



Integration of Speed Control and Parking Management in SUMO for Urban Traffic Optimization

Cristian Iza^a, Joana Iza^b, Juan-Luis Posadas-Yagüe^c, and José-Luis Poza-Lujan^c

^a Doctoral School, Universitat Politècnica de València, Valencia 46022, Spain

^b Escuela Superior Politécnica de Chimborazo, Riobamba 060104, Ecuador

^c Research Institute of Industrial Computing and Automatics, Universitat Politècnica de València, Camino de Vera, s/n 46022, Valencia, Spain

✉ ciza@alumni.upv.es, joana.iza@epoch.edu.ec, jposadas@disca.upv.es, jopolu@upv.es

KEYWORDS

speed control;
parking management;
lane activation;
control via networks;
performance indicators;
Simulation in SUMO;
adaptive control

ABSTRACT

This study presents an intelligent system for dynamic urban traffic management, focused on the automatic adjustment of speed limits and the activation of lanes as parking areas to optimize urban traffic. Using realistic scenarios simulated in the SUMO (Simulation of Urban MObility) platform and controlled by a neural network implemented in Python, the system effectively responds to variable conditions such as vehicle density, traffic flow, and other factors. To evaluate the achieved optimization, the results are compared using various traffic performance indicators, including congestion levels, average speed, travel time, among others. Through adaptive real-time behavior, the system achieves greater traffic fluidity, reduced congestion, and better use of available infrastructure. The designed control system significantly improves performance compared to uncontrolled scenarios: on average, it reduces travel time by 16 %, CO₂ emissions by 13 %, waiting time by 12 %, and accident probability by 15 %, while traffic flow increases by 11 %. The results show that this approach can effectively complement traditional traffic light control and adapt to diverse urban contexts, positioning itself as a promising solution to improve mobility in modern cities.



1. Introduction

The rapid growth of the vehicle fleet and urban expansion have posed new challenges for traffic management in cities. Traffic congestion, lost travel time, and inefficient use of road space are recurring problems in many metropolitan areas. In response, various strategies aimed at improving mobility have emerged, among which intelligent transportation systems (ITS) stand out. These systems integrate technologies such as sensors, communication networks, and artificial intelligence algorithms (Ehlers, 2019).

In this context, the use of artificial neural networks (ANNs) has gained relevance as a tool for real-time decision-making. These networks can learn complex patterns from large volumes of data, allowing them to predict traffic behaviors and propose automatic adjustments to road conditions (Rahman et al., 2020). This approach aims not only to reduce congestion and improve traffic flow, but also to optimize the use of existing infrastructure, avoiding major physical investments. Furthermore, the integration of these intelligent systems enables an adaptive and continuous response to changes in vehicle demand, thus contributing to the improvement of the urban transportation system (Bagheri and Sarvi, 2023). Speed control and parking management are fundamental components in urban traffic planning. Both aspects influence vehicle behavior, the configuration of road infrastructure, and the distribution of traffic flow within the city. Studying them helps to understand how mobility related decisions impact the urban environment. This study implements a neural network developed in Python, applied to various scenarios simulated in SUMO, for speed and parking control based on urban traffic variables.

The manuscript is structured into five sections: the introduction contextualizes the problem; the background section presents the theoretical foundations of intelligent transportation systems and neural networks; the experimental design describes the design process; the results are critically presented and analyzed; and finally, the study's conclusions are outlined.

2. Related Work

To address this topic, the following related works are presented based on this article. In 2025, SUMO-MCP, a platform that integrates traffic simulation tools with SUMO, enabling the generation of traffic scenarios, demand creation, and comparative analysis. Although its focus is autonomous traffic simulation, the platform also includes complementary features such as the definition of parking areas, which can be useful in certain urban management cases (Ye et al., 2025).

In 2024, an intelligent parking system based on machine learning was developed, using learning automata and an end-to-end approach for managing spaces in off street parking lots. To validate the model and analyze its performance, the study conducted the experimental part using the Arena simulation software, with a maximum parking time under simulated conditions not exceeding 4 hours (Chen, 2024). In 2024, an adaptive highway speed control system based on deep reinforcement learning was proposed, using real-time data to dynamically adjust speed limits. The system was designed to operate in road environments with high variability in traffic demand, considering variables such as vehicle density, average speed, and environmental conditions to continuously adapt the model's decisions (Sun and Yu, 2024).

In 2023, an intelligent system based on temporal difference reinforcement learning was proposed for autonomous vehicles that optimize the average speed in an urban network simulated with SUMO.

The system uses real time data to adjust its parameters and responds to variations in traffic flow during the simulation (Reza et al., 2023). In 2023, a traffic light control system based on artificial neural networks was developed, using data from adaptive traffic light cycles generated by SUMO to train the network. The study focused on minimizing vehicle delay during peak hours in the city of Volgograd (Samarakoon and Sivakumar, 2022).

In 2022, it was proposed to develop a model based on infrastructural agents and reinforcement learning (RL) for the dynamic management of urban spaces, including areas designated for parking, where curbs act as agents regulating the use of parking spaces. This hybrid approach combines partial observations with full data to support decision-making in simulated urban environments (Yin et al., 2022). Also in 2021, traffic management and control measures on highways, such as variable speed control and route guidance, were evaluated using SUMO. The study showed that these strategies can improve traffic flow and reduce congestion under high demand conditions (Ye et al., 2021). In 2020, Attend Light was introduced, a universal model based on attention mechanisms and deep reinforcement learning for traffic signal control. This system uses two attention mechanisms to adapt to different intersection configurations and traffic light phases, achieving significant improvements in traffic efficiency compared to traditional methods (Farrag et al., 2020).

Although the reviewed works are relevant and closely related to the present study, it is important to note that none of them specifically address the combination of speed control and dynamic management of parking lanes within the same system using neural networks in simulated urban environments. Most focus primarily on parking control or isolated speed management. In this regard, our proposal aims to fill a gap in the literature by integrating these two aspects of traffic control into a single system, through an adaptive approach based on machine learning. The reviewed methodologies serve as a starting point and source of inspiration, providing a reference framework that supports and contextualizes our contribution to the field of intelligent urbanity.

3. Experimental Design

To evaluate the effectiveness of intelligent traffic regulation, several simulation scenarios were designed using the SUMO (Simulation of Urban Mobility) environment. In these scenarios, dynamic adjustments were made to speed limits and parking availability, aiming to improve traffic flow efficiency and reduce congestion.

