Automated intelligent-agent optimisation of per-lane variable speed limits

Kandiri, Amirreza; Nogal, Maria; Martinez-Pastor, Beatriz; Teixeira, Rui

DOI
10.1016/j.asoc.2025.113554

Publication date
2025

Document Version
Final published version

Published in
Applied Soft Computing

Applied Soft Computing Journal 181 (2025) 113554

Contents lists available at ScienceDirect

Applied Soft Computing

journal homepage: www.elsevier.com/locate/asoc

Automated intelligent-agent optimisation of per-lane variable speed limits

Amirreza Kandiri [a,*], Maria Nogal [b], Beatriz Martinez-Pastor [a], Rui Teixeira [c]

[a] School of Civil Engineering, University College Dublin, Ireland
[b] Faculty of Civil Engineering & Geosciences, Delft University of Technology, Netherlands
[c] Department of Civil, Structural and Environmental Engineering, Trinity College Dublin, Ireland

HIGHLIGHTS

- A novel two-stage method is introduced for variable speed limit control.
- The estimator can predict the travel time by almost 97 % accuracy.
- The proposed method reduces the mean travel time by up to 55 %.
- It reduces the mean waiting time by up to 69 %.
- It finds the near-optimum solution in almost 10 s.

ARTICLE INFO

Keywords:
Intelligent transportation systems
Traffic management
Intelligent agents
Variable speed limits
Optimisation
Travel time

ABSTRACT

Recent advancements in intelligent transportation systems and data analytics within transportation systems present a significant opportunity to enhance operational efficiency. In this context, the pivotal role of intelligent agents in achieving real-time optimisation for traffic management is highlighted. Such agents can predict and decide autonomously and can be trained to understand the underlying complexities of the traffic in real-time. In this paper, an innovative framework to perform real-time traffic optimal management decisions is proposed. Its rationale uses a fusion of data observations and simulation to enable an autonomous agent capable of accurate adaptive traffic management. A Case Study of application is developed using the M50 motorway in Dublin, where the speed limits are applied as adaptive parameters for optimal traffic management. Results show that the intelligent agent can autonomously predict travel times and decide in real-time the optimal speed limits to impose on a motorway when signs of congestion are found. The agent can reduce the mean travel time of a time interval by up to 55 % and the mean waiting time by up to 69 % in a situation of congestion. The average travel times of the studied M50 junction have significantly improved, showing the potential of autonomous agents in enhancing real-time optimal traffic management.

1. Introduction

Traffic congestion is one of the most critical issues in current metropolitan areas, where it has been aggravated by the continuously increasing number of road vehicles [1]. Traffic congestion causes an increase in pollution, losses of energy and time, and increases stress and anxiety in drivers [2,3]. Moreover, congested traffic negatively affects both public and private services and logistics [4]. Congestion control is an undeniable need, and common actions to improve traffic efficiency include expanding transport infrastructure, creating incentives to use public transportation, or imposing limits or policies to discourage the use of private vehicles. Traffic performance can also be improved by adapting the traffic network to known operational conditions. To do so, dynamic traffic management approaches that use real-time travel information systems [5], dynamic routing [6,7], and ramp metering [8,9] have been used before. Among these, Variable Speed Limit (VSL) control [10,11] is one of the promising strategies that use real-time traffic information systems to control the mainstream traffic on freeways, increase the operational capacity of the traffic network, and reduce its congestion [12,13].

The literature on VSL can be broadly divided into two main classes: (a) Continuous Models, where time and road are treated as continuous

* Corresponding author.
E-mail address: Amirreza.Kandiri@ucdconnect.ie (A. Kandiri).

https://doi.org/10.1016/j.asoc.2025.113554
Received 17 January 2025; Received in revised form 26 May 2025; Accepted 26 June 2025
Available online 5 July 2025

variables with corresponding continuous values assigned to them [14,15]. For a comprehensive review of traffic controllers, readers can refer to [16]. (b) Discrete Models, where time and road are represented with discrete values. Among discrete models, Reinforcement Learning (RL) and Model Predictive Control (MPC) are the most commonly employed techniques [17].

MPC methods integrate predicted data with real-time information about the system's current state to guide speed controllers throughout the traffic network. These methods aim to achieve specific objectives, such as minimising travel time [18] or improving traffic flow [19,20]. They often rely on macroscopic traffic models like METANET [21] or the Cell Transmission Model (CTM) [22,23] to estimate traffic performance metrics such as travel time. However, conventional traffic models are not yet capable of accounting for per-lane VSL in their computations, making simulation a necessary alternative. Additionally, MPC is typically paired with optimisation algorithms to achieve a global optimum for the selected traffic performance metric. However, incorporating simulation as a predictive component within an optimisation algorithm significantly increases computational complexity, rendering it impractical for real-time traffic management applications.

In RL-based approaches, an autonomous agent learns through interaction with the environment. The agent changes the environment, observes their outcomes, and subsequently decides to control traffic. This interaction, if trained in a real environment, can disturb the traffic in the training phase. The main goal of an RL-based agent is to make decisions that maximise pre-defined rewards, such as a reduction in travel time [24] or carbon dioxide emissions [25]. One of the main drawbacks of RL-based methods is related to rewards. In traffic, since the travel time is not to be obtained until the vehicles arrive at their destination, reinforcement learning approaches need to deal with delayed rewards, which presents a significant operational challenge [26,27]. These make the training of the agent a long and computationally intensive process [28,29]. Alternative rewards, such as the current vehicles' speed [30,31] have been used before to overcome this issue. Reinforcement learning also iteratively searches for improvement in traffic operation conditions. However, it cannot guarantee optimum solutions. A straightforward example of this would be the case where the agent gets stuck in a local optimum regarding system performance and where changes to system parameters may not produce a reward and incentivise it to explore other potentially more efficient operational conditions further.

The main limitations of RL include the potential to disrupt traffic during the training phase when trained in real-world environments, its inability to consistently achieve the global optimum, and the challenges associated with handling delayed rewards. In contrast, MPC does not suffer from these issues; however, it requires an accurate predictive model with low computational overhead to be viable for real-time implementation. These limitations highlight the need for integrated approaches that effectively balance predictive accuracy and real-time optimisation to address the gaps in current methods.

In the present work, by combining the ideas of MPC and RL, an Intelligent Agent (IA) called Adaptive Speed Limit Controller (ASLC) is proposed. It benefits from both methodologies' advantages while addressing their main drawbacks. The ASLC adopts the MPC approach in prediction as an integral part of real-time optimal decisions through the IA. The IA is presented as an agnostic agent that is expected to function with any data-driven predictive method (DDPM) if the surrogate traffic function is compatible with the DDPM assumptions (e.g., complexity or smoothness of the function to be approximated). Hence, a DDPM such as a Support Vector Machine (SVM), a decision tree, or an Artificial Neural Network (ANN) is expected to be equally functional with appropriate training. The present implementation uses an ANN as an example of how the proposed ASLC performs. Ensuring the prediction accuracy of the ANN depends on its structure, including the number of hidden layers and neurons in each layer. To optimise the ANN's structure and enhance prediction accuracy, various optimisation algorithms such as nature-inspired randomised metaheuristics like Particle Swarm Optimisation (PSO) [32], Gray Wolf Optimisation (GWO) [33], Salp Swarm Algorithm (SSA) [34], or others can be employed during training. These algorithms help the ANN understand the complex traffic dynamics in practical operational scenarios. In the present work, this part of ASLC is referred to as the Estimator.

Subsequently, the ASLC employs interaction with the environment to infer potential changes to it. To mitigate issues of learning complexity and potential disruptions to traffic performance while it has not yet learned how to manage it, the ASLC employs a strategy that integrates information from the estimator and high-fidelity simulations instead of direct environmental interaction. This part of the ASLC is referred to as the decision-maker. It excels in exploring unknown and unrecorded traffic states through learning high-fidelity simulations. Additionally, it uses another optimisation algorithm (in this paper, SSA is applied) to find the optimum speed limits in real-time while accounting even for operational conditions that have not been observed in the system yet. It is noted that through the ASLC estimator (as the estimator is trained to know all it can about the system being handled), it can obtain results of each potential speed limit variation in virtually zero time. Hence, this exploration of the traffic speed configuration that generates the minimum traffic cost, i.e., the optimum, can take place in real-time, which is referred to here as real-time optimum or real-time optimality.

