AI-Driven Traffic Signal Control System to Reduce CO2 Emissions*

Oleksander Ryzhanskyi[1], Vitaliy Pavlyshyn[1], Pavlo Radiuk[1,*], Eduard Manziuk[1], Oleksander Barmak[1] and Iurii Krak[2,3]

[1] Khmelnytskyi National University, 11, Institutes str., Khmelnytskyi, 29016, Ukraine
[2] Taras Shevchenko National University of Kyiv, 64/13, Volodymyrska str., Kyiv, 01601, Ukraine
[3] Glushkov Cybernetics Institute, 40, Glushkov Ave., Kyiv, 03187, Ukraine

Abstract

Rapid urbanization and the resulting traffic congestion have made the reduction of transport-related carbon-dioxide (CO2) emissions a pivotal goal for sustainable-city initiatives. This paper introduces an artificial-intelligence (AI) traffic-signal control system that exploits deep reinforcement learning (DRL) to achieve dynamic, emission-aware phase scheduling at signalized intersections. Unlike conventional fixed-time or actuated controllers, the proposed framework continuously perceives multi-modal state information, queue length, average speed, arrival rates and projected emergency-vehicle trajectories, through a digital twin fed by microscopic traffic simulation. A convolutional DRL agent, trained with proximal-policy optimization and shaped by a composite reward that penalizes total CO_2, queue growth and delay while rewarding throughput, learns an adaptive policy that balances environmental and mobility objectives. In controlled experiments mirroring a mid-size European arterial, the AI agent lowers cumulative CO2 emissions by 18% without sacrificing pedestrian service levels. The learned policy exhibits robust generalization under varying traffic demand patterns and mild sensor noise, suggesting strong transferability.

Keywords

Traffic light control, Deep reinforcement learning, SUMO, CO_2 emissions, Hilly topography of the road

1. Introduction

Optimizing traffic light systems to curb carbon dioxide (CO2) emissions is a pressing challenge, essential not only for current urban sustainability but also for the long-term preservation of ecosystems and the global biosphere [1]. Achieving this objective is not straightforward due to several competing constraints, such as road gradient variations, prioritization for emergency vehicles, and the need for fair treatment of drivers, pedestrians, and other road users [2, 3].

One of the most promising techniques for intelligent traffic signal management is reinforcement learning (RL), particularly its deep learning-enhanced variant, deep reinforcement learning (DRL) [4, 5]. DRL algorithms are capable of dynamically adapting to fluctuating traffic patterns by continuously interacting with the environment to learn optimal decision-making policies that maximize cumulative rewards [6, 7]. These learned strategies aim to enhance traffic flow and reduce vehicular emissions.

The primary contribution of this study is the development of a DRL-based simulation framework for determining traffic light control policies at a single intersection, with a targeted goal of minimizing CO2 emissions.

The remainder of the article is organized as follows: the “Related Works” section surveys recent efforts in traffic signal control aimed at emission reduction. The “Materials and Methods” section introduces the proposed model tailored for local road intersections. Finally, “Results and Discussion” presents the simulation outcomes and their interpretation.

The Second International Conference of Young Scientists on Artificial Intelligence for Sustainable Development (YAISD), May 8-9, 2025, Ternopil, Ukraine

* Corresponding author.
alex@ualeaders.com (O. Ryzhanskyi); vitaliy@ualeaders.com (V. Pavlyshyn); radiukp@khmnu.edu.ua (P. Radiuk); manziuk.e@khmnu.edu.ua (E. Manziuk); barmako@khmnu.edu.ua (O. Barmak); iurii.krak@knu.ua (I. Krak)
0009-0000-4664-5195 (O. Ryzhanskyi); 0009-0003-0702-9359 (V. Pavlyshyn); 0000-0003-3609-112X (P. Radiuk); 0000-0002-7310-2126 (E. Manziuk); 0000-0003-0739-9678 (O. Barmak); 0000-0002-8043-0785 (I. Krak)

© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR-WS.org/Vol-3974/paper02.pdf

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073

2. Related Works

This section synthesizes prior investigations into traffic light signal control focused on reducing CO_2 emissions. A common thread among these studies is the widespread adoption of the SUMO simulation model [8]. Researchers have employed a variety of optimization strategies, such as fuzzy logic [9], game theory (utilizing Nash equilibrium) [10], extreme learning machines [11], genetic algorithms [12], and an array of modeling techniques [9, 12, 14]. Nevertheless, RL and its derivatives dominate the field [13–15].