The control of parameters and adjustments was implemented using Python, allowing dynamic and real time management of traffic conditions. This approach also facilitated the evaluation of different regulation strategies under various traffic conditions, providing a comprehensive insight into the optimization of the urban transportation system. The process developed for the implementation of the proposed system is detailed below.

3.1. Traffic Indicators

Traffic indicators are fundamental metrics that allow for the quantification of road system behavior and the objective evaluation of the impact of implemented strategies or control systems. These indicators facilitate the identification of bottlenecks, the measurement of operational efficiency, and the analysis of performance under different traffic conditions. Additionally, they provide a solid foundation for informed decision making, the planning of mobility policies, and the continuous improvement

of infrastructure. When used appropriately, they enable the optimization of the urban transportation system. The indicators implemented in our system are as follows:

3.1.1. Queue Length (QL)

It is the total physical distance typically occupied (in meters or number of vehicles) by stopped or very slowly moving vehicles waiting to pass a specific point, commonly a traffic light, a congested intersection, a construction zone, or a bottleneck (Sabour et al., 2021). Specifically, it is measured in meters (m) and is calculated using (1):

$$QL = N \times L_v \quad (1)$$

Where N represents the number of stopped vehicles, and L_v is the average vehicle length.

3.1.2. Waiting Time at Traffic Lights (WT)

It is the average duration that a vehicle remains stopped while waiting for the green light at a signalized intersection (Skorobogatchenko et al., 2023). It is measured in seconds (s), and it is calculated using (2). This indicator reflects traffic congestion at intersections and helps assess the efficiency of traffic signal timing.

$$WT = \sum (t_s - t_e) / N \quad (2)$$

Where t_s is the time at which a vehicle stops, t_e is the time at which the vehicle resumes movement, and N represents the total number of vehicles.

3.1.3. CO₂ Emissions (ECO₂)

This refers to the amount of carbon dioxide (and other greenhouse gases or local pollutants such as NO_x, PM_{2.5}) released into the atmosphere by vehicles. CO₂ is the main greenhouse gas associated with transportation (Kheterpal et al., 2018). It is measured in grams of CO₂ per kilometer (gCO₂/km), and it is calculated using (3):

$$ECO_2 = \sum e_i / T \quad (3)$$

Where e_i represents the CO₂ emission of vehicle i , and T corresponds to the total observation time. This allows for evaluating the total emissions produced during the traffic observation period.

3.1.4. Traffic Flow (TF)

It refers to the total number of vehicles successfully passing through a specific point or section of a road during a given time interval (Guo et al., 2022). Specifically, it is measured in vehicles per hour (veh/h), and it is calculated using (4):

$$TF = N / T \quad (4)$$

Where N is the number of vehicles passing a point, and T is the time interval observed, providing a measure of traffic volume at that location.

3.1.5. Average Travel Time (ATT)

It is the average time that vehicles take to travel along a specific route or a segment of the road network (Hussain et al., 2020). It is measured in seconds (s), and it is calculated using (5):

$$ATT = \Sigma (t_{end} - t_{start}) / N \quad (5)$$

Where t_{end} is the time the vehicle reaches its destination, t_{start} is the trip start time, and N represents the number of vehicles completing the trip.

3.1.6. Accident Damage Probability (ADP)

This refers to the likelihood that, once an accident occurs, it results in injuries (minor or severe) or fatalities. It is directly related to the kinetic energy involved (speed) and the type of conflict. It can also be interpreted as the average severity of accidents. The accident damage probability is measured as a dimensionless value between 0 and 1, or expressed as a percentage (%) (Mousavi et al., 2017), and it is calculated using (6):

$$ADP = N_a / N_t \quad (6)$$

Where N_a represents the number of risk or conflict events, and N_t is the total number of vehicles observed or simulated.

3.2. Management Strategies

These strategies corresponded to the interventions or policies implemented in the system, whose impact was evaluated in relation to different traffic indicators. These strategies were specifically designed to modify the operational conditions of the simulated environment, thereby allowing a deep and detailed analysis of their effect on the overall behavior of the transportation system under different scenarios.

This facilitated the identification of significant changes in the system's performance and dynamics, providing valuable information for decision-making and the optimization of traffic management. Moreover, by systematically applying these strategies, the study was able to quantify improvements in traffic flow, safety, and efficiency, highlighting the practical relevance of the proposed control measures.

3.2.1. Speed Management

Speed management refers to the implementation of measures aimed at regulating or influencing vehicle speeds. These measures may include fixed or variable speed limits, control systems such as cameras or radars, traffic calming elements like speed bumps or chicanes, and traffic light coordination to create «green waves» that promote constant speeds. The overall objective of these actions is to improve road safety, facilitate smoother traffic flow, reduce speed variability, and, in some cases, optimize the capacity of road infrastructure (Beil et al., 2022).

3.2.2. Parking Management

These are strategies related to the provision, regulation, and control of vehicle parking, both on-street and off-street. They include aspects such as the availability and location of parking spaces,

time limits for parking, and the monitoring and enforcement of illegal parking, such as double parking. The overall objective of these actions is to reduce the time drivers spend searching for parking which generates additional traffic prevent obstructions to traffic flow caused by improperly parked vehicles and efficiently manage the use of road space (Amiour et al., 2022). Proper implementation of these strategies can significantly improve overall traffic efficiency and urban mobility.

3.3. Correlation between Indicators and Managements

A 5-point Likert scale was used to evaluate the expected impact on key traffic performance indicators following the effective implementation of two strategies: speed management and parking management (Beil et al., 2022). The use of a Likert scale table allows for a structured evaluation of the estimated impact of traffic management strategies on key performance indicators. It enables the quantification of expert judgment in a clear and analyzable format. The scale was defined as follows: 0: Very Favorable (Clear improvement of the indicator).

- 1. Favorable (Moderate improvement)
- 2. Neutral (No significant impact)
- 3. Unfavorable (Slight deterioration)
- 4. Very Unfavorable (Considerable deterioration of the indicator)

Table 1 presents the estimated impact on each indicator resulting from the successful implementation of each type of strategy.