The proposed methodology's large computational cost resides in the estimator's definition and the evaluation of high-fidelity traffic simulations. In these, the applied DDPM (in this case, an ANN) learns all the functionality of the traffic operation, fusing real data and simulation while offline and when there is no need for VSL. The real data imparts information on recorded traffic performance, and the simulations allow it to learn the functionality of the traffic conditions and the trends of its traffic performance function in a large spectrum of possible scenarios of operation that have not been recorded in the system. As remarked, this cost can occur offline, before real-time management, or concurrently as part of ongoing IA enhancement (*e.g.*, enriching the estimator with further recorded data).

It is also noted that most studies in the literature consider a single speed limit for all lanes [29,35,36], meaning the speed limit changes along the length of the road but remains consistent across all lanes. However, the proposed IA in the current work can find optimum speed limits for each lane individually. While issues with individual lane speeds may be raised, such as safety, the results show that the IA can be highly adaptive in decision-making and work with as many degrees of freedom or decision variables (DVs) as possible. This is mainly due to the capability of the IA to learn all operational scenarios allowed in the network offline. Practicality issues of multi-lane adaptation are also discussed, but it is noted that the main purpose of the presented framework is to remark on how agents can enhance real-time traffic management by achieving real-time optimality, potentially revolutionising current practices in this field. Existing literature explores various vehicle types, from human-driving vehicles (HDVs) to fully connected autonomous vehicles (CAVs). Nonetheless, studies show that increasing the penetration rate of CAVs can benefit the implementation of VSL [37]. On the other hand, enforcing the speed limit using CAVs is easier due to recent technologies such as intelligent speed adaptation [38], and this has synergies with the data-driven management approach of the ASLC. In this study, 100 % HDVs are used.

The main contribution of the present work is to propose a methodology that optimises the per-lane VSL in real-time by using data collected from loop detectors without disrupting traffic, while aiming to attain the global optimal decision in variable speeds that can be achieved under a set of functional traffic conditions. For this purpose, the present work is organised as follows: Section 2 presents the idea of real-time optimality; Section 3 thoroughly explains the approach to the ASLC; Section 4 demonstrates the ASLC in a case study that uses SUMO; Section 5 presents and discusses the results. Finally, Section 6 outlines the main conclusions and identifies opportunities for further exploration

in applying IA in traffic management.

2. Problem definition: real-time optimality with intelligent agents

The current implementation tackles the problem of optimising traffic flow within a multi-lane motorway, where traffic moves in a single direction. The goal is to determine the optimum speed limit for each lane that minimises the mean travel time in discrete time-intervals and to do this in so-called real-time. There are two key aspects to this idea of finding optimality in real-time that are remarked in this work: I) Finding an optimum or quasi-optimum solution to a decision-making problem that involves speed limits in traffic, and II) doing it in virtually zero-time [39].

Optimality refers to achieving the best solution for an objective function, with reference to DVs, under a set of constraints. The most effective way to achieve it consistently is to use an optimisation algorithm. The most elementary way to achieve an optimum solution is to use mathematical programming. Nonetheless, in many engineering problems, randomised optimisation algorithms [40] have provided successful results when compared to other optimisation techniques [27–29]. Implementing an optimisation algorithm requires defining an objective function. The objective of the ASLC is to minimise the mean travel time. Fig. 1 shows how the proposed IA addresses, in real-time, the goal of achieving optimal decisions with respect to traffic operation.

The IA method involves two steps to evaluate the average travel time for a given time-interval. First, it uses high-fidelity data (e.g., real or calibrated high-fidelity SUMO data). Second, it uses a low-fidelity model (i.e., DDPM through an ANN in this study). While high-fidelity models are anticipated to be more accurate, they typically require extensive analysis efforts, rendering them impractical for real-time optimisation. In fact, optimisation using high-fidelity models is usually characterised by its high demand in terms of analysis time [39]. On the other hand, lower-fidelity data obtained in low-fidelity models can also achieve adequate performance in accuracy (when approximating the larger fidelities) in significantly less time if defined or trained and validated with adequate high-fidelity data [41].

As presented in Fig. 2, the fundamental assumption of the IA definition is that a lower-fidelity model can be used to solve a problem of optimisation in real-time, since its analysis cost is negligible. Nonetheless, several questions remain, such as how to define a highly accurate low-fidelity representation for complex issues like traffic and travel time, or which models are suitable to perform an optimisation that commonly involves many calls of the objective function, in real-time. Metamodels [39,42] have shown significant potential due to their ability to accurately surrogate even highly non-linear or non-smooth functions. Among these, ANNs are particularly promising because of their flexibility to incorporate numerous layers and neurons. An issue, however, is that of defining the ANN in a format that has enhanced accuracy. In the present implementation, an optimisation algorithm is applied to find the structure that more accurately surrogates the mean travel time, serving as the IA's objective function. In practice, this is an unsupervised definition of the estimator.

To find an optimum or quasi-optimum solution for the DVs (i.e., speed limits), the optimisation algorithm can now call the travel time estimator several times at a negligible cost (less than 1 s). Such an approach would be impractical with high-fidelity models due to their characteristic larger analysis times, even with parallel or other advanced processing techniques.

The ASLC compiles information on mean travel time as the performance measure that can be the basis of an optimisation scheme of the VSL. It essentially consists of an ANN with two main types of input variables: Prediction Support (PS) and DV. The main rationale is that, for a given traffic system state, speed limits act as the DVs. All other data, such as information from detectors, serves as supportive input to understand system behaviour, i.e., PS, even though this data itself cannot be adjusted.

It is noted that this approach assumes that a High-Fidelity simulation model is an appropriate representation of High-Fidelity real operational data. This assumption is commonly used in traffic management, such as agent-based simulation models. As a result, a high-fidelity simulator can effectively approximate potential operational scenarios not previously recorded or implemented. Hence, the ASLC can test and learn the inherent functionalities of available decisions offline. As shown in Fig. 1, it learns the impact of different decisions (speed limits) when information about input traffic conditions is available and calibrates these as information becomes available.

3. ASLC methodology

In this work, the IA is developed as a multi-lane variable speed-limit controller to individually set the optimum speed limit for each lane in a motorway. Fig. 3 illustrates how ASLC and its various components function and interact with the motorway traffic systems. The ASLC comprises two modules: I) ASLC estimator and II) ASLC decision-maker, both utilising a DDPM, with ANN being applied in this study and an optimisation algorithm (specifically SSA in this study). It is noted that the SSA is applied to optimise the network and to find the optimal decisions in the operational phase, but the methodology proposed is not exclusive to it. Its main rationale is that offline-online training of independent decision-makers can lead to real-time optimality in motorways, and for that, real-time optimisation should take place. To achieve this, data is fused offline to explore system functionality, and the decision-making process relies on a minimum response time decision scheme.

3.1. Definition of the ASLC Estimator

The main goal of the estimator is to predict the mean travel time for each potential speed limit combination. It is also noted that the estimator is defined to be a global approximator of the traffic performance,

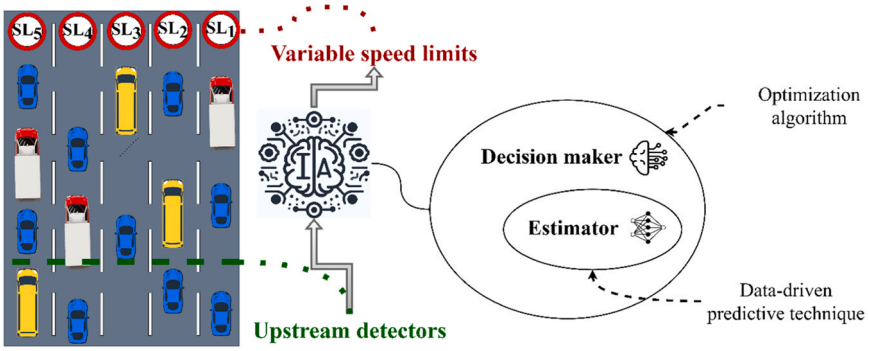

Fig. 1. Multi-lane VSL controlling in real-time.