For example, reference [9] developed a three-phase fuzzy traffic control framework utilizing the SUMO model, refining fuzzy rules to achieve superior CO_2 emission reductions. Meanwhile, study [10] distinguished itself by factoring pedestrians into traffic light management, applying a game-theoretic method to synchronize and reduce wait times for vehicles and pedestrians alike, with SUMO simulations confirming decreases in delays and CO_2 emissions.

In [11], a dual-algorithm approach was proposed, integrating a passive extreme learning machine with periodic mini-batch learning (PB-ELM) for traffic forecasting and a maximum pressure control (MPA) algorithm for traffic light regulation, validated through SUMO for enhanced congestion relief. Reference [12] devised an intersection management framework that optimizes both traffic flow and energy use within a vehicle-infrastructure setting, employing a genetic algorithm. Simulation outcomes revealed reductions in total queued vehicles, maximum queue sizes, and average travel times by 77.81%, 33.33%, and 10.95%, respectively.

Using SUMO, study [16] illustrated that adaptive traffic lights markedly lower fuel use and emissions compared to fixed systems, achieving reductions in travel time by 82.84% and CO_2 emissions by 51.2%. In [17], an intelligent traffic system prioritizing emergency vehicles was introduced, capable of adjusting signal timings based on vehicle classification (e.g., “emergency” versus “passenger”).

Work [18] tackled traffic light arrangements to ease congestion, proposing models to optimize signal timing graphs for specific goals [19], with SUMO assessing their efficacy. Study [13] utilized RL algorithms and SUMO to refine urban traffic management in Taipei, yielding enhancements in traffic capacity and travel duration. In [14], the D3QN (Dual Deep Q Network) algorithm effectively managed traffic light timings across varying traffic densities, outperforming conventional methods like Webster’s fixed-time control in SUMO trials.

Focusing on isolated intersections, [20] optimized traffic light signals using a double deep Q-network (DDQN), an RL variant. Notably, [15] stressed the importance of tailoring reward functions in RL to align agent actions with CO_2 reduction objectives, testing various reward structures with SUMO and its emission model to gauge performance sensitivity. Reference [21] offered a novel traffic light control alternative by incorporating predictive noise modeling with RL and a SeqtoSeq-LSTM framework [22], significantly improving noise levels, CO_2 emissions, and fuel efficiency over traditional systems.

Further insights come from works [23–25], which showcase successful DQN applications to analogous challenges. Study [23] introduced a “pressure” concept for regional signal coordination, supported by extensive SUMO experiments, including a real-world Manhattan scenario with 2510 traffic lights. Reference [24] provided a toolkit for modern DQN implementations, while [25] leveraged value-based RL for online learning with interpretable policy functions [26]. These findings highlight the demand for intelligent traffic light systems that reduce CO_2 emissions while addressing vehicle and pedestrian queues, emergency vehicle access [27], and other factors, necessitating further exploration. The prevalent use of RL and SUMO underscores their promise.

Thus, this work is aimed at advancing sustainable development by curbing CO_2 emissions through an intelligent traffic signal control model tailored for complex terrains, balancing the needs of drivers, pedestrians, and emergency services. This will be accomplished by devising a simulation and DQN model for local intersection traffic light control to minimize CO_2 emissions.

3. Materials and Methods

In this study, we delineate the elements of a holistic approach crafted by the authors to address traffic light control at local intersections, targeting CO_2 emission reductions (see Figure 1). A key attribute of this method is the creation of a unified metric and solution that, beyond lowering CO2 emissions, accounts for complex road topographies (e.g., slopes); drivers' desires for swift destination arrivals; pedestrians' needs, especially near crowded locales like schools and malls; rapid transit for emergency services; the presence of heavy vehicles; and road emergencies. Moreover, this approach is engineered for adaptability to encompass additional priorities as they arise.

