta = \sum_{i=2}^{k-1} \cos^{-1} \left(\frac{\vec{v}_{i-1} \cdot \vec{v}_i}{|\vec{v}_{i-1}| \cdot |\vec{v}_i| + \epsilon} \right)$$

where ϵ is a small value added to avoid division by zero. The number of explored nodes N_e during the A* search process is another critical metric, as it reflects the computational cost of the pathfinding algorithm:

$$N_e = \text{Cardinality of closed set}$$

These metrics collectively offer a thorough evaluation of the navigation path, weighing the trade-offs among path efficiency, smoothness, and processing overhead. This module functions for benchmarking the FL-P4AV system and for adaptive model tuning throughout the training process. The FL-P4AV framework offers a modular and decentralized approach to privacy-preserving and personalized path planning within collaborative autonomous vehicle networks. Each vehicle functions autonomously using a lightweight local model, which is improved via federated learning and safeguarded by differential privacy techniques. The system incorporates a semantic-aware A* search, allowing vehicles to calculate context-sensitive paths in dynamically changing environments. Probabilistic obstacle movement facilitates real-time re-planning, thereby enhancing adaptability. A path quality analyzer systematically assesses each route based on various performance metrics. This architecture enhances scalability, efficiency, and data security within multi-agent navigation systems.

3.3 Data Flow and Inter-Module Communication

The data flow in the FL-P4AV framework is designed to be modular, privacy-conscious, and cyclical, facilitating effective collaboration among autonomous vehicles while safeguarding sensitive information. Every vehicle initiates its process by engaging with the Environment Grid Initialization module to acquire a grid filled with obstacles, a source, and a destination. The car gathers information about its surroundings from what it sees and sends it to the Lightweight Local Model to be trained. The model that was taught guesses the semantic costs of grid cells. The Semantic-Aware A* Path Planner then uses these costs to make a starting route. There are different levels of privacy that are

used by the Privacy-Preserving Noise Injection tool to handle the local model parameters before they are shared. Then, these changes with noise are sent to the Federated Aggregation Engine, which uses the received parameters to figure out a global model. The new global model is sent back to all the cars that are taking part and mixed with their own local models.




Fig. 5. Federated Communication & Privacy Flow Diagram

While the vehicle traverses the grid, the Dynamic Obstacle Handling module refreshes the environment in real time, necessitating the re-invocation of the path planner when required. The Path Quality Analyzer assesses the resultant path, calculating performance metrics and transmitting the outputs to the Visualization Engine. This inter-module communication guarantees synchronized functionality across all system components and facilitates resilient, adaptive, and secure path planning in decentralized vehicular networks. Figure 5 depicts the comprehensive data flow and inter-module communication framework of the FL-P4AV system. The graphic illustrates the iterative learning process, starting with environmental sensing and culminating with federated model enhancement. Vehicles autonomously gather environmental characteristics and develop localized models, which are safeguarded through differential privacy prior to transmission to the central aggregator. The aggregator integrates these updates into a comprehensive model and disseminates them to vehicles, facilitating collaborative enhancement while maintaining the confidentiality of raw data. Arrows illustrate the two-way communication between local agents and the aggregator, whereas internal flows signify the interactions among the grid, planner, evaluator, and learning modules. The process emphasizes the close integration of privacy, learning, planning, and evaluation within the FL-P4AV ecosystem.

4. Results and Discussion

Extensive simulations were undertaken utilizing both synthetic settings and real-world datasets to assess the performance and scalability of the proposed FL-P4AV framework. Urban grid layouts were specifically obtained using the Waymo Open Dataset [31] and high-resolution

stereo images from the KITTI Dataset (2012/2015) [32], offering realistic environmental representations. Simulation experiments were carried out using different grid sizes: 20×20 , 30×30 , and 50×50 . The grids were filled with a corresponding number of static obstacles, varying from 50 in the 20×20 configuration to as many as 120 in the 50×50 arrangement. In each scenario, five autonomous vehicles were deployed, featuring randomized local starting positions and a shared destination point. The federated learning procedure was conducted across 10 communication iterations in each simulation. Local models, developed using semantically extracted data such as obstacle density, Euclidean distance to the objective, and grid-normalized coordinates, were consolidated to create a unified global model. These predictions were added to an A* algorithm that is aware of semantics to make paths that are sensitive to context and protect privacy. Comparative study of different grid sizes showed that FL-P4AV keeps planning efficiently and model convergence consistency even when the environment gets more complicated. This shows that it is very scalable and flexible for real-time navigation tasks.