Table 1. Comparison of Traffic Indicators between Speed Management and Parking Management

Traffic Indicator	Speed Management	Parking Management
Queue Length	1	1
Waiting Time at Traffic Lights	1	2
CO ₂ Emissions	1	1
Traffic Flow (Optimization)	1	1
Average Travel Time	1	1
Accident Damage Probabilities	3	1

The table presents the expected impact of two traffic management strategies: speed management and parking management, evaluated using a Likert scale from 0 to 5, where low values indicate a positive and near optimal effect, and high values reflect a higher level of problems or negative impact. According to the results, both strategies are effective in reducing queue length, waiting time at traffic lights, CO₂ emissions, optimizing traffic flow, and decreasing average travel time, with values close to 1, indicating a significant improvement in these indicators.

Speed management particularly stands out in reducing queue length and waiting time by enabling traffic light synchronization and traffic calming, promoting smoother and more stable traffic flow. Meanwhile, parking management helps free up road space and prevents obstructions caused by

improperly parked vehicles, which also positively impacts congestion reduction and improves traffic flow efficiency (Sun and Yu, 2024).

However, a critical aspect that emerges from the table is the increased likelihood of accident damage associated with speed management, with a value of 3 on the Likert scale, reflecting a moderately high level of risk. This aligns with the reality that increasing speed can raise the kinetic energy involved in collisions, thereby increasing the severity of accidents. In contrast, parking management has a positive impact on road safety, with a rating of 1, as it avoids conflicts and dangerous parking situations that could lead to accidents.

These results suggest that, although speed management can improve multiple operational indicators, its implementation must be carefully balanced with complementary safety measures to mitigate accident risks. Together, both approaches contribute to an overall improvement of the traffic system, provided they are implemented in a coordinated and adaptive manner, considering the trade-offs between efficiency and safety (Azfar and Ke, 2024).

3.4. Input Variables

For the development of the simulated scenarios, several input variables were considered to characterize the traffic state and guide control decisions. Each is described below along with its importance in the traffic system. These variables capture key aspects of urban traffic. By integrating these inputs, the system can make adaptive decisions that respond to both local and network-wide conditions, allowing the reinforcement learning agent to learn effective control strategies that balance efficiency, safety, and operational.

3.4.1. Number of Vehicles

This variable allows estimating traffic density in real time. A higher number of vehicles may trigger reductions in speed limits to prevent congestion or limit access to parking areas to prioritize traffic flow. In contrast, with fewer vehicles, speed limits can be increased, and more parking spaces can be enabled.

3.4.2. Average Speed

The average traffic speed provides a direct measure of the level of vehicle flow. Low values may indicate congestion, which would justify further reducing speed limits to avoid dangerous maneuvers or restricting parking to free up road space. Conversely, a high average speed suggests smooth flow, allowing speed limits to be maintained or even increased and parking availability expanded.

3.4.3. Time of Day

Time directly influences traffic patterns, as there are peak periods (such as workday start and end hours) and times of low demand. During peak hours, more restrictive speed limits and reduced parking spaces may be applied to favor traffic flow. During off-peak hours, both measures can be relaxed.

3.4.4. Day of the Week

Traffic behavior varies between weekdays and weekends. On weekdays, the system can prioritize mobility by reducing parking availability and adjusting speed limits to better manage high demand. On weekends, greater parking availability and wider speed limits can be allowed due to lower vehicle pressure. It reflects traffic and parking differences between weekdays and weekends.

3.4.5. Traffic Light Timings

The durations assigned to each traffic light phase directly influence vehicle flow. Adjusting them based on the number of vehicles and average speed allows for improved synchronization between intersections and reduces unnecessary waiting. Under high demand conditions, timings can be extended for the most congested roads, while during low traffic periods, the cycle can be optimized to avoid delays.

In SUMO, traffic light phases are represented by character strings where each letter indicates the state of a group of lights: «G» means green, «y» yellow, and «r» red. For example, «GGrr» means the first group has a green light and the second group a red light; «rryy» indicates red for the first group and yellow for the second; «yyrr» shows yellow in the first group and red in the second; and «rrGG» means red for the first group and green for the second. These SUMO specific variables reflect the current state of the traffic light phases in the simulation and are key to controlling vehicle flow at intersections.

3.5. Output Variables

Two fundamental output variables are presented that enable dynamic management of urban traffic. These variables facilitate real time adaptation of key parameters to improve traffic flow, safety, and reduce congestion. Their analysis is essential to optimize the overall performance of the transportation network.

3.5.1. Adjusted Speed

The output variable «adjusted speed» represents the optimal speed that the system recommends for vehicles based on current traffic conditions and the implemented management strategies. This variable is fundamental to regulating vehicular flow, avoiding congestion and minimizing the risk of accidents. Properly adjusting speed allows for a more uniform and safe circulation, reducing both queue formation and pollutant emissions associated with sudden braking and acceleration. In systems like SUMO, this variable translates into the dynamic modification of speed limits on the edges of the simulated scenario, adapting to the real needs of the road environment.

3.5.2. Parking

The output variable «parking» indicates the availability or restriction of specific zones where vehicles can park within the system. Efficient parking management is key to reducing congestion caused by vehicles searching for parking space or stopping in inappropriate areas, which negatively impacts traffic flow. By controlling the use of space designated for parking, the system contributes to a better distribution of traffic and reduces waiting times and emissions. In simulated environments such as SUMO, this variable can activate or deactivate parking areas, facilitating the evaluation of the impact of different management policies on urban dynamics.

The selected output variables, vehicle speed and parking management, were chosen for their ability to directly influence key urban mobility indicators in a measurable and controllable way. Unlike other potential control variables, these allow the system to act dynamically and adaptively on both traffic flow efficiency and the spatial distribution of vehicles, providing a high impact on traffic performance while maintaining manageable operational complexity. This combination ensures that the reinforcement learning-based system can effectively and scalably optimize urban mobility.

In addition, the coordinated control of speed and parking helps balance vehicle distribution and reduce local congestion. By adjusting speed limits and managing parking availability according to

real-time demand, the system improves traffic flow, reduces unnecessary stops, and contributes to lower fuel consumption and emissions.

3.6. Model Used

The implemented model corresponds to an artificial neural network applied through deep reinforcement learning, specifically using the Proximal Policy Optimization (PPO) algorithm. This approach combines an actor critic architecture where one neural network makes action decisions (actor) while another estimates the state value (critic). PPO stands out for ensuring stable policy updates by restricting changes between iterations through a probability clipping mechanism, which improves training stability.