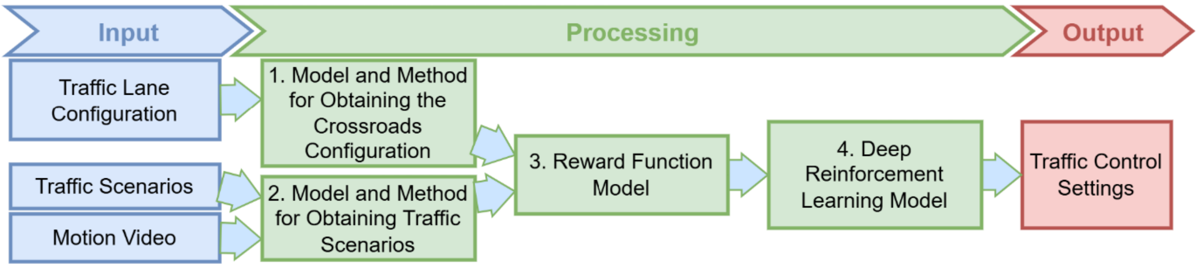

Figure 1: Steps in obtaining and using optimal traffic light signal controls using deep learning with reinforcement for local regulated road intersections.

Note that for traffic signal control (TSC), one of the most widely used solution tools is the reinforcement learning algorithm, namely its modern implementation Deep Reinforcement Learning (DRL), since DRL can learn to adapt to changing traffic demands [5, 6, 28]. Another tool used in this work is the traffic modeling package Simulation of Urban Mobility (Eclipse SUMO or simply SUMO v1.21.0 [29]) with built-in CO_2 emission models (e.g., Handbook Emission Factors for Road Transportation – HBEFA 3.1 [30]) and complex road terrain models.

Figure 1 shows the decomposition of the proposed approach into stages. Note that this article presents research results for stages 3 and 4. Other stages will be considered in future studies.

To test the approach's feasibility, it is proposed to simulate traffic through an intersection using **Algorithm 1**.

Algorithm 1.
Simulating the movement of cars through intersections.

Input: intersection properties and traffic scenarios
Initialize event: *Evnt* about the appearance of a new car, relative to the start of the simulation with the following properties: lane from which the car is declared; lane into which the car is leaving; speed of the car
Set time counter at 0
For simulation
    // Check if there is an event that corresponds to the current simulation time. If there is such an event, add the corresponding car to the simulation on the specified lane with the specified parameters:
    If *Evnt.Time* = current simulation time **Then** Add the car to the lane
    // Update for each vehicle inside the simulation:
If the vehicle is near the intersection **Then:**
        // Checking the condition of the traffic light:
        If the green signal allows movement in its direction **Then:**
            A car crosses an intersection into a lane
        If the signal is red **then:**
            The car stops
    If the car is not near the intersection **then:**
        If the path is free **then:**
            The car continues to move at the current speed
        If Obstacle Ahead Then**:**
            The car stops
    // Updating the state of the traffic light (at each moment of time we check if it is time to change the phase of the traffic light)
    If it's time to change the phase of the traffic light **Then:**
        Moving the traffic light to the next phase

// Checking the condition for the end of the simulation cycle
If All Cars have completed their movement and there are no more events to process **Then:**
    Simulation completes
Otherwise:
    Increase the time counter by Δt
// Updating the properties for the lane:
Set properties for the lane
// Updating properties for all lanes
Set properties for all lanes ($TotalCO2, QueueLength, TotalWait, TotalAvgSpead$)
End For Simulation

In *Stage* 3, it is proposed to synthesize a reward function for the intersection to reduce CO_2 emissions, considering the properties of the observation environment.

Note that the motion simulation iteration is performed every Δt, but the reward is calculated not for each iteration, but after $N\Delta t$, i.e. the system is controlled for a predetermined number of iterations and the state of the system changes during this period at certain points in time.

Therefore, the value of r_t (reward at certain steps of the simulation) is proposed to be defined as an integral indicator consisting of the following intermediate components:

$$r_t = \sum_i w_i r_i, \tag{1}$$

where w_i is a weighting factor determined empirically, r_i is a fine or reward.

In *Stage* 4, it is proposed to synthesize a DQN model to achieve the goal of reducing CO_2 emissions, taking into account the properties of the observation environment and the integral loss function. It should be noted that the study hypothesized that optimal solutions for traffic light control should be sought not for a specific intersection, but for a sequence of intersections.

The study proposes to use the independent DQN (IDQN) model for a sequence of intersections and vehicle traffic scenarios. The essence of IDQN is that we will use several networks – each intersection's own DQN. An approach is proposed where each agent learns independently and simultaneously its own policy, considering other agents as part of the environment.

In IDQN, each agent observes the observation space of its intersection and chooses a separate action and receives a reward. Since the agent "sees" the observation space only from its own perspective, IDQN can be implemented by assigning each agent its own observation history of actions. In deep reinforcement learning, this can be implemented by giving each agent "its own" DQN to perform the observations and produce the actions.