Federated learning rounds are implemented using a secure, privacy-preserving aggregation algorithm. Each vehicle injects differential privacy (DP) noise into its model weights and biases before communication. Gaussian noise with variance σ^2 is applied to ensure the privacy of the local data while still enabling effective global model convergence. Metrics collected during the simulation include path length, total distance, planning time, number of nodes explored, number of inflection points, and total turning angle. These metrics are logged and visualized across multiple vehicles and rounds for comparative evaluation. A summary of simulation parameters is presented in Table 1.

Table 1. Simulation Configuration Parameters

Parameter	Value
Grid Size	$20 \times 20, 30 \times 30, 50 \times 50$
Number of Vehicles	5
Number of Obstacles	50
FL Rounds	10
Local Epochs per Round	5
Differential Privacy Noise	$\sigma = 0.1$
Feature Vector Dimension	4 (x, y, density, distance)
Path Metrics Evaluated	Length, Distance, Angle, Turns, Nodes, Time

All simulations were performed on a local machine using an Intel Core i7 (11th Gen, 2.8 GHz) CPU, 16 GB of RAM, and the Ubuntu 22.04 LTS operating system. The complete FL-P4AV framework was developed in Python 3.10 utilizing modules including NumPy for numerical analysis, pandas for data management, and matplotlib for result presentation. The

federated learning, privacy injection, and A* semantic path planner modules were developed from scratch to retain fine-grained control over data flow and parameter tuning. Each experimental run involved multiple vehicles, repeated over 10 federated learning rounds, and used custom scripts to monitor convergence, privacy behavior, and metric evolution under dynamic grid conditions. The summary of the hardware and software setup is provided in Table 2.

Table 2. Hardware and Software Specifications

Component	Specification
Operating System	Ubuntu 22.04 LTS
Processor	Intel Core i7, 11th Gen, 2.8 GHz
RAM	16 GB
Programming Language	Python 3.10
Libraries Used	NumPy, pandas, matplotlib
Simulation Rounds	10
FL/DP Framework	Custom Implementation (Python-based)

4.1 Environment Visualization and Dynamic Replanning

Figure 6 presents the full sequence of environment states and corresponding path planning behaviors under the FL-P4AV framework, visualized over seven key simulation frames. This comprehensive illustration demonstrates how the autonomous vehicles intelligently adapt their paths in response to changing environmental conditions while preserving privacy and decentralization.

Subfigure (a) depicts the initial static environment setup using a 20×20 grid. Static obstacles (black), source (red), and destination (green) are placed, forming the foundational scenario. Subfigure (b) shows the planned paths of five autonomous vehicles (V0–V4) at timestep $t = 0$, generated using the semantic-aware A* algorithm guided by local models refined via federated learning. Subfigures (c) to (g) correspond to timesteps $t = 1$ to $t = 5$, where a percentage of obstacles are moved dynamically. In each step, vehicles perceive the environmental change, and affected agents recompute their navigation paths using updated local semantic cost maps. This sequential visualization demonstrates the system's ability to maintain real-time awareness and reactivity. Vehicles independently identify obstructed paths and autonomously re-plan without needing centralized control or raw data sharing, which is particularly critical in privacy-sensitive and distributed applications. The proposed FL-P4AV framework outperforms traditional static or centralized path planning methods in several ways. First, it ensures continual adaptability as vehicles re-plan efficiently in response to dynamic obstacles, unlike conventional A* which would either fail or repeat full re-computation. Second, the semantic-aware cost functions,

learned locally and improved collaboratively via federated learning, enable vehicles to anticipate and avoid high-risk zones even before direct obstruction, demonstrating predictive intelligence. Finally, privacy is preserved without degrading system performance, and the simulation sequence (Figure 6 a–g) provides visual proof of real-time resilience, improved path continuity, and intelligent behavior at both local and global levels.