This technique allows the agent to learn optimal policies in complex and dynamic environments, such as vehicle traffic management in a simulated scenario (Sabour et al., 2021). The entire system was developed in the Python programming language, leveraging libraries such as TensorFlow and Stable Baselines3, which facilitated both the model implementation and its integration with the SUMO simulation platform.

3.6.1. Model Architecture

The neural network architecture used in your system is a fully connected network, also known as an MLP (Multi-Layer Perceptron) as shown in Figure 1.




Figure 1. Neural network architecture for controlling speed and parking in urban traffic

Composed of two hidden layers of fixed size. The network input consists of a vector of 8 normalized features representing the traffic state on a given edge, including: number of vehicles, average speed, traffic light encoding in one hot format (green, yellow, red), remaining duration of the traffic light phase, day of the week, and time of day. Each hidden layer has 64 neurons with ReLU activation functions, which allows the model to capture nonlinear relationships in the data.

The output is divided into three branches: one to predict the continuous action (e.g., speed adjustment), another to decide whether to enable a zone as parking, and a third output that predicts the value of the state required for PPO training (Alekszejenkó and Dobrowiecki, 2023). Altogether, this architecture enables the agent to learn a mixed (continuous and discrete) policy and a value function to optimize urban traffic management.

3.6.2. Input Normalization

The normalization of the inputs was a fundamental step to ensure efficient and stable convergence during the training of the neural network (Guastella et al., 2023). In this case, the variables representing the traffic state such as the number of vehicles, average speed, remaining duration of the traffic light phase, day of the week, and time of day were scaled so that their values fell within comparable and bounded ranges. For example, the number of vehicles was normalized by dividing by a presumed maximum value (50 vehicles), the average speed was scaled based on an estimated maximum speed (13.9 m/s, equivalent to 50 km/h), and the phase duration was normalized to a range between 0 and 1 considering an upper limit (60 seconds).

Categorical variables such as the day of the week and the hour were transformed into normalized continuous values by dividing by their possible maximum (6 for days of the week indexed from 0 to 6, and 23 for hours of the day). Additionally, the traffic light state was encoded into a three element one hot vector indicating the presence of green, yellow, and red lights.

3.6.3. Output Normalization

Output normalization in this model was implemented to ensure that the actions generated by the neural network remained within valid and physically interpretable ranges in the context of traffic control (Guastella et al., 2023). In particular, the action related to speed was modeled as a continuous variable limited between 0 and 2, where the network produced an initial value through the sigmoid function, which was then scaled by multiplying it by 2 to fit the desired range. This output was sampled from a normal distribution with a fixed standard deviation, allowing some variability in the action to encourage exploration during reinforcement learning training.

On the other hand, the action of enabling or disabling parking was represented as a binary variable obtained by applying a sigmoid function followed by a Bernoulli distribution, which generated a probability for the state 0 or 1. It facilitated the interpretation of the network and stabilized the learning process.

3.6.4. Model Training

The neural network was trained using the Proximal Policy Optimization (PPO) algorithm, an advanced reinforcement learning method that optimizes the agent's policy by controlling the magnitude of changes to ensure stability. At each step, experience data including states, actions taken, log probabilities of those actions, rewards, and estimated values were collected. This data was stored in a buffer and processed in fixed-size batches to update the model.

During the update, Generalized Advantage Estimation (GAE) was calculated to estimate how much better or worse an action was compared to the expected value, helping reduce variance in the gradient estimation. The model computed new action probabilities and compared them with the old ones using a ratio, which was clipped within a specified range to avoid large updates that could destabilize learning. The loss function combined a policy loss based on this clipped ratio, a value loss to adjust the value estimation function, and an entropy term to encourage exploration (Hakeem et al., 2016).

Finally, the Adam optimizer adjusted the network weights through gradient backpropagation over this composite loss function. The model parameters and buffers were periodically saved to preserve progress. This process was repeated iteratively to improve the policy of the agent controlling speed and parking enabling in the SUMO traffic simulation. Algorithm 1 shows the training of a sequential neural network with dense and dropout layers, using ReLU activation and the Adam optimizer, with early stopping to prevent overfitting.

Algorithm 1. PPO Core Training Loop

Require: batch of (states, actions, old_log_probs, rewards, values, dones), model, optimizer

Ensure: updated model parameters

1. $\text{advantages} \leftarrow \text{compute_advantages}(\text{rewards}, \text{values}, \text{dones})$
 2. $\text{advantages} \leftarrow \text{normalize}(\text{advantages})$
 3. $\text{targets} \leftarrow \text{values} + \text{advantages}$
 4. $(\text{vel_logit}, \text{park_logit}, \text{predicted_values}) \leftarrow \text{model}(\text{states})$
 5. $\mu \leftarrow \text{sigmoid}(\text{vel_logit}) \times 2.0$
 6. $\text{dist_vel} \leftarrow \text{Normal}(\mu, 0.1)$
 7. $\text{dist_park} \leftarrow \text{Bernoulli}(\text{park_logit})$
 8. $\text{log_prob_new} \leftarrow \text{log_prob}(\text{dist_vel}, \text{actions}[:,0]) + \text{log_prob}(\text{dist_park}, \text{actions}[:,1])$
 9. $\text{ratios} \leftarrow \exp(\text{log_prob_new} - \text{sum}(\text{old_log_probs}))$
 10. $\text{policy_loss} \leftarrow -\text{mean}(\min(\text{ratios} \times \text{advantages}, \text{clipped_ratios} \times \text{advantages}))$
 11. $\text{value_loss} \leftarrow \text{MSE}(\text{predicted_values}, \text{targets})$
 12. $\text{entropy} \leftarrow \text{mean}(\text{entropy}(\text{dist_vel}) + \text{entropy}(\text{dist_park}))$
 13. $\text{total_loss} \leftarrow \text{policy_loss} + 0.5 \times \text{value_loss} - 0.01 \times \text{entropy}$
 14. $\text{optimizer.zero_grad}()$
 15. $\text{total_loss.backward}()$
 16. $\text{optimizer.step}()$
-

3.6.5. Action Generation and Reward Calculation

The traffic control model makes two main decisions based on the current state of a road segment: speed adjustment and parking management. Both decisions are based on the output of a neural network and use probabilistic functions to encourage exploration during learning. The calculation of actions in the PPO neural network was performed by processing a state vector composed of the number of vehicles, average speed, remaining time of the traffic light phase, time of day, and day of the week, all previously normalized to facilitate learning.