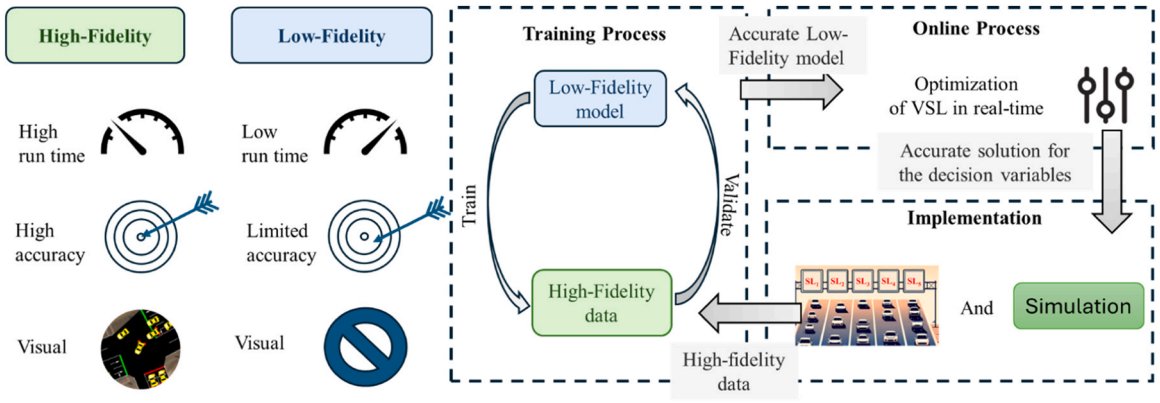

Fig. 2. High and low-fidelity models' properties.

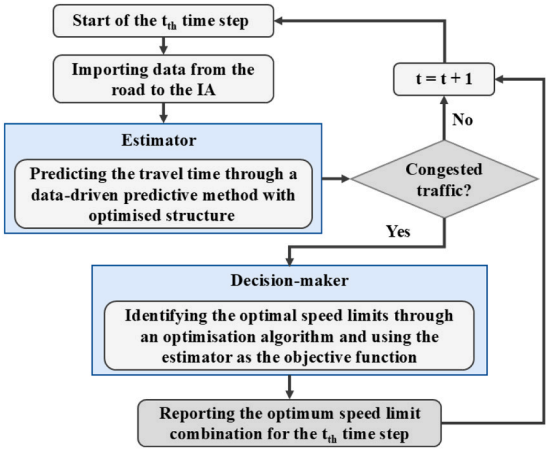

Fig. 3. The ASLC's overall workflow.

and this is key for the performance of the decision-maker. In the literature, various machine-learning approaches such as neural networks [43,44], support vector regression [45], random forest [46], and the gradient boosting regression tree method [47], among others, have been employed to predict travel time. Interested readers are referred to Oh et al. [48] and Bai et al. [49] for further reference works. In this study, an ANN is used as the estimator. Detailed explanations on ANN are presented in Appendix A.

However, an ANN's accuracy relies on the number of its hidden layers and the number of nodes in each one of these layers [50]. A small change in these properties can make a remarkable difference in the ANN's accuracy. Common practice uses trial-and-error to fit the ANN, which can be acceptable for simple problems. In complex problems such as the ones that involve traffic and multiple speed limits, an automated solution is required. If three potential hidden layers and 20 potential nodes for each of those layers are considered, the total search area for ANN combinations considers a maximum of 8000 combinations. In such circumstances, trial and error is effort-consuming if the goal is to find the best combination for the ANN's structure. Thus, an optimisation algorithm is used to find the optimum ANN structure in an automatic form and without supervision. Hence, the estimator is a hybrid machine-learning technique consisting of a nature-based swarm intelligence-based optimisation algorithm and ANN (OANN). Detailed explanations on OANN and how it optimises the ANN's structure are presented in Appendix B.

Moreover, to present the accuracy of the predictions, mean absolute percentage error (MAPE), mean absolute error (MAE), Pearson correlation coefficient (R), and RMSE are also applied.

$$R = \frac{T\sum_{j=1}^{T} A_j E_j}{\left(T\sum_{j=1}^{T} A_j^2 - \left(\sum_{j=1}^{T} A_j\right)^2\right)\left(T\sum_{j=1}^{T} E_j^2 - \left(\sum_{j=1}^{T} E_j\right)^2\right)}, \quad (1)$$

$$RMSE = \sqrt{\frac{1}{T}\sum_{j=1}^{T}(A_j - E_j)^2}, \quad (2)$$

$$MAE = \frac{1}{T}\sum_{j=1}^{T}|A_j - E_j|, \quad (3)$$

$$MAPE = \frac{1}{T}\sum_{j=1}^{T}\left|\frac{A_j - E_j}{A_j}\right|, \quad (4)$$

where E_j, A_j, and T are the ANN's estimated output, actual values of the output, and the number of data records in each dataset, respectively. After obtaining the objective values for each individual, the fittest individual is identified by comparing their objective function against those of the food source. If it has a superior objective function value, it replaces the food source and is used as the new food source for the next iteration. Subsequently, the new positions of the population are updated accordingly. This process of finding individuals with the best fitness function continues until termination conditions are met, which in this study is the maximum number of iterations. Finally, that individual indicates the architecture of the estimator.

3.2. ASLC decision-maker

The task of the decision-maker is to find the optimum speed limits for each lane individually for a specified time segment when necessary (see Fig. 3). At the beginning of each time step, the ASLC uses the estimator to assess whether traffic congestion will occur. Congestion is identified if the predicted mean travel time exceeds a pre-defined threshold. When congestion is detected, the decision-maker initiates the optimisation process to optimise speed limits and minimise mean travel time. To have a better understanding of the function of ASLC, the input vector (X) is divided into two subsets: PSs containing data that are collected from the traffic and are not changeable, and DVs, which are being changed by the ASLC decision-maker to minimise the mean travel time i.e., individual speed limits of each lane. Here, each individual's position consists of the DV set. This optimisation takes place in the decision-making variables of the ANN and the decision-maker's optimisation process is illustrated in Fig. 4.

After setting up the initial values, the first generation of the population, comprising speed limit combinations, is randomly generated. Using the estimator, each individual's objective function (i.e., mean travel time) is calculated. The fittest individual is compared to the food source and replaces it if it has a better objective function (lower travel time). Then, the next generation of the population is generated, and this entire process is repeated until the last iteration is reached. Finally, the

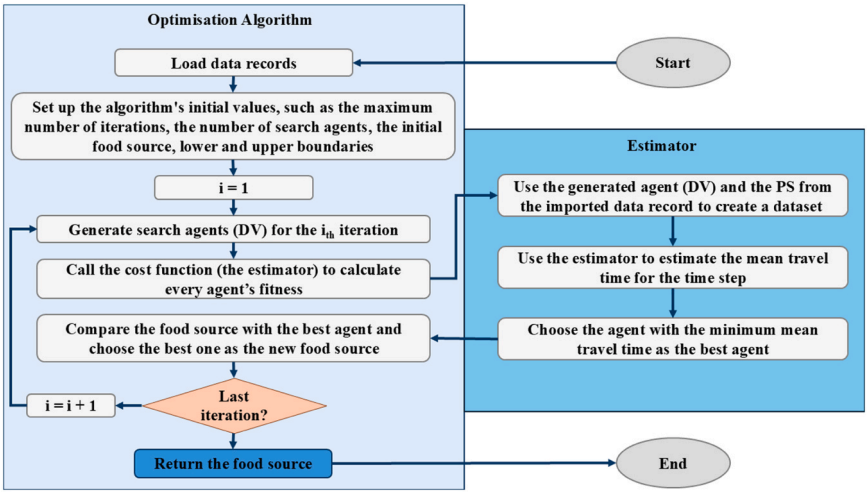

Fig. 4. Decision-maker's optimisation process.

food source will be returned as the optimum solution for travel speed limits.