4.2 Federated Model Convergence

The convergence behavior of the FL-P4AV framework is illustrated through the final aggregated global model parameters computed after ten rounds of federated learning. Each vehicle trains a lightweight linear model and contributes its differentially private parameters to a central aggregator. The aggregator subsequently calculates the average of all incoming weights and biases to generate a cohesive global model, ensuring that no raw data is sent during the process. The global model parameters are illustrated in Figure 7 and further upon in Table 3. Four weights are associated with input features: normalized X and Y positions, obstacle density, and distance to the goal, along with a bias term. The highest magnitude weight is significantly linked to obstacle density, affirming its essential role in influencing navigation cost. The convergence of these values indicates a unified understanding among agents and enhances the model's applicability in distributed contexts.




Fig. 7. Final Aggregated Global Model Parameters


Table 3. Final Global Model Parameters

Parameter	Value
w0	0.108244
w1	0.067984
w2	0.823011
w3	0.493697
Bias	0.460564


4.3 Quantitative Evaluation of Planned Paths

A series of quantitative metrics were collected and visualized for each autonomous vehicle to validate the robustness and adaptability of the FL-P4AV framework. Factors include path length, planning time, number of inflection points, and total turning angle. Each metric is displayed using a unique chart type to improve interpretability and highlight system performance from various analytical viewpoints. Figure 8 uses a line plot to show the planned path length per vehicle. The near-uniformity of these lines suggests consistent performance across all vehicles, indicating that personalized local models are still harmonized effectively through federated learning. Figure 9, a scatter plot of planning time, confirms that all vehicles achieved extremely low computational latency (under 0.003s), even in dynamic conditions, demonstrating the efficiency of the semantic-aware A* planner.


Figure 10 presents a horizontal bar chart that illustrates the quantity of inflection points, reflecting the number of directional changes undertaken by each vehicle. The regulated range of inflections suggests more streamlined trajectories with reduced complexity. Figure 11, presented as an area plot, illustrates the total turning angle for each path. Lower angular values indicate fewer sudden bends, which improves navigational smoothness and energy efficiency—both important markers of optimal autonomous driving performance. The visual results substantiate that the proposed FL-P4AV framework generates stable, low-latency, smooth, and adaptive paths, outperforming traditional static or centralized planners. Through federated learning and semantic modeling, vehicles achieve a balance between personalized intelligence and collaborative optimization. As presented in Table 4, each vehicle achieves comparable path quality metrics, with minimal variation in path length, turning complexity, and planning time. These consistent results across distributed agents confirm the stability and collaborative intelligence of the FL-P4AV system.




a) Initial Environment Setup with Static Obstacles and Goal Locations.




b) Planned Paths of Vehicles (V0-V4) at Timestep $t = 0$




c) Dynamic Environment Update and Path Replanning at $t = 1$




d) Dynamic Environment Update and Path Replanning at $t = 2$



e) Dynamic Environment Update and Path Replanning at $t = 3$



f) Dynamic Environment Update and Path Replanning at $t = 4$



g) Dynamic Environment Update and Path Replanning at $t = 5$

Fig. 6. Sequence of Environment States and Path Replanning in FL-P4AV Across Dynamic Timesteps