So, we will form IDQN models for intersection sequences and traffic scenarios:

$$\begin{gathered} M = \{m_1, m_2, \dots, m_N\}, \\ m_i = \begin{pmatrix} InputLayer, Conv2d, ReLU, Flatten, Linear1, ReLU, Linear2, ReLU, \\ Linear3, DiscreteActionValueHead \end{pmatrix}, \\ i = 1,2, \dots, N, \end{gathered} \tag{2}$$

where N – is the number of intersections in the sequence, $InputLayer$ – is the agent's observation space (at each observation time) in the form of a matrix ($C \times H \times W$), where C – is the number of channels, H – is the number of lanes that the traffic lights regulate, W – is the number of properties for the lanes, $Conv2d$ – is a convolutional layer that detects spatial dependencies between lanes, takes into account the relationships between the properties of the lanes and prepares data for fully connected layers that make decisions about changing the phases of the traffic light; convolution with a kernel (2×2), stride (stride) 1 and no padding (padding) is used, $ReLU$ – is a transfer function (ReLU(x) = max(0, x) i.e. this function replaces all values less than zero with zero and leaves values greater than zero unchanged), $Flatten$ – is a layer that converts the output of the $Conv2d$ layer into a one-dimensional vector of size $64 * (H-1) * (W-1)$, which is then fed to the input of the fully connected layers, $Linear1$, $Linear2$, $Linear3$are fully connected layers: the first layer converts a vector of dimension $64 * (H-1) * (W-1)$ into a vector of 64 elements, the second layer supports dimension 64, the third linear layer outputs a vector whose size is equal to the number of actions, $DiscreteActionValueHead$ – is a layer that processes the output of the last linear layer, ensuring that the final Q values for discrete actions (traffic light phase switching) are obtained.

Next, we present an **Algorithm 2** for training IDQN models m_i (2).

Algorithm 2.
Training Reinforcement Deep Learning Models $\boldsymbol{m_i}$.

1: ***Input***: properties of the intersection and traffic light and properties of the traffic scenario at each moment in time of the simulation according to ***Algorithm 1***
2: **For** i=1… N
3: Forming the input matrix $InputLayer$ for m_i
4: The input matrix *InputLayer* passes through the convolutional layer *Conv2d* to extract features and learn patterns between traffic lanes, and then through the *Flatten* layer and fully connected layers (*Linear1*, *Linear2*, *Linear3*), which form a vector of *Q*-values.
5: Each element of the vector *Q* corresponds to the expected reward for a particular action, so its size is equal to the number of possible actions.
6: Based on the obtained *Q*-values, an action is selected – that is, the action where *Q* is maximum. The observation space determines the number of possible actions. It is used to form the last linear layer of the network. Each element of this layer corresponds to a *Q*-value for a specific action.
7: **End For**
8: Apply the Reward function – the reward is received from the environment after performing the action and is used to correct *Q*-values during training.
9: Imprint: Trained Models m_i

Thus, *M*-networks learn to predict the future reward for each action, allowing for optimal decisions. Note that during training *M* uses: reward (from the environment), next state (calculated after performing the action) obtained from motion simulation, discount factor $\gamma \in [0,1]$ (if $\gamma \sim 1$, then the agent focuses on long-term benefits if $\gamma \sim 0$ the agent prefers immediate rewards), which shows the difference in the importance of future and current rewards, and target *Q*-values. We also note that the agent's observation space for an intersection will be a matrix of all lanes and their properties.

4. Results and Discussion

To validate the proposed approach, a simple intersection with one lane in each direction was used (Figure 2).

Through this intersection, the flow of vehicles with different CO_2 emissions were simulated using ***Algorithm 1***. 12 routes were defined (from each direction to each possible target direction), and traffic flows were generated in four-time intervals (0–25000, 25000–50000, 50000–75000, 75000–100000 seconds). For each flow, the following were specified: route (in direction of travel), simulation start and end times, intensity, initial speed, lane position, and lane selection.

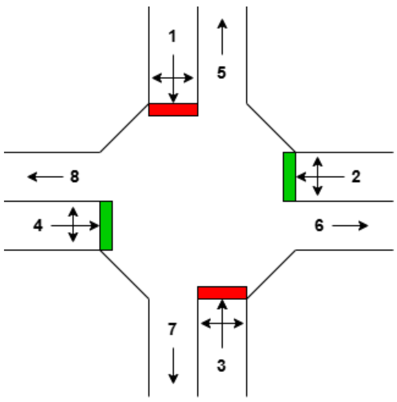

Figure 2: Simple intersection model.