4. Case study

The ASLC is demonstrated in a simulation study conducted on the M50 motorway in Dublin. The M50 is a C-shaped motorway linking the north and south of Dublin, encircling the city (see Fig. 5). As the busiest motorway in Ireland, it faces significant congestion issues [51]. According to existing data, 45.1 million passages were recorded using the M50 motorway in 2021 [52].

This motorway has been extensively studied in traffic research. For example, Corbally et al. [53] utilised machine learning techniques to estimate the duration of incidents on the M50. Similarly, Kandiri et al. [54] focused on travel time prediction by developing a feature selection framework, using the M50 as a case study. Rogers and Darcy [55] explored how toll booths affect traffic flow on the motorway, while Kemp and O'Mahony [56] examined delays in travel time predictions, measured their impact, and proposed a model to address these challenges. These studies highlight the critical importance of the M50 in understanding and improving traffic dynamics in Ireland. Moreover, De Paor et al. [57] introduced a rule-based approach to increase the traffic capacity of M50. A high-fidelity traffic simulation of the motorway using the Simulation of Urban MObility (SUMO) is applied in the present work. Fig. 6(a) and (b) show the segment under study on the M50. This road segment is also marked by the blue ellipse in Fig. 5. The analysed segment comprises five lanes with a standard speed limit of 100 km/h. In this study, only one direction of M50 (i.e., southbound) is considered.

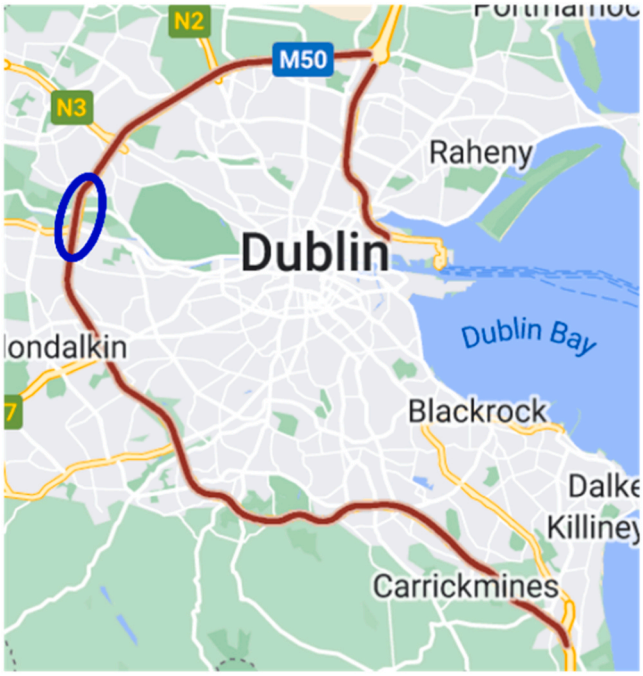

Fig. 5. M50 motorway. Source: Google Maps.

4.1. Calibration and validation of the simulation

The model of M50 generated in SUMO that is used in this study was developed by Gueriau and Dusparic [58]. The model includes HDV and CAV with automation levels 2 and 4, which means automation with the need for human supervision and fully automated driving, respectively. The authors developed six different scenarios with different HDV and CAV ratios. In this study, scenario A is used, which has 100 % HDV and no CAVs are considered. Demand patterns and volumes were developed based on real data, averaged over daily observations from 2012 to 2019. This demand pattern is visualised in Fig. 6. (c). A total of 144 distinct traffic demand scenarios were simulated over varying time intervals. These demands ranged from approximately 100 to 2500 vehicles per hour per lane (veh/h/lane). The experiment enabled the evaluation of agent performance under three traffic conditions: (1) free flow (below 1800 veh/h/lane), where vehicles travel at maximum speed with minimal interaction; (2) saturated (between 1800 and 2200 veh/h/lane), corresponding to the point of maximum flow; and (3) congested (above 2200 veh/h/lane), where both speed and flow decrease due to high density.

Moreover, the SUMO parameters for calibrating the simulation are also adopted from the same study. These parameters are presented in Table 1.

SUMO provides a simulation framework for microscopic traffic modelling, where vehicle dynamics result from the interplay of longitudinal and lateral models. Longitudinal dynamics are governed by a car-following model that can be adjusted to replicate a range of driving styles and vehicle capabilities. Lateral dynamics are computed through a lane-changing model, typically involving decision-making processes influenced by the perception of adjacent lanes. The interested readers are referred to [58] for an extensive explanation of the model and its parameters. Fig. 7 schematically shows the under-study area and the VSL area.

To validate the simulation's results, another dataset collected in the

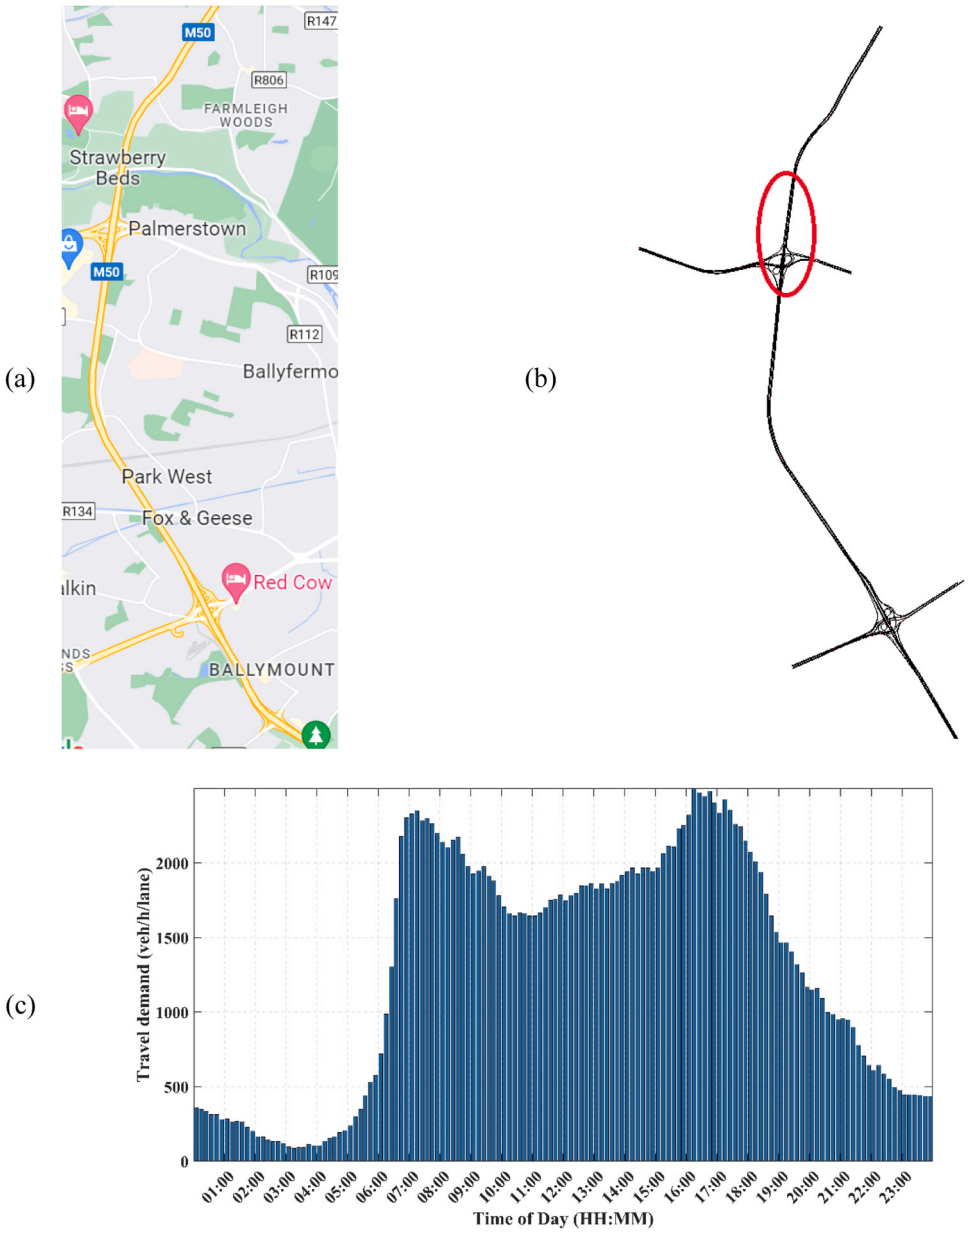

Fig. 6. Under-study part in the map (a) vs in the simulator (b), and the travel demand in 24 h (c). This area is from the toll, located between Blanchardstown and Palmerstown, to Chapelizod Bypass. Source: Google Maps.

Table 1
Sumo parameters.

Parameter	Value
Car-following model	Krauss [59]
Speed deviation (%)	0.1
Time headway (s)	1.2
Min gap (m)	2.5
Max accel. ($m/^{s2}$)	2.5
Deceleration (m/s^2)	7.5
Max decel. (m/s^2)	9
Imperfection	0.5
Lane-changing model	SUMO lane-change model [60]
Cooperation	0.5
Anticipation	0.5

summer of 2023 is used. In this process, each day of the week is classified as a distinct class to account for variations in traffic patterns. To address the limited data availability and the rapidly changing weather conditions in Dublin, days with similar weather patterns were selected to form an aggregated historical dataset for each class. Holidays were excluded to eliminate anomalies in traffic patterns. This resulted in a dataset consisting of average values of 10 days in each class. These datasets were used to define the simulations and an unseen day in each class was used to validate the simulation results. Table 2 presents the results of this validation.

The results indicate that for flow, the MAPE values remain below 5 % on all days except Thursday, which has an MAPE of 7.8 %. However, this is still within an acceptable range. It is noted that this order of error is in line with the errors found in the literature. For instance, the lowest MAPE achieved in [61] was 6.5 %. The MAPE reported in [62] for 10-minute time intervals is 8.88 %. This value ranged from 4.32 to 8.99 in different sites in [63].

Correlation coefficients for flow are consistently 0.92 or higher, demonstrating a strong alignment between simulated and observed data, particularly on Sunday. For speed, the MAPE values are slightly higher than those for flow, but remain below 10 %, indicating good

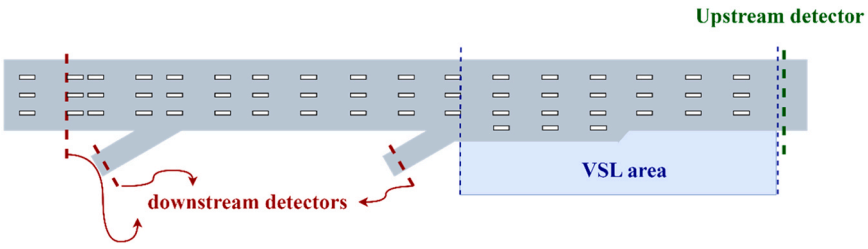

Fig. 7. The under-study segment of M50.

Table 2
Error indicators in validating the simulation.

	Flow		Speed	
	MAPE (%)	R	MAPE (%)	R
Monday	3.02	0.97	6.65	0.95
Tuesday	4.3	0.95	7.76	0.96
Wednesday	4.59	0.94	7.24	0.97
Thursday	7.8	0.92	5.56	0.98
Friday	3.65	0.97	6.7	0.95
Saturday	4.17	0.98	8.28	0.94
Sunday	3.36	0.99	6.75	0.95
mean	**4.41**	**0.96**	**6.99**	**0.96**

predictive accuracy. However, performance is slightly less precise on Saturdays. The correlation coefficients for speed consistently exceed 0.94. The average MAPE and correlation coefficient (R) values for flow are 4.41 % and 0.96, respectively, while for speed, they are 6.99 % and 0.96. These results confirm that the simulation delivers high accuracy and strong correlations overall.

4.2. Implementation of the ASLC

This section provides a detailed, step-by-step guide for implementing the proposed framework using the M50 motorway in Dublin as a case study. In the current study, the traffic time is divided into 10-minute time segments. Therefore, the ASLC assesses every 10 min, at the start of each time segment, the prediction for the current mean travel time in accordance with the knowledge it has on the network and optimises, taking action in decisions, the speed limit combinations if required. In this study, besides the collected information from detectors, i.e., the number of cars entering the road segment during the time step, the data from the past time steps, i.e., the mean travel time for the past five time segments and the number of cars that are already in the road segment controlled at the start of the present time-segment are considered to predict the state of the traffic in the next 10 min. Further insight into this methodology is provided in the remainder of this section.

The estimator is provided by the ANN with an optimised structure and aims to predict the mean travel time for each step in global operational conditions. That is, it is trained to learn all the features of traffic performance in any set of possible speed limit scenarios. To implement the estimators, the following inputs are used: (i) speed limit of each lane ($s_1,\ldots,s_5$), that is, the DV set of inputs. Considerations of the speed limits for each lane are intuitive, as for the estimator, these are key parameters that are expected to affect the mean travel time. (ii) The number of vehicles entering the network at the beginning of the time segment (EC), and (iii) the number of vehicles already present in the network at the start of the time segment (IC). Both variables were found to be influential during implementation [54], as they reflect the vehicle density at the onset of the time interval. (iv) The mean travel-time of the past five time-intervals ($MT_{ti-1},\ldots,MT_{ti-5}$). In traffic, there is a temporal correlation in the time segments. Further analysis revealed that both the current travel times and the vehicle dynamics within the controlled network segment collectively have a significant influence on performance dynamics. Therefore, to gain a better understanding of the mean travel time, the estimator and the ASLC must know the mean travel time of past time steps to predict what will happen in future time steps accurately. Hence, the mean travel time for the past five time steps is also provided. It should be noted that inputs ii to iv serve as the PS set of inputs.

The M50 motorway has five lanes in the under-study area, and the speed limits are assumed to be allowed between 10 and 40 m/s, i.e., 36 and 144 km/h. The decision space is kept large to encourage exploration and gain a better understanding of the agent's capabilities. Therefore, the search area for decisions (i.e., the degrees of freedom left for the IA to operate) is a five-dimensional space with each dimension ranging from 36 to 144 km/h. At the start of a time step the estimator can predict the travel time and if it shows a congested traffic, the decision-maker gets activated. Then, the decision-maker gets the PS set of inputs (see Section 5.2) from the traffic network. Subsequently, the optimisation algorithm in the decision-maker generates different values for the DV set of inputs (the speed limit of each lane) to find the optimum combination of speed limits under the specified conditions. As the estimator is a global approximator, it has learned the relationship between these input features and the output, the mean travel time, *a priori*. Then, after reporting the optimum speed limit combination, the same process will be repeated for the next time step (see Fig. 3) whenever required.

4.3. Experimental setting

The travel demand dataset, averaged over daily observations from 2012 to 2019, is first imported into the calibrated SUMO to define the vehicles. This dataset comprises 144 data records ranging from 180 vehicles per time interval to 4362 vehicles per time interval, each representing a 10-minute time interval. Then, by changing the per-lane VSL from 36 to 144 km/h, a comprehensive dataset including 142,336 data records is produced. This dataset is then randomised, and 2000 generated records are used to train the ASLC. The rest of this dataset is used as an unseen dataset to validate the ASLC's performance (see Section 5). A small percentage of the dataset is used for training to avoid high computational costs and overfitting of the ANN. Fig. 8 illustrates the histogram of the comprehensive dataset for EC, IC, and mean travel time. As shown in Fig. 8(c), a peak is observed at 200 s in the mean travel time. This is due to the fact that, during a 24-hour day, a large period of time has free-flow traffic, e.g., between midnight and 5 am.

5. Results and discussion

The ASLC was applied in a 24-hour traffic scenario used as a reference case study, and the results are presented in this Section. First, the ASLC performance as a global approximator of this junction is discussed, and then its adaptive decision-making capabilities are analysed.

5.1. ASLC as a global approximator of traffic performance

The following set of initial parameters is applied to train the ASLC estimator, see Table 3. The values of the maximum number of iterations, number of search agents, maximum number of hidden layers, and

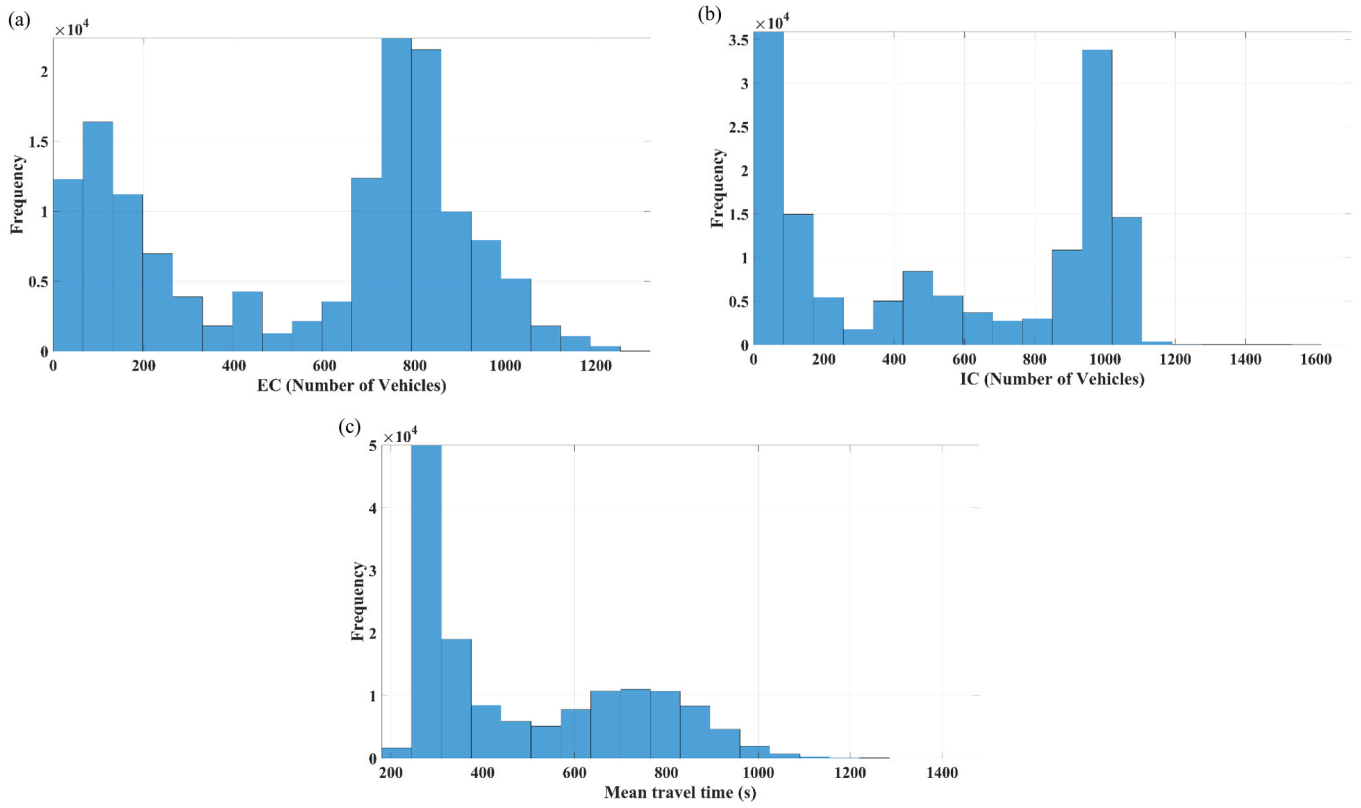

Fig. 8. The histogram of a) EC, b) IC, and c) mean travel time.

Table 3
The estimator's initial parameters.

Parameter	Value
Maximum number of iterations	100
The number of search agents in each iteration	30
Maximum number of hidden layers	3
Maximum number of neurons in each hidden layer	20
Activation function in hidden layers	Hyperbolic tangent sigmoid
ANN training algorithm	Levenberg-Marquardt
Cost function	RMSE

maximum number of neurons in each layer are selected based on the computational budget. In Section 5.3, the robustness of these choices is discussed. The *tanh* function is applied due to its nonlinear nature and ability to maintain a controlled, consistent learning process, which is less prone to issues such as vanishing or exploding gradients or erratic updates. [64,65]. The Levenberg–Marquardt algorithm is well-suited for traffic prediction due to its ability to model complex, non-linear patterns with high accuracy and fast convergence. It efficiently adjusts network parameters, handles irregular traffic data when combined with pre-processing, and consistently delivers reliable, low-error forecasts [66, 67].

Throughout the training phase, a total number of 3000 individual search agents, go through a search area of 8000 different combinations to achieve the final ANN (see Fig. 9). The goal of this stage of the approach is to find a best-known optimum structure for the ANN that has the lowest error among the design of experiments considered for the ANN structure. The proposed method finds the ANN with the lowest error in each iteration, compares it with the lowest reached error in previous iterations, and saves the structure with the global lowest error. At the end of each iteration, only one ANN, i.e. the most accurate one, is kept as the best known solution by the algorithm. The optimum structure has 12 neurons in the input layer (one neuron for each input), two

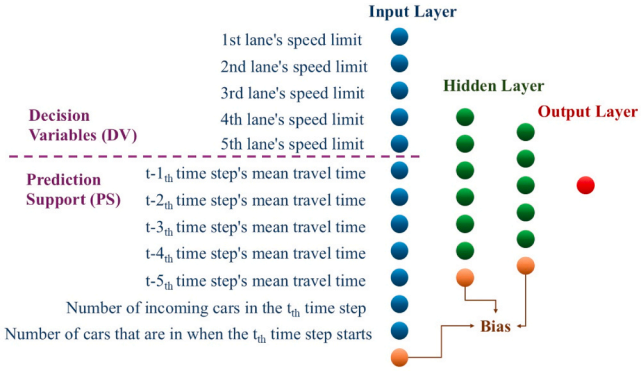

Fig. 9. The estimator's structure.

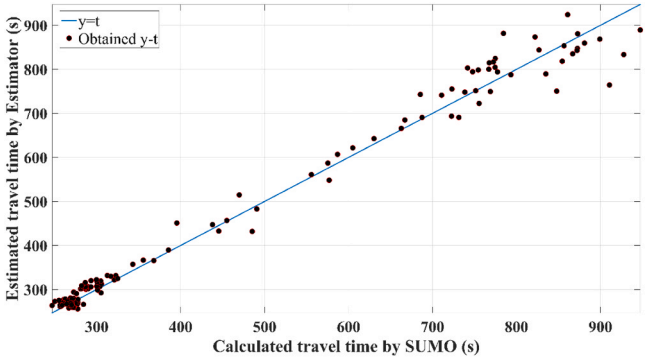

Fig. 10. Estimator's cross-validation.

hidden layers with six and five neurons in the first and the second layers, respectively, and one neuron in the output layer (for the output). Moreover, to validate the ASLC estimator, an unseen dataset is given to it. Fig. 10 shows the cross-validation of the estimator on an unseen dataset, and Table 4 represents the error indexes. As can be seen, the results have an R-value of 98 %, which shows the large correlation between the estimator's and the simulation's outputs. Furthermore, the RMSE is 34.6 s and the MAPE is approximately 3.3 %. These error indexes indicate accurate performance in relation to the travel times considered, noting that this is the estimator of global traffic performance in the section studied (i.e., with combinations of variable speeds in a 5-dimensional space). The training scenarios have a wide range of domain values from almost free-flow traffic, *i.e.,* time steps with low mean travel times, to congested traffic, *i.e.,* time steps with large mean travel times. This wide range is potentially a challenge for training an accurate ANN. Moreover, the effect of having different speed limits for each lane should be taken into consideration, since it makes the relation between inputs and outputs more challenging to predict. The mathematical equation of the final ANN is as follows:

$$y = F_3(LW_2(F_2(LW_1 \times (F_1(IW \times X + b_1 \quad)) + b_2)) + b_3), \tag{5}$$

$$X = DV \cup PS, \tag{6}$$

$$DV = \{s_1, s_2, s_3, s_4, s_5\}, \tag{7}$$

$$PS = \{EC, IC, MT_{ti-1}, \ldots, MT_{ti-5}\}, \tag{8}$$

where y is the mean travel time, F_i is the activation function in the *i-th* layer, *X* is the input vector, *IW* is the input layer weights, LW_1 and LW_2 are the first and the second hidden layer weights. b_1, b_2, and b_3 are the input layer, the first hidden layer, and the second hidden layer's biases, respectively. IW_1, IW_2, b_1, b_2, and b_3 are presented in Appendix D.

5.2. Application of the pair estimator – decision-maker

After training the ASLC on this case study, the IA is applied on a 24-hour simulation to research how it can control the traffic flow. The optimisation problem solved by the ASLC decision-maker is as follows:

Minimise

$$y = F_3(LW_2(F_2(LW_1 \times (F_1(IW \times X + b_1 \quad)) + b_2)) + b_3), \tag{9}$$

Subject to:

$$36 \leq s_1, \ldots, s_5 \leq 144 \tag{10}$$

Fig. 11 illustrates a comparison between a) the travel time of each vehicle and b) the mean travel time in each time step in the standard and optimally controlled network. As can be seen, there are two peak time-intervals in 24 h, and the vehicles' travel times dramatically decreased in both due to the ASLC management of traffic. The overall biggest reduction in mean travel time for a time step occurs in the second peak and has a value of 55 %, and the biggest reduction in a vehicle's travel time also occurs in the second peak and has a value of 69 % (see Table 5).

It can be argued that mean travel time may not be the best parameter to analyse the IA's performance, as the results indicate that, to maintain a more consistent traffic flow, it is necessary to reduce the speed limit in lanes where vehicles drive fastest, especially under approximately free-flow conditions. These obstacles prevent congested lanes from clearing because vehicles attempting to exit the junction into non-congested lanes are unable to do so, hindered by the fast-moving cars already occupying those lanes. Therefore, in this study, another parameter- the waiting time- is applied to provide a clearer assessment of the IA's performance. Waiting time represents the sum of time that a vehicle's speed is under 0.36 km/h, or in other words, that shows how long a vehicle spends almost stationary. Fig. 12 illustrates a comparison between waiting time in the standard network and the optimised network in two ways: a) for each vehicle and b) as the mean value in each time step. According to Fig. 12, waiting times drop sharply after this junction is operated by the ASLC speed limits. The maximum reduction in a vehicle's waiting time occurs during the second peak, with a relative decrease of 93 percent. Similarly, the maximum reduction in mean waiting time within a time-interval also occurs during the second peak, showing an 88 % reduction (see Table 6).

It is interesting to observe that along with the reductions in mean travel time and mean waiting time as a result of the ASLC, there is a slight increase in the finishing time steps at the end of each peak. One reason for this phenomenon is that the ASLC has previously lowered speed limits in certain lanes during earlier time-intervals. This action forces the faster vehicles to drive slower than they would in the standard network. Therefore, some of them enter one of the previous time segments and finish their trip through the junction later, thereby increasing the mean travel time and mean waiting time for that time-interval. In other words, traffic is delayed until the very end of peak times, when traffic flow is lower and has the capacity for the extra demand. This issue highlights areas for further research in ASLC applications, particularly exploring the relationship between the length of the time segments and

Table 4
The estimator's performance.

ANN structure	R-value	MAPE (%)	RMSE (s)
12–6–5–1	98 %	3.3	34.6

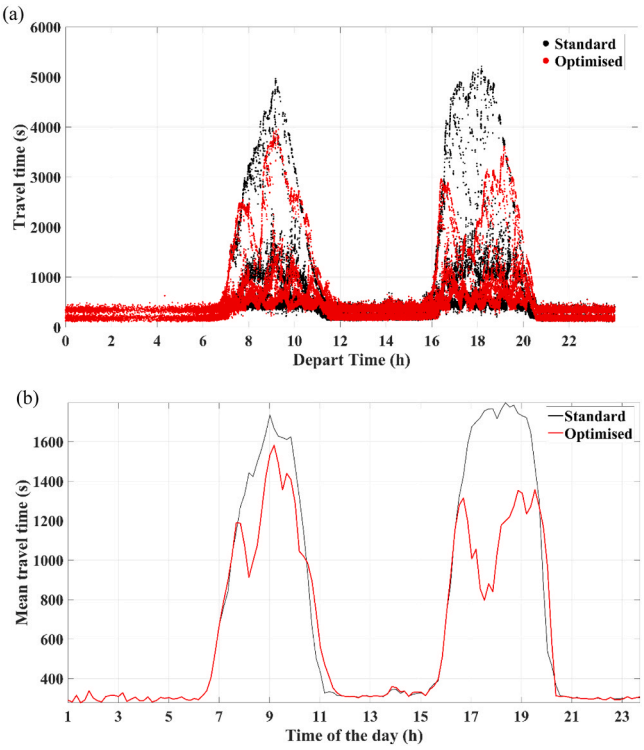

Fig. 11. Travel times in the standard (black) and optimised (red) motorway for a) each vehicle and b) the mean value for each time step.

Table 5
Travel time reduction.

First peak	Highest reduction in mean travel time of a time step (%)	The highest reduction in travel time of a vehicle (%)
First peak	37	60
Second Peak	55	69

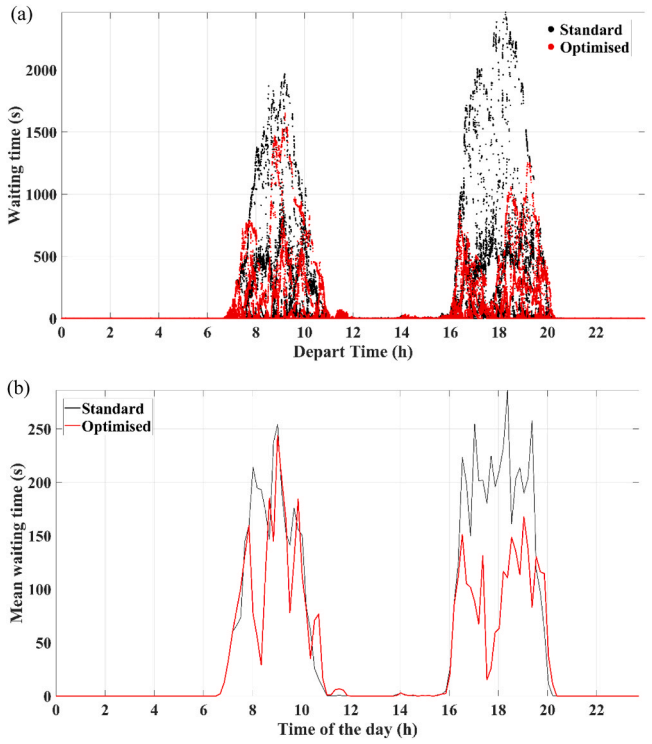

Fig. 12. a) Vehicles' waiting time in the standard (black) and optimised (red) motorway. b) Time steps mean waiting time in the standard (black) and optimised (red) motorway.

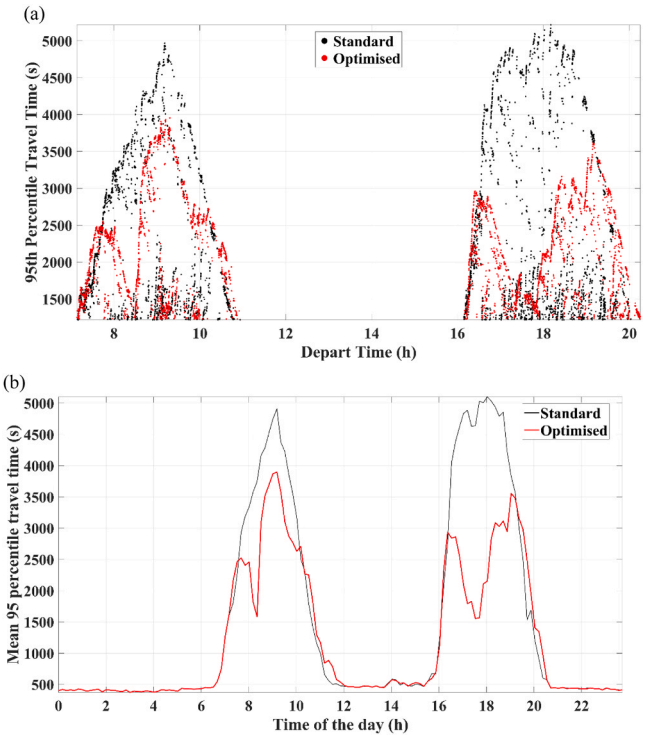

Fig. 13. 95th percentile values of travel time: a) scatter and b) time steps' mean.

Table 6
Waiting time reduction.

First peak	Highest reduction in mean waiting time of a time step (%)	Highest reduction in waiting time of a vehicle (%)
First peak	85	91
Second Peak	88	93

IA activation.

It is noted that traffic congestion is related to the most delayed vehicles, and this may be a small trade-off for the application of intelligent and autonomous traffic control through speed limits. In the context of traffic congestion control, the slowest vehicles are the target of the ASLC. The 95th percentile values of travel time (see Fig. 13) and waiting time (see Fig. 14) can then provide insight into the traffic performance of the most delayed cars since they only consider the 5 % most delayed vehicles for the greatest delays. As it can be inferred, the reduction in the mean 95th percentile travel time and waiting time after optimisation is larger than the reduction in overall mean travel time and waiting time. This indicates that the ASLC significantly enhances performance for the slowest vehicles as a means to reduce congestion. The biggest reductions in mean 95th percentile travel time in the first and second congestion peaks are 58 and 69 percent, respectively, while the same value for the waiting time in the first and the second peaks is 86 and 94 percent, respectively. Similarly, most of the time, congestion is reduced in the slowest lanes in motorways due to the ASLC slowing down the vehicles in the fastest lanes, facilitating easier lane change for vehicles in the slowest lanes if they wish to bypass junctions.

To conclude, it is also relevant to emphasise the difference in the two congestion peaks, where the first is relatively more complex for the ASLC and which may indicate a case of operation over capacity regardless of the traffic speed limit parameters.

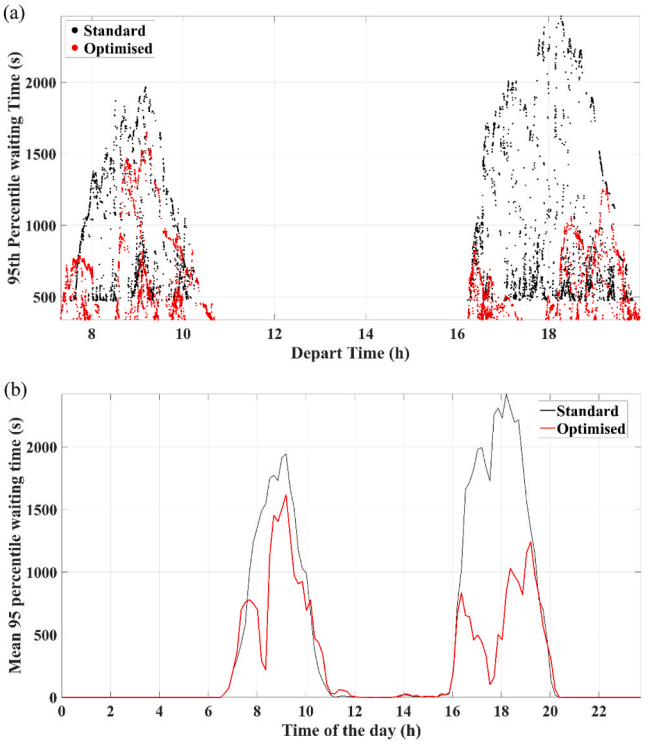

Fig. 14. 95th percentile values of waiting time: a) scatter and b) time steps' mean.

The analysis of the speed limits' performance amongst solutions also provides insight into this approach. In this illustrative example, four time intervals, i.e., 8:00 – 8:10 (congested traffic), 9:30 – 9:40 (saturated traffic), 17:50 – 18:00 (congested traffic), and 19:00 – 19:10 (free flow)

are used as reference, where a random sample of speed limit combinations is applied with results shown in Fig. 15. In this example, every possible combination of speed limits is implemented, and the simulation is run based on these combinations. Each lane has four possible speed limits, ranging from 36 km/h to 144 km/h, and the road segment features five lanes, resulting in a total of 1024 different scenarios. In the first example, i.e., 8:00 – 8:10, the mean travel time for the uncontrolled traffic is 1442.31 s, whereas for controlled traffic it is 912.7 s. The recommended combination is among the best-known solutions, i.e. 3.2 % lowest mean travel times. However, this specific example is also used to show that there are considerations to improve in the approach to find the global optimum, such as the reduction of the estimator error that implicitly makes the solution a metaheuristic. Nonetheless, the performance is close to the optimum, although it is obtained in virtually zero decision-time.

The recommended speed limits for the one hour from 8:30 am to 9:30 am are represented in Table 7. The big difference in speed limits in two adjacent lanes poses safety risks and may be impractical to implement. Large changes in the speed limits of a specific segment of a lane can be challenging and dangerous to enforce. The speed limit of the fourth lane in the first time-interval was increased from 36 km/h to 140 km/h in the following time-interval. These issues are more appropriately addressed from a policy perspective rather than methodologically restricting the agent. This means that the agent can propose solutions without constraints, and policies regarding speed limit oscillations and large changes can be implemented in it. It was also observed that in congested traffic conditions, the larger speed limits are not as important as the lower speed limits (the large ones may be unfeasible to adhere to, and from an optimisation perspective, represent a solution that may not present relevant sensitivity to changes in a lane if it has a traffic jam). Nevertheless, in this study, the agent's performance is also tested under safety constraints as well.

Table 7
Recommended speed limits by the ASLC decision-maker.

Time-intervals	Speed limits (km/h)				
	First lane	Second lane	Third lane	Fourth lane	Fifth lane
08:30–8:40	100	108	40	36	78
08:40–8:50	72	33	107	140	38
08:50–9:00	108	36	108	36	62
09:00–9:10	76	96	41	105	69
09:10–9:20	70	114	48	113	74
09:20–9:30	78	95	52	102	70

5.3. Safety considerations

Speed limit safety constraints can vary depending on the case study context, national traffic regulations, and operational considerations. Consequently, the acceptable magnitude of speed changes across time or lanes differs across the literature. A simplified safety comparison is now studied. In a recent study [68], the maximum allowable temporal change in the speed limit for the same lane across consecutive time intervals was defined as 10 mi/h, which corresponds to approximately 16 km/h. [69] established a 20 km/h maximum difference between the speed limits of adjacent lanes, representing a spatial change constraint. These two constraints are now incorporated into the proposed methodology, and Fig. 16 shows the results of the mean travel