Fig. 8. Path Length Analysis across Vehicles




Fig. 9. Planning Time Efficiency among Vehicles




Fig. 10. Comparative Turning Complexity of Planned Paths




Fig. 11. Cumulative Turning Behavior in Vehicle Navigation

Table 4. Path Metrics Summary

Vehicle	Path Length	Total Distance	Infection Points	Turning Angle (°)	Planning Time (s)	Nodes Explored
0	16.0	20.97	5.0	225.00	0.0022	83.0
1	16.0	20.14	11.0	495.00	0.0024	112.0
2	17.0	20.72	12.0	540.00	0.0028	203.0
3	17.0	21.14	10.0	495.00	0.0026	190.0
4	15.0	19.142	7.0	315.00	0.0021	90.0

4.4 Differential Privacy and Aggregation Behavior

To preserve individual data confidentiality within the federated learning cycle, the FL-P4AV framework integrates differential privacy (DP) by injecting noise into each vehicle's model parameters before sharing. This technique ensures that model updates are statistically indistinguishable, thereby preventing adversarial inference about local datasets. Figure 12 illustrates the magnitude of noise injected across four model weights for five participating vehicles. The observed variations are the result of independently applied Gaussian noise $\sigma = 0.1$, affirming that each update maintains a privacy-preserving boundary without affecting learning convergence significantly.




Fig. 12. Differential Privacy Noise Injected per Vehicle




Fig. 13. Aggregated Model Weights Across Federated Rounds

Following noise injection, these anonymized updates are aggregated at the federated coordinator node to produce global model parameters. Figure 13 presents the convergence behavior of four major weights across ten rounds of federated learning. Each curve shows a smooth trajectory toward the final learned values, indicating stable and collaborative model refinement across rounds. The observed stability, even in the presence of privacy-related disturbances, validates the effectiveness of the federated learning strategy implemented in FL-P4AV. This also confirms the efficacy of lightweight local models in learning context-aware navigation cost estimations and enhancing global intelligence. The results highlight the framework's capacity to maintain a balance between security and utility: differential privacy mechanisms preserve model accuracy, and the aggregation process reliably yields generalizable weights appropriate for distributed autonomous navigation.

4.5 Comparative Analysis

To evaluate the effectiveness of the proposed FL-P4AV framework, we conducted a comparative analysis against two standard approaches: Static A* path planning and Federated Learning without Differential Privacy (FL-NoDP) [30]. All models were tested under identical simulation conditions

using the same grid, vehicle count, and environmental settings. Table 5 summarizes the results across four performance indicators: path length, number of inflection points, total turning angle, and planning time. As illustrated in Figure 14, FL-P4AV consistently delivers superior results across all metrics. It generates shorter paths with fewer directional changes and lower angular deviation, indicating smoother and more efficient navigation. Additionally, FL-P4AV maintains the lowest planning time, highlighting its real-time applicability.

Table 5. Comparative Evaluation of Path Planning Methods (Averaged over 5 Vehicles)

Method	Avg. Path Length	Inflection Points	Turning Angle (°)	Planning Time (s)
Static A*	25.2	7.8	144.3	0.0091
FL-NoDP	21.4	5.2	113.8	0.0027
FL-P4AV	20.8	4.6	108.5	0.0025




Fig. 14. Comparative Analysis of FL-P4AV vs Baseline Methods

Notably, even with privacy-preserving noise added to model updates, FL-P4AV outperforms the FL-NoDP model [30], proving that the integration of differential privacy does not compromise system performance. In contrast, Static A* lacks semantic understanding [13], resulting in suboptimal paths with increased complexity and computational overhead. These findings confirm that FL-P4AV not only supports decentralized and privacy-aware learning but also achieves practical gains in navigational accuracy, efficiency, and adaptability in dynamic environments.