Thus, vehicles were generated at different times and with different intensities, following specific routes through the intersection.

Then, the IDQN model trained by ***Algorithm 2*** on the generated traffic was compared with known similar algorithms [23–25], according to the criteria: (a) CO_2 emissions, (2) the total waiting time of all cars, and (3) the total time required for all cars to pass through the intersection.

For the experimental study presented below, the following component reward functions were used: (1): $TotalCO2$ – total emissions CO_2, $QueueLength$ – total length of the queue at the intersection, $TotalWait$ – total waiting time at the intersection, $TotalAvgSpead$ – average speed at the intersection. Accordingly: $-w_1r_1 = -1 \cdot \frac{TotalCO2}{NumLanes}/10E3$, $-w_2r_2 = -0.1 \cdot \frac{QueueLength}{NumLanes}/10E2$, $-w_3r_3 = -0.1 \cdot \frac{TotalWait}{NumLanes}/10E2$, $-w_4r_4 = 0.1 \cdot \frac{TotalAvgSpead}{NumLanes}/10E2$.

Next, we present the obtained results. Figure 3 shows a comparison of CO_2 emissions for four approaches (approaches [23–25] and ours) after 100 training cycles. The vertical axis shows the CO_2 level in mg, and the horizontal axis shows the training cycles of the neural network. As can be seen from the graph, for our proposed approach, CO_2 emissions are the lowest.

Next, we present the results obtained. Figure 3 shows a comparison of CO_2 emissions for four approaches (ours and approaches [23–25] after 100 training cycles.

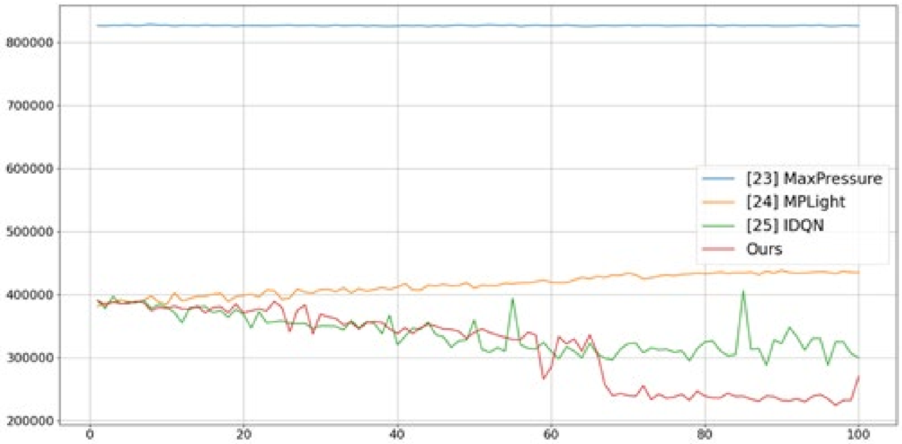

Figure 3: Comparison of CO_2 emissions for different approaches.

The vertical axis shows the CO_2 level in mg, and the horizontal axis shows the neural network training cycles. From Figure 3, for our proposed approach, CO_2 emissions are the lowest.

Figure 4a compares the total waiting time of all cars for the four approaches after 100 training cycles. The vertical axis shows the total waiting time of the full simulation in seconds, and the horizontal axis shows the neural network training cycles. As can be seen from the graph, the proposed approach performs at the same level as other approaches.

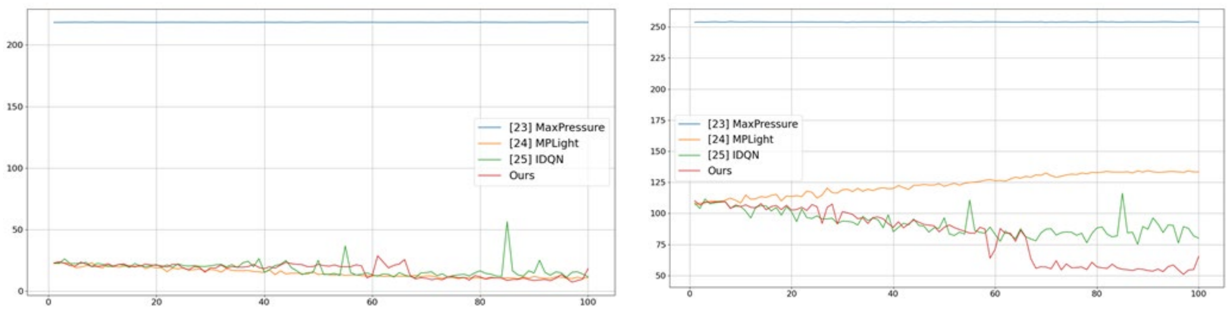

(a) (b)