Through its actor critic architecture, the neural network generated two main outputs: a continuous action representing the adjusted traffic speed, and a discrete action determining whether to enable or disable parking in designated areas. The continuous action was interpreted as a multiplier applied to the maximum allowed speed, thus allowing the increase, decrease, or maintenance of vehicle speed based on the current traffic state.

On the other hand, the discrete action corresponded to a binary value where exceeding a predefined threshold enabled the use of parking, while a lower value kept the area closed to avoid congestion. This mechanism allowed the network to make dynamic and adaptive decisions to optimize traffic flow and space availability, adjusting to temporal conditions and vehicular load in the simulated environment. In Table 2, a summary of this process is shown.

During each step of the simulated environment, the neural network-based agent receives a total reward RRR, which guides learning toward more efficient policies. This reward is composed of a weighted combination of several components, each representing a desired objective in urban traffic management. The main terms that make up this reward function are detailed in Table 3 below.

Table 2. Description and Interpretation of Traffic Control Actions

Action	Type	Description	Interpretation
Adjusted Speed	Continuous	Numeric value indicating the factor by which the maximum allowed speed is multiplied	Allows increasing, reducing, or maintaining the traffic speed
Parking	Discrete	Binary value (0 or 1) indicating whether the parking area is enabled or disabled	1. Enables parking; 0. Keeps the area closed to avoid congestion

Table 3. Reward Function Components and Control Objectives

Reward Component	Criterion Description	Control Objective	Weight
Average Speed	Proportional to how close the current speed is to the expected maximum speed	Encourage smooth speeds and avoid stops	+1.0
Congestion	Penalty proportional to the number of vehicles accumulated in the queue	Reduce congestion in the road network	-1.5
Improper Parking Usage	Penalty when parking is enabled on high traffic road segments	Prevent parking from interfering with traffic flow	-1.0
Abrupt Changes in Decisions	Penalty when speed or parking state changes abruptly between steps	Promote stable and smooth decisions	-0.2

3.6.6. Model Hyperparameters

The main hyperparameters of the neural network, which determine the agent's behavior and learning capacity in the traffic control system, are summarized in Table 4.

Table 4. Hyperparameters used in the training of the traffic control system based on PPO

Hyperparameter	Value	Description
Learning rate	0.0003	Learning rate of the optimizer
Gamma	0.99	Discount factor for future rewards
Lambda	0.95	GAE parameter for estimated advantage
Clip range	0.2	PPO clipping for policy stability
Batch size	64	Size of the update batches
n_steps	2048	Steps per policy update
Seed	42	Random seed for reproducibility
Entropy coefficient (ent_coef)	0.01	Coefficient for entropy bonus to encourage exploration
Value function coefficient (vf_coef)	0.5	Coefficient for value function loss in total loss
Max gradient norm (max_grad_norm)	0.5	Clip gradients to stabilize training

3.6.7. Output Adjustment in SUMO

After obtaining the corresponding output data, the parameters in the SUMO scenario are adjusted to reflect the decisions made by the control system. These adjustments may include changes in speed limits, the enabling or disabling of parking zones, or any other parameter relevant to traffic dynamics. This feedback mechanism allows for a direct comparison between the optimization results and key traffic indicators, such as average speed, traffic flow, queue length, and pollutant emissions.

The control system uses a neural network in Python to process real-time traffic data and sends decisions to SUMO via TraCI. This allows the simulation to adapt dynamically, enabling effective testing of traffic management strategies in a controlled environment.

3.7. Parameters for SUMO Scenarios

The scenarios were designed to represent different traffic conditions, varying factors such as vehicle density, time of day, and parking configurations. Each one simulates different traffic situations to analyze the effectiveness of the implemented strategies for managing vehicle flow and parking, to identify improvements and optimize system performance in real environments. Three scenarios were created, and these are:

- 2 x 2 Intersection Scenario (Scenario 1)
- 3 x 3 Intersection Scenario (Scenario 2)
- 4 x 4 Intersection Scenario (Scenario 3)

The first scenario consists of a 2x2 intersection with 2 lanes per edge (one for parking and one for dynamic speed adjustments) and 4 traffic lights operating under a fixed phase cycle: GGrr (30 s), yrrr (5 s), rrGG (30 s), and rrry (5 s), with a constant speed between 40 and 50 km/h. In this controlled environment, traffic flow was uniform, allowing for precise evaluation of the individual effects of each strategy.

The second scenario simulates a 3x3 Manhattan style urban grid, with each of the 9 intersections featuring 2 lanes per direction and 7 traffic lights placed at key junctions. In this case, traffic flow was variable but structured, reflecting realistic mobility in medium density urban areas.

The third scenario increases the complexity to a 4x4 grid with 16 intersections, also with 2 lanes per direction and 12 strategically located traffic lights. Here, traffic was highly dynamic and heterogeneous, simulating typical urban congestion conditions. In all scenarios, a uniform traffic light cycle and a constant speed between 40 and 50 km/h were maintained, enabling controlled comparisons of traffic behavior across different scales, complexities, and demand levels. The day can be a weekday, weekend, or holiday, and the time is classified as normal or peak. The designed scenarios are presented in Figure 2.




Figure 2. Traffic scenarios in SUMO to evaluate the performance of the proposed control strategies

4. Results and Discussion

This section critically analyzes the results from the simulated scenarios to evaluate the performance of urban traffic management strategies. The impact of implementing different input variables across the scenarios was examined, identifying the most relevant ones and optimizing the balance between complexity and accuracy. For this purpose, the most representative edges of each scenario were selected. Additionally, the system's sensitivity and robustness under various conditions were assessed.

4.1. Results for Scenario 1

This experimental study evaluated the impact of local interventions on a road segment consisting of two lanes: one subject to dynamic speed adjustments and another designated for parking. The neural network used as input variables the number of vehicles, the average speed, and the traffic light phases, considering traffic conditions and road structure to model the system's dynamic state. Based on these inputs, the network assessed congestion and vehicle flow in real time and determined speed limit adjustments as well as the activation or deactivation of parking zones. This approach enabled the optimization of traffic flow and the minimization of delays, maintaining a balance between efficiency and safety in the adaptive management of urban traffic control. The results corresponding to the neural network outputs are summarized in Table 5.