5. Discussions

The proposed FL-P4AV framework delivers a decentralized and privacy-preserving approach to collaborative path planning for autonomous ground vehicles. It combines federated learning, semantic feature modeling, and differential privacy to allow intelligent and secure

navigation in dynamic environments. Unlike traditional centralized or static planners, FL-P4AV enables each vehicle to learn from localized experience while benefiting from global intelligence. From a computational perspective, the system introduces acceptable complexity given its scalability and accuracy. The semantic-aware A* path planner used in each vehicle operates with a time complexity of $\mathcal{O}(E + V \log V)$, where V is the number of grid cells and E is the number of edges between valid cells. The training of the lightweight linear model incurs a cost of $\mathcal{O}(n \cdot d)$ per local epoch, where n is the number of training samples and d is the feature dimension (in our case, $d = 4$). Federated aggregation requires $\mathcal{O}(V \cdot d)$ per round for averaging model parameters across V vehicles. The differential privacy step, which adds Gaussian noise, has a constant-time complexity $\mathcal{O}(d)$ but introduces trade-offs in convergence accuracy depending on the noise scale σ . As supported by Figures 7–12 and Table 5, FL-P4AV maintains strong accuracy, path smoothness, and low planning latency, even with injected DP noise. The modularity of the architecture allows easy scalability to larger grid sizes or more vehicles, although increased communication cost in federated rounds may necessitate optimization via gradient compression or asynchronous updates. The real-time simulations from $t = 0$ to $t = 5$ validate the system’s responsiveness to dynamic obstacle changes. FL-P4AV’s model convergence behavior (Figure 13) and low per-vehicle noise injection levels (Figure 12) confirm its readiness for deployment in multi-agent environments requiring both autonomy and data protection.

6. Conclusion

This study introduces FL-P4AV, an innovative framework that utilizes federated learning to facilitate privacy-preserving and personalized path planning for autonomous ground vehicles. FL-P4AV integrates lightweight local intelligence, semantic-aware A* pathfinding, and differential privacy to facilitate secure and adaptive navigation while preserving the confidentiality of sensitive raw data. The framework underwent validation via comprehensive simulations in dynamic environments, demonstrating that each vehicle effectively adapted to real-time obstacle alterations while ensuring minimal planning latency and smooth trajectories. Comparative results indicate that FL-P4AV surpasses traditional A* and non-private federated systems in terms of path quality and computational efficiency. Furthermore, the use of differential privacy attained robust data security with little performance compromises. FL-P4AV provides a scalable and pragmatic approach for collaborative multi-agent navigation, serving as a viable basis for future applications in smart city logistics, intelligent transportation, and autonomous fleet coordination.

Funding: The authors did not receive support from any organization for the submitted work.