Figure 4: (a) Comparison of the total waiting time of all cars, (b) Comparison of the total time it takes for all cars to travel through the intersection

Figure 4b compares the total time it takes for all cars to travel through the intersection. The vertical axis plots time in seconds, and the horizontal axis plots neural network training cycles. As can be seen from the graph, the proposed approach performs slightly better than the others.

Table 1 shows comparative results for different approaches to the same motion simulation.

The results obtained confirm the ability of the proposed intelligent model to obtain optimal traffic light signal control for a local intersection using simulation modeling and DQN to reduce CO_2 emissions.

Table 1
Performance comparison of different methods.

Approach	Average duration of travel, sec	Average Queue length	CO_2, mg
MaxPressure [23]	253.86	0.11	1449971108.74
MPLight [24]	133.42	3.34	1050450518.34
IDQN [25]	80.08	1.68	944216575.78
Our	65.48	1.05	854412287.22

Table 2 shows the whole test results.

Table 2
The whole test results.

Traffic indicators	Approach	1st learning iteration	100th learning iteration
Avg Duration, sec	MaxPressure [23]	253.94	253.86
	MPLight [24]	122.95	133.42
	IDQN [25]	91.37	80.08
	Our	81.75	65.48
Avg Waiting time, sec	MaxPressure [23]	218.35	218.35
	MPLight [24]	15.01	11.18
	IDQN [25]	17.95	11.58
	Our	16.97	18.41
Total CO_2 Emissions, mg	MaxPressure [23]	143643128985.10	1449971108.74
	MPLight [24]	106248701043.38	1050450518.34
	IDQN [25]	102006837304.70	944216575.78
	Our	94901416855.32	854412287.22
Avg Queue length	MaxPressure [23]	0.11	0.11
	MPLight [24]	3.43	3.34
	IDQN [25]	2.27	1.68
	Our	2.02	1.05

The results in Table 2 show the ability of our approach to reduce CO_2 emissions by 18%. The limitations might include the individual DQN for each intersection, the lack of synchronization between intersections, and the non-transparency of decision-making in the DQN model.

Conclusions

This study presents an intelligent traffic light control approach based on DRL, specifically employing a DQN, to reduce CO2 emissions while considering the interests of all road users. The primary contribution lies in the development and simulation-based validation of a DQN model and a tailored loss function that enables adaptive traffic signal control at a local intersection. Through extensive simulation experiments using realistic traffic patterns, the proposed model achieved a measurable reduction in CO2 emissions by approximately 18%, highlighting the potential of DRL in addressing environmental and operational challenges in urban traffic management. The limitations include the need to train a separate DQN model for each intersection, the absence of coordination between adjacent intersections affecting overall traffic flow, and the lack of interpretability in the model's decision-making process, which hinders transparency, trust, and practical deployment.

Future work will focus on addressing these challenges by incorporating real-time input data from surveillance cameras and other sensors to enable online learning and context-aware signal control.

Acknowledgements

This research was supported by the U_CAN –"Towards carbon neutrality of Ukrainian cities" project, under Grant Agreement No. 101148374, funded by the Horizon Europe program.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] General Assembly United Nations, Progress towards the Sustainable Development Goals. Report of the Secretary-General, A/79/79-E/2024/54, General Assembly. Economic and Social Council, New York City, NY, USA, 2024. URL: https://unstats.un.org/sdgs/files/report/2024/SG-SDG-Progress-Report-2024-advanced-unedited-version.pdf.

[2] O. Ryzhanskyi, E. Manziuk, O. Barmak, I. Krak, N. Bacanin, An approach to optimizing CO2 emissions in traffic control via reinforcement learning, in: Proc. of the 5th International Workshop on Intelligent Information Technologies & Systems of Information Security with CEUR-WS, CEUR-WS.org, Aachen, 2024, pp. 137–155. URL: https://ceur-ws.org/Vol-3675/paper10.pdf.