Table 5. Results for scenario 1 applying the neural network

Edge	Traffic Light Phase	Vehicles	Average Speed	Adjusted Speed	Parking
E1	GGrr	6	48	80 (increased)	No
E2	rrGG	9	41	0	Yes
E3	yyrr	7	43	43 (normal)	No
E4	GGrr	5	50	79 (increased)	No
E5	rrGG	10	40	0	Yes
E6	rryy	6	47	41 (normal)	No
E7	GGrr	8	42	20 (reduced)	No
E8	rrGG	7	45	20 (reduced)	No
E9	GGrr	5	49	80 (increased)	No
E10	yyrr	10	39	0	Yes

4.2. Results for Scenario 2

In the current model, any lane within each road segment can be dynamically selected for control functions, either as a speed regulation lane or as a parking zone, based on the input variables and real time traffic conditions. For this implementation, the neural network uses as input variables the average speed, the number of vehicles present, the traffic light phase, and the time of day, allowing it to capture both the operational state of the system and the typical temporal variations in traffic flow. In this way, the system makes adaptive decisions that optimize the use of road space and improve the overall efficiency of urban traffic. The corresponding results from the neural network outputs are presented in Table 6.

Table 6. Results for scenario 2 applying the neural network

Edge	Traffic Light Phase	Vehicles	Average Speed	Time of Day	Adjusted Speed	Parking
E1	GGrr	6	48	Normal	80 (increased)	No
E2	rrGG	9	42	Peak	20 (reduced)	No
E3	yyrr	7	44	Normal	44 (normal)	No
E4	rryy	5	50	Peak	80 (increased)	No
E5	rrGG	10	38	Peak	0	Yes
E6	rryy	7	41	Normal	0	Yes
E7	rrGG	8	43	Peak	20 (reduced)	No
E8	GGrr	6	45	Normal	0	Yes
E9	GGrr	5	49	Normal	80 (increased)	No
E10	yyrr	7	39	Normal	0	Yes
E11	rryy	6	46	Normal	46 (normal)	No
E12	rrGG	5	47	Peak	47 (normal)	No

4.3. Results for Scenario 3

In this model, each lane within a road edge can be dynamically assigned to specific control functions, either as a speed regulation lane or as an area designated for parking, depending on real time traffic conditions and available input variables. For this implementation, the neural network considers as input variables the average speed, the number of vehicles present, the traffic light phase, the time of day, and, as an additional enhancement, the day of the week. This inclusion allows the system to capture recurring patterns of vehicle behavior associated with weekdays or weekends, thereby improving the model's predictive capacity. Thanks to this architecture, the model can make adaptive decisions that optimize the use of road space and contribute to a more efficient management of urban traffic. The corresponding results from the neural network outputs are presented in Table 7.

As the road network topology and environmental complexity increases, the neural network itself progressively incorporates a greater number of input variables. This evolution enables it to adapt more precisely to dynamic contexts by integrating both immediate operational data and temporal or contextual information. The expansion of its input structure reflects the model's ability to enrich its understanding of the system, improving decision making performance and optimizing intelligent urban traffic management. This behavior can be observed across the three analyzed scenarios: in Scenario 1, with a simple road network, a basic input set is sufficient for stable environments; in Scenario 2, the inclusion of the time of day allows the system to recognize daily patterns and adjust strategies based on hourly demand.

And in Scenario 3, adding the day of the week strengthens the system's predictive capabilities, enabling it to anticipate variations typical of weekdays, weekends, or holidays. This progression demonstrates the neural network's ability to scale its intelligence as the demands of the traffic environment increase.

4.4. Comparison of Performance Indicators Across Scenarios

This section presents a detailed comparison of the main traffic indicators against the system outputs. The analysis allows identifying how these indicators vary under different traffic control conditions. To

Table 7. Results for scenario 3 applying the neural network

Edge	Traffic Light Phase	Vehicles	Average Speed	Time of Day	Day	Adjusted Speed	Parking
E1	rryy	7	41	Normal	weekday	0	Yes
E2	rrGG	9	42	Peak	weekday	20 (reduced)	No
E3	GGrr	5	49	Normal	weekend	80 (increased)	No
E4	GGrr	5	50	Peak	Holiday	80 (increased)	No
E5	rryy	5	47	Peak	weekday	47 (normal)	No
E6	Yyrr	6	48	Peak	weekend	0	Yes
E7	GGrr	7	39	Normal	Holiday	0	Yes
E8	GGrr	8	43	Peak	weekday	20 (reduced)	No
E9	yyrr	7	44	Normal	Holiday	44 (normal)	No
E10	rryy	6	46	Normal	weekday	46 (normal)	No
E11	rrGG	6	45	Normal	Holiday	0	Yes
E12	rrGG	10	38	Normal	weekend	0	Yes
E13	GGrr	6	50	Normal	weekend	80 (increased)	No
E14	rrGG	7	44	Peak	weekday	20 (reduced)	No

this end, two factors were considered: the traffic scenario, defined by the topology of the road network (Scenarios 1, 2, and 3), and the control strategy, with two levels (no control and control using a neural network). The combination of these factors resulted in six experimental cases, each representing a system output under specific conditions. A general average of all edges was calculated for each indicator in each scenario.

This comparative approach facilitates understanding the effectiveness and limitations of each implemented strategy.

Figure 3 compares the average travel times across three traffic scenarios under normal control and neural network control. In Scenario 1, the neural network slightly improves travel times. In Scenario 2, the greatest improvement is observed, with clear reductions in both time and variability. In Scenario 3, although the median decreases, congestion limits the benefits. Overall, the neural network is more effective under moderate traffic conditions for this indicator.

Figure 4 compares waiting times across three traffic scenarios under normal control and neural network control. In Scenario 1, the neural network significantly reduces both the median and the dispersion of waiting time, showing a clear improvement over normal control. In Scenario 2, the median remains similar, but greater variability is observed with the use of the neural network. In Scenario 3, waiting time increases considerably under intelligent control, both in median and in extreme values. Overall, the neural network improves this indicator under low traffic conditions, but in highly congested networks it may produce adverse effects.