Declarations

Conflict of interest The authors declare no conflict of interest

References

- [1] P. Bhattacharya et al., "Coalition of 6G and Blockchain in AR/VR Space: Challenges and Future Directions," *IEEE Access*, vol. 9, pp. 169709–169735, 2021 [CrossRef] [Google Scholar] [Publisher Link]
- [2] D. Wasif, T. J. Moore, and J.-H. Cho, "RESFL: An Uncertainty-Aware Framework for Responsible Federated Learning by Balancing Privacy, Fairness and Utility in Autonomous Vehicles," *arXiv preprint arXiv:2503.16251*, 2025 [CrossRef] [Google Scholar] [Publisher Link]
- [3] Y. Li, X. Tao, X. Zhang, J. Liu, J. Xu, "Privacy-Preserved Federated Learning for Autonomous Driving," in *IEEE Transactions on Intelligent Transportation Systems*, 2021, pp. 8423–8434 [CrossRef] [Google Scholar] [Publisher Link]
- [4] Y. Chen, Z. Qin, and W. Shi, "Drone Traffic Monitor Aided Federated Learning for Autonomous Road Vehicles," *Frontiers in Communications and Networks*, vol. 2, p. 58, 2020. [CrossRef] [Google Scholar] [Publisher Link]
- [5] Y. Liu et al., "Privacy-Preserving Real-Time Action Detection in Intelligent Vehicles Using Federated Learning-Based Temporal Recurrent Network," *Sensors*, vol. 24, no. 14, p. 2820, 2024. [CrossRef] [Google Scholar] [Publisher Link]
- [6] Abenezer Golda et al., "Privacy and Security Concerns in Generative AI: A Comprehensive Survey," *IEEE Access*, vol. 12, pp. 14089–14115, 2024 [CrossRef] [Google Scholar] [Publisher Link]
- [7] Reza Shahbazian et al., "Integrating Machine Learning into Vehicle Routing Problem: Methods and Applications," *IEEE Access*, vol. 12, pp. 13883–13895, 2024 [CrossRef] [Google Scholar] [Publisher Link].
- [8] Y. Tian, Y., Shi, Y., Zhang, Y., & Tian, Q, "Personalized Federated Learning with Correlated Differential Privacy for Autonomous Driving," *Sensors*, vol. 25, no. 1, p. 178, 2025 [CrossRef] [Google Scholar] [Publisher Link].
- [9] Alqubaysi T, Asmari AFA, Alanazi F, Almutairi A, Armghan A. Federated Learning-Based Predictive Traffic Management Using a Contained Privacy-Preserving Scheme for Autonomous Vehicles. *Sensors*. 2025 [CrossRef] [Google Scholar] [Publisher Link]
- [10] Jakob Weber, Markus Gurtner, Benedikt Alt, Adrian Trachte, Andreas Kugi, "Privacy-Preserving Federated Learning for Feedforward Control in Multi-Agent Autonomous Systems," *arXiv preprint arXiv:2503.02693*, 2025 [CrossRef] [Google Scholar] [Publisher Link].
- [11] Robert Akinie, Nana Kankam Gyimah, Mansi Bhavsar, John Kelly, "Fine-Tuning Federated Learning-Based Intrusion Detection Systems for Transportation IoT," *arXiv preprint arXiv:2502.06099*, 2025 [CrossRef] [Google Scholar] [Publisher Link].
- [12] Wang, B., Han, Y., Wang, S., Tian, D., Cai, M., Liu, M., & Wang, L. (2022). A Review of Intelligent Connected Vehicle Cooperative Driving Development. *Mathematics*, 10(19), 3635, 2022 [CrossRef] [Google Scholar] [Publisher Link].
- [13] Gosselin, R., Vieu, L., Loukil, F., & Benoit, A. (2022). Privacy and Security in Federated Learning: A Survey. *Applied Sciences*, 12(19), 990, 2022 [CrossRef] [Google Scholar] [Publisher Link].
- [14] Y. Ming et al., "A Survey of Path Planning Algorithms for Autonomous Vehicles," *SAE Int. J. Commer. Veh.*, vol. 14, no. 1, pp. 97–109, 2021 [CrossRef] [Google Scholar] [Publisher Link].
- [15] M. Rahmati, "Federated learning for privacy-preserving AI in human–robot collaboration for smart manufacturing," *J. Intell. Manuf. Spec. Equip.*, 2025 [CrossRef] [Google Scholar] [Publisher Link].
- [16] Mahmud, S. A., Islam, N., Islam, Z., Rahman, Z., & Mehedi, S. T. (2024). Privacy-Preserving Federated Learning-Based Intrusion Detection Technique for Cyber-Physical Systems. *Mathematics*, 12(20), 3194, 2024 [CrossRef] [Google Scholar] [Publisher Link].
- [17] Yanan Ma, Senkang Hu, Zhengru Fang, Yun Ji, Yiqin Deng, Yuguang Fang, "Sense4FL: Vehicular Crowdsensing Enhanced Federated Learning for Autonomous Driving," *arXiv preprint arXiv:2503.17697*, 2025 [CrossRef] [Google Scholar] [Publisher Link].
- [18] Muzi Peng, Jiangwei Wang, Dongjin Song, Fei Miao, Lili Su., "Privacy-Preserving and Uncertainty-Aware Federated Trajectory Prediction for Connected Autonomous Vehicles," *arXiv preprint arXiv: 2303.04340*, 2023 [CrossRef] [Google Scholar] [Publisher Link].
- [19] M. D. Tezerjani, D. Carrillo, D. Qu, S. Dhakal, A. Mirzaeinia, and Q. Yang, "Real-time Motion Planning for autonomous vehicles in dynamic environments," *arXiv:2406.02916 [cs.RO]*, 2024 [CrossRef] [Google Scholar] [Publisher Link].
- [20] Qin, H., Shao, S., Wang, T., Yu, X., Jiang, Y., & Cao, Z, Review of Autonomous Path Planning Algorithms for Mobile Robots. *Drones*, 7(3), 211, 2023 [CrossRef] [Google Scholar] [Publisher Link].
- [21] V. Bulut, "Path planning for autonomous ground vehicles based on quintic trigonometric Bézier curve," *J. Braz. Soc. Mech. Sci. Eng.*, vol. 43, no. 2, pp. 1–14, 2021 [CrossRef] [Google Scholar] [Publisher Link].
- [22] S. Ragothaman, M. Maaref, and Z. M. Kassas, "Autonomous Ground Vehicle Path Planning in Urban Environments Using GNSS and Cellular Signals Reliability Maps: Simulation and Experimental Results," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 57, no. 4, pp. 2575–2586, 2021 [CrossRef] [Google Scholar] [Publisher Link].
- [23] Yin, J., Hu, Z., Mourelatos, Z., Gorsich, D., Singh, A., & Tau, S. (2023). Efficient Reliability-Based Path Planning of Off-Road Autonomous Ground Vehicles Through the Coupling of Surrogate Modeling and RRT*. 24, 15035–15050, 2023 [CrossRef] [Google Scholar] [Publisher Link].
- [24] Global Off-Road Path Planning of Unmanned Ground Vehicles Based on the Raw Remote Sensing Map. (2023). *SAE Technical Paper Series*, 2023, [CrossRef] [Google Scholar] [Publisher Link]