[3] V. Pavlyshyn, E. Manziuk, O. Barmak, I. Krak, R. Damasevicius, Modeling environment intelligent transport system for eco-friendly urban mobility, in: Intelligent Information Technologies \& Systems of Information Security 2024, CEUR-WS.org, Khmelnytskyi, 2024, pp. 118–136. URL: https://ceur-ws.org/Vol-3675/paper9.pdf.

[4] I. Krak, V. Kuznetsov, S. Kondratiuk, L. Azarova, O. Barmak, P. Padiuk, Analysis of deep learning methods in adaptation to the small data problem solving, in: S. Babichev, V. Lytvynenko (Eds.), Lecture Notes in Data Engineering, Computational Intelligence, and Decision Making, Vol. 149: LNDECT, Springer International Publishing, Cham, 2023, pp. 333–352. doi:10.1007/978-3-031-16203-9_20.

[5] A. Haydari, Y. Yilmaz, Deep reinforcement learning for intelligent transportation systems: A survey, IEEE Trans. Intell. Transp. Syst. 23.1 (2020) 11–32. doi:10.1109/tits.2020.3008612.

[6] C. Louw, L. Labuschagne, T. Woodley, A comparison of reinforcement learning agents applied to traffic signal optimisation, SUMO Conf. Proc. 3 (2022) 15–43. doi:10.52825/scp.v3i.116.

[7] B. Setiyono, D. R. Sulistyaningrum, S. Soetrisno, D. W. Wicaksono, Multi vehicle speed detection using euclidean distance based on video processing, Int. J. Comput. 18.4 (2019) 431–442. doi:10.47839/ijc.18.4.1613.

[8] P. A. Lopez, M. Behrisch, L. Bieker-Walz, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, L. Lücken, J. Rummel, P. Wagner, E. Wiessner, Microscopic traffic simulation using SUMO, in: 2018 21st International Conference on Intelligent Transportation Systems (ITSC), IEEE, New York, NY, USA, 2025, pp. 2575–2582. doi:10.1109/ITSC.2018.8569938.

[9] T. D. Chala, L. T. Kóczy, Intelligent fuzzy traffic signal control system for complex intersections using fuzzy rule base reduction, Symmetry 16.9 (2024) 1177. doi:10.3390/sym16091177.

[10] A. Wang, K. Zhang, M. Li, J. Shao, S. Li, Game theory-based signal control considering both pedestrians and vehicles in connected environment, Sensors 23.23 (2023) 9438. doi:10.3390/s23239438.

[11] N. Faqir, C. Loqman, J. Boumhidi, Combined extreme learning machine and max pressure algorithms for traffic signal control, Intell. Syst. With Appl. 19 (2023) 200255. doi:10.1016/j.iswa.2023.200255.

[12] B. Großwindhager, A. Rupp, M. Tappler, M. Tranninger, S. Weiser, B. K. Aichernig, C. A. Boano, M. Horn, G. Kubin, S. Mangard, et al., Dependable internet of things for networked cars, Int. J. Comput. 16.4 (2017) 226–237. doi:10.47839/ijc.16.4.911.

[13] I. William, S. Kozhevnikov, M. Sontheimer, S. Chou, Enhancing urban traffic management in Taipei: A reinforcement learning approach, in: 2024 Smart City Symposium Prague (SCSP), IEEE, New York, NY, USA, 2024, pp. 1–6. doi:10.1109/scsp61506.2024.10552687.