Figure 5 compares queue lengths across the three traffic scenarios. In Scenario 1, the neural network reduces both the median and the dispersion, indicating an improvement in road space distribution. In Scenario 2, although the median slightly increases with the neural network, a decrease in minimum values and a more uniform distribution are observed. In Scenario 3, queue length increases




Figure 3. Results of the comparative analysis of traffic scenarios simulated in SUMO based on average travel time




Figure 4. Results of the comparative evaluation of traffic scenarios simulated in SUMO based on vehicle waiting time

under intelligent control, both in the median and in extreme values, reflecting a greater vehicle accumulation. Overall, neural network control proves effective under low traffic conditions, but its performance declines in highly congested contexts for this indicator.

Figure 6 compares CO₂ emissions across the three traffic scenarios. In Scenario 1, the use of the neural network reduces both the median and the dispersion of emissions, indicating a positive effect on environmental efficiency. In Scenario 2, a significant improvement is observed: the neural network lowers both the median and the extreme values, reflecting smoother circulation and fewer stops. In Scenario 3, despite high congestion, the intelligent control manages to reduce emissions compared to




Figure 5. Results of the comparative assessment of traffic scenarios simulated in SUMO based on queue length




Figure 6. Evaluation of CO₂ emission levels under different traffic conditions of the various scenarios simulated in SUMO

normal control, with a lower median and reduced variability. Overall, the neural network shows favorable performance in all scenarios for this indicator, contributing to an effective reduction in pollutant emissions.

Figure 7 compares accident damage probabilities (ADP) across the three traffic scenarios. In Scenario 1, the neural network slightly reduces both the median and the dispersion, indicating a moderate positive effect. In Scenario 2, the improvement is more evident: intelligent control reduces both the median and extreme values, suggesting more stable and safer driving. In Scenario 3, however, the probability of accidents increases under neural network control, with higher values and greater variability than under normal control. Overall, the intelligent system improves this indicator in low and medium congestion networks but may be counterproductive under severe traffic conditions.




Figure 7. Evaluation of the traffic indicator accident probability in different scenarios simulated in SUMO

Figure 8 shows the comparison of traffic flow across the three analyzed scenarios, under normal conditions and with neural network control. In all cases, an increase in traffic flow is observed when the neural network is implemented, suggesting an optimization of traffic without compromising system stability. This positive trend remains consistent across the different scenarios, reflecting the potential of intelligent control to improve vehicular circulation through adaptive regulation that prevents congestion and promotes greater efficiency in the monitored segments.




Figure 8. Evaluation of traffic flow performance under different traffic conditions in the scenarios simulated in SUMO

Intelligent control based on neural networks shows variable performance depending on the scenario. In low-traffic networks, it maintains stable and efficient conditions, contributing to smooth traffic management. Under moderate traffic conditions, it optimizes urban mobility by reducing travel times, emissions, and accidents. In congested networks, it provides significant improvements such as reduced emissions and travel times, demonstrating its potential to enhance efficiency in diverse

contexts. This highlights the importance of adapting intelligent control to the operational context of the road network to maximize its benefits.

Our proposed work, which integrates speed control and parking management within a single system in SUMO, presents certain limitations when considering its implementation outside the simulated environment. According to the Technology Readiness Levels (Héder, 2017), the current validation was carried out in a simulation environment (TRL 2–3), under controlled traffic conditions and with complete data. The transfer of the system to controlled real-world environments (TRL 6–7) introduces additional variability, such as heterogeneous driver behavior, environmental conditions, and potential sensor errors, which may affect the stability of the learned policies.

Implementation in fully operational environments (TRL 8–9) requires ensuring robustness, safety, fault tolerance, seamless integration with existing infrastructure, and user acceptance. Therefore, although the simulation results are promising, a progressive validation plan and gradual adaptation are recommended to reliably transfer the observed benefits to real urban scenarios.

5. Conclusions

The results obtained show that the integration of speed control and dynamic parking regulation within a single intelligent management system has a positive impact on urban traffic performance. This integrated system responds efficiently to changing traffic conditions by dynamically adjusting both vehicle speeds and parking lane availability, allowing for smoother and more effective management of urban space. Under low traffic conditions, the system maintains stable and efficient circulation, avoiding congestion and ensuring optimal use of available infrastructure. In moderate traffic scenarios, the system demonstrates its ability to reduce travel times, minimize emissions, and decrease accidents by optimizing both speed and parking utilization in real time. This dual intervention allows for better distribution of vehicles on the roads and smoother circulation. In more congested traffic situations, the system remains effective, achieving notable improvements in reducing emissions and travel times, highlighting its capacity to adapt to varying congestion levels while continuing to optimize vehicle flow. Thanks to the integration of both control strategies, the system ensures better traffic distribution across the road network, preventing congestion points and facilitating the movement of vehicles more efficiently. This approach underscores the importance of having a traffic management system that goes beyond a single aspect, such as parking control or speed regulation, and combines multiple adaptive strategies to improve the overall efficiency of the urban transportation system.

As future work, it is suggested to apply this system in different urban infrastructures, such as residential or commercial areas, to assess its adaptability. Additionally, exploring the integration of public transport prioritization or emergency management systems could further optimize traffic flow in critical situations.

Funding

This research work has been carried out thanks to the support of the Ministry of Science and Innovation of Spain (MICINN), Project CICYT PRECON-I4: «Models and Platforms for Predictable, Safe, and Reliable Industrial Computer Systems» PID2021-124502OB-C41.