- [25] Han, Z., Chen, P., Zhou, B., & Yu, G. (2024). Dual-Layer Path Planning for Unmanned Ground Vehicles Based on Probabilistic Roadmap and Proximal Policy Optimization. 1–6, 2024, [CrossRef] [Google Scholar] [Publisher Link]
- [26] J. Yin, Z. Hu, Z. P. Mourelatos, D. Gorsich, A. Singh, and S. Tau, "An Efficient Surrogate Modeling Method for Reliability-Based Global Path Planning of Off-Road Autonomous Ground Vehicles," in Proc. ASME 2023 Int. Des. Eng. Tech. Conf. Comput. Inf. Eng. Conf., vol. 3B: 49th Des. Autom. Conf. (DAC), Boston, MA, USA, Aug. 20–23, 2023 [CrossRef] [Google Scholar] [Publisher Link].
- [27] F. B. Romero, I. E. Soriano, R. Marotta, and L. De Matteis, "Optimizing Motion Planning for Autonomous Vehicles: Path Planning, Curvature Analysis, and Steering Control," in 2024 Int. Symp. Power Electron., Elect. Drives, Autom. Motion (SPEEDAM), Napoli, Italy, 2024, pp. 569-573 [CrossRef] [Google Scholar] [Publisher Link].
- [28] Chen, R., Hu, J., & Xu, W. (2022). An RRT-Dijkstra-Based Path Planning Strategy for Autonomous Vehicles. Applied Sciences, 12(23), 2022 [CrossRef] [Google Scholar] [Publisher Link].
- [29] Zhang, H., Liu, K., Gao, F., & Zhao, F. (2023). Research on Path Planning and Path Tracking Control of Autonomous Vehicles Based on Improved APF and SMC.2023, [CrossRef] [Google Scholar] [Publisher Link].
- [30] Chen Z, Zheng H, Liu G. AWDP-FL: An Adaptive Differential Privacy Federated Learning Framework. Electronics. 2024; 13(19):3959, 2024 [CrossRef] [Google Scholar] [Publisher Link].
- [31] Waymo Open Dataset. [Online]. Available: <https://waymo.com/open/>
- [32] KITTI Vision Benchmark Suite – Stereo and Flow Datasets (2012/2015). [Online]. Available: <https://www.kaggle.com/datasets/klemenko/kitti-dataset>