[14] Z. Wang, L. Xu, J. Ma, Carbon dioxide emission reduction-oriented optimal control of traffic signals in mixed traffic flow based on deep reinforcement learning, Sustainability 15.24 (2023) 16564. doi:10.3390/su152416564.
[15] M. Schumacher, C. M. Adriano, H. Giese, Challenges in reward design for reinforcement learning-based traffic signal control: An investigation using a CO2 emission objective, SUMO Conf. Proc. 4 (2023) 131–151. doi:10.52825/scp.v4i.222.
[16] P. C. Sherimon, V. Sherimon, J. Joy, A. M. Kuruvilla, G. Arundas, Efficient deep learning methods for detecting road accidents by analyzing traffic accident images, Int. J. Comput. 23.3 (2024) 440–449. doi:10.47839/ijc.23.3.3664.
[17] A. Sachenko, V. Kochan, V. Turchenko, Intelligent distributed sensor network, in: IMTC/98 Conference Proceedings. IEEE Instrumentation and Measurement Technology Conference. Where Instrumentation is Going (Cat. No.98CH36222), IEEE, New York, NY, USA, 1998, pp. 60–66. doi:10.1109/IMTC.1998.679663.
[18] H. Abohashima, A. Eltawil, M. Gheith, A mathematical programming model and a firefly-based heuristic for real-time traffic signal scheduling with physical constraints, IEEE Access 9 (2021) 128314–128327. doi:10.1109/access.2021.3112600.
[19] E. Manziuk, O. Barmak, I. Krak, O. Mazurets, T. Skrypnyk, Formal model of trustworthy artificial intelligence based on standardization, in: Proc. of the 2nd International Workshop on Intelligent Information Technologies & Systems of Information Security with CEUR-WS, CEUR-WS.org, Aachen, 2021, pp. 190–197. URL: https://ceur-ws.org/Vol-2853/short18.pdf.
[20] T. A. Haddad, D. Hedjazi, S. Aouag, A new deep reinforcement learning-based adaptive traffic light control approach for isolated intersection, in: 2022 5th International Symposium on Informatics and its Applications (ISIA), IEEE, New York, NY, USA, 2022, pp. 1–6. doi:10.1109/isia55826.2022.9993598.
[21] C. Ounoughi, G. Touibi, S. B. Yahia, EcoLight: Cco-friendly traffic signal control driven by urban noise prediction, in: Database and Expert Systems Applications. DEXA 2022, Vol. 13426: Lecture Notes in Computer Science, Springer International Publishing, Cham, 2022, pp. 205–219. doi:10.1007/978-3-031-12423-5_16.
[22] I. Sutskever, O. Vinyals, Q. V. Le, Sequence to sequence learning with neural networks, in: Proceedings of the 28th international conference on neural information processing systems - volume 2, MIT Press, Cambridge, MA, USA, 2014, pp. 3104–3112. doi:10.5555/2969033.2969173.
[23] C. Chen, H. Wei, N. Xu, G. Zheng, M. Yang, Y. Xiong, K. Xu, Z. Li, Toward a thousand lights: Decentralized deep reinforcement learning for large-scale traffic signal control, Proc. AAAI Conf. Artif. Intell. 34.04 (2020) 3414–3421. doi:10.1609/aaai.v34i04.5744.
[24] J. Ault, J. Hanna, G. Sharon, Learning an interpretable traffic signal control policy, in: Proc. of the 19th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS ‘20), International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, USA, 2020, pp. 88–96. URL: https://ifaamas.org/Proc./aamas2020/pdfs/p88.pdf.
[25] S. Vladov, L. Scislo, V. Sokurenko, O. Muzychuk, V. Vysotska, S. Osadchy, A. Sachenko, Neural network signal integration from thermogas-dynamic parameter sensors for helicopters turboshaft engines at flight operation conditions, Sensors 24.13 (2024) 4246. doi:10.3390/s24134246.
[26] P. Radiuk, O. Barmak, E. Manziuk, I. Krak, Explainable deep learning: A visual analytics approach with transition matrices, Mathematics 12.7 (2024) 1024. doi:10.3390/math12071024.
[27] J. Ault, G. Sharon, Reinforcement learning benchmarks for traffic signal control, in: Proc. of the Thirty-fifth Conference on Neural Information Processing Systems (NeurIPS 2021) Datasets and Benchmarks Track, 2021, pp. 1–11. URL: https://people.engr.tamu.edu/guni/papers/NeurIPS-signals.pdf.
[28] V. Mnih, K. Kavukcuoglu, D. Silver, E. Al, Human-level control through deep reinforcement learning, Nature 518.7540 (2015) 529–533. doi:10.1038/nature14236.
[29] P. Alvarez Lopez, A. Banse, M. Barthauer, M. Behrisch, B. Couéraud, J. Erdmann, Y.-P. Flötteröd, R. Hilbrich, R. Nippold, P. Wagner, Simulation of urban mobility (SUMO), Software, v. 1.21.0, Zenodo, 2024. doi:10.5281/zenodo.13907886.

[30] D. Krajzewicz, M. Behrisch, P. Wagner, R. Luz, M. Krumnow, Second generation of pollutant emission models for SUMO, in: Modeling Mobility with Open Data, Lecture Notes in Mobility, Springer International Publishing, Cham, 2015, pp. 203–221. doi:10.1007/978-3-319-15024-6_12.