References

- Alekszejenkó, L., & Dobrowiecki, T. (2023). SUMO simulations for federated learning in communicating autonomous vehicles: A survey on efficiency and security. *SUMO User Conference Proceedings*, 4(1), 115–121. <https://doi.org/10.52825/scp.v4i.221>
- Amiour, Y., Waygood, E. O. D., & Van den Berg, P. E. W. (2022). Objective and perceived traffic safety for children: A systematic literature review of traffic and built environment characteristics related to safe travel. *International Journal of Environmental Research and Public Health*, 19(5), 2641. <https://doi.org/10.3390/ijerph19052641>
- Azfar, T., & Ke, R. (2024). Traffic co-simulation framework empowered by infrastructure camera sensing and reinforcement learning. *arXiv preprint*, arXiv:2412.03925. <https://doi.org/10.48550/arXiv.2412.03925>
- Bagheri, M., & Sarvi, M. (2023). Implementing artificial neural network-based gap acceptance models in the simulation model of a traffic circle in SUMO. *Transportation Research Record*, 2677(1), 1–12. <https://doi.org/10.1177/03611981231167420>
- Beil, C., Kendir, M., Ruhdorfer, R., & Kolbe, T. H. (2022). Dynamic and web-based 4D visualization of streetspace activities derived from traffic simulations and semantic 3D city models. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-4/W2-2022, 29–36. <https://doi.org/10.5194/isprs-annals-X-4-W2-2022-29-2022>
- Chen, L. (2024). Development of a parking system based on learning automata. *Multimedia Tools and Applications*, 83(13), 61165–61180. <https://doi.org/10.1007/s11042-023-18031-9>
- Ehlers, S. F. G. (2019). Traffic queue length and pressure estimation for road networks with geometric deep learning algorithms. *arXiv preprint*, arXiv:1905.03889. <https://doi.org/10.48550/arXiv.1905.03889>
- Farrag, S., El-Hansali, M. Y., Yasar, A., & Shakshuki, E. M. (2020). Simulation-based evaluation of using variable speed limit in traffic incidents. *Procedia Computer Science*, 175, 340–348. <https://doi.org/10.1016/j.procs.2020.07.049>
- Guastella, D. A., Montero-Porras, E., Morales-Hernández, A., & Bontempi, G. (2023). Traffic modeling with SUMO: A tutorial. *arXiv preprint*, arXiv:2304.05982. <https://doi.org/10.48550/arXiv.2304.05982>
- Guo, J., Cheng, L., & Wang, S. (2022). CoTV: Cooperative control for traffic light signals and connected autonomous vehicles using deep reinforcement learning. *arXiv preprint*, arXiv:2201.13143. <https://doi.org/10.48550/arXiv.2201.13143>
- Hakeem, A., Gehani, N., Ding, X., Curtmola, R., & Borcea, C. (2016). On-the-fly curbside parking assignment. *Proceedings of the 8th International Conference on Mobile Computing, Applications and Services (MobiCASE 2016)*, 1–10. <https://doi.org/10.4108/eai.30-11-2016.2267101>
- Héder, M. (2017). From NASA to EU: The evolution of the TRL scale in public sector innovation. *The Innovation Journal: The Public Sector Innovation Journal*, 22(2), 1–23. https://web.archive.org/web/20171011071816/https://www.innovation.cc/discussion-papers/22_2_3_heder_nasa-to-eu-trl-scale.pdf
- Hussain, A., Wang, T., & Jiahua, C. (2020). Optimizing traffic lights with multi-agent deep reinforcement learning and V2X communication. *arXiv preprint*, arXiv:2002.09853. <https://doi.org/10.48550/arXiv.2002.09853>

- Kheterpal, N., Vinitsky, E., Behrisch, M., Bieker-Walz, L., Wu, C., & Bayen, A. M. (2018). Flow: Deep reinforcement learning for control in SUMO. *EPiC Series in Engineering*, 2, 134–151. https://flow-project.github.io/papers/Flow_Deep_Reinforcement_Learning_for_Control_in_SUMO.pdf
- Mousavi, S. S., Schukat, M., & Howley, E. (2017). Traffic light control using deep policy-gradient and value-function based reinforcement learning. *arXiv preprint*, arXiv:1704.08883. <https://doi.org/10.48550/arXiv.1704.08883>
- Rahman, M. R., Akhter, S., & Sumit, S. H. (2020). A SUMO based simulation framework for intelligent traffic management system using deep-neuro-fuzzy model. *Journal of Traffic and Logistics Engineering*, 8(1), 1–5. <https://www.jtle.net/uploadfile/2020/0529/20200529050106994.pdf>
- Reza, S., Ferreira, M. C., Machado, J. J. M., & Tavares, J. M. R. S. (2023). A citywide TD-learning based intelligent traffic signal control for autonomous vehicles: Performance evaluation using SUMO. *Expert Systems*, 42(1), e13301. <https://doi.org/10.1111/exsy.13301>
- Sabour, S., Rao, S., & Ghaderi, M. (2021). DeepFlow: Abnormal traffic flow detection using Siamese networks. *arXiv preprint*, arXiv:2108.12016. <https://doi.org/10.48550/arXiv.2108.12016>
- Samarakoon, G. S., & Sivakumar, T. (2022). Microscopic simulation of parking violations in curbside with-flow bus priority lanes using SUMO Traffic Control Interface (TraCI). *SUMO User Conference Proceedings*, 2023(1), 45–52. <http://dl.lib.uom.lk/handle/123/19416>
- Skorobogatchenko, D., Zhokhov, V., Astafurova, O., & Fantrov, P. (2023). Neural network-based optimization of traffic light regulation of a transport hub with data fetched during simulation in SUMO package. In A. G. Kravets, M. V. Shcherbakov, & P. P. Groumpos (Eds.), *Creativity in Intelligent Technologies and Data Science* (pp. 337–350). Springer. https://doi.org/10.1007/978-3-031-44615-3_23
- Sun, K., & Yu, H. (2024). Reinforcement learning for freeway lane-change regulation via connected vehicles. *arXiv preprint*, arXiv:2412.04341. <https://doi.org/10.48550/arXiv.2412.04341>
- Ye, C., Xiong, G., Shang, J., Dai, X., Gong, X., & Lv, Y. (2025). SUMO-MCP: Leveraging the Model Context Protocol for autonomous traffic simulation and optimization. *arXiv*. <https://arxiv.org/abs/2506.03548>
- Ye, Q., Feng, Y., Han, J., Stettler, M., & Angeloudis, P. (2021). A reinforcement learning-based adaptive control model for future street planning: An algorithm and a case study. *arXiv preprint*, arXiv:2112.05434. <https://doi.org/10.48550/arXiv.2112.05434>
- Yin, S., Cui, Z., & Wang, Y. (2022). Reinforcement learning for curbside space management with infrastructure autonomy and mixed vehicle connectivity. In *2022 IEEE 25th International Conference on Intelligent Transportation Systems (ITSC)* (pp. 3276–3282). IEEE. <https://doi.org/10.1109/ITSC55140.2022.9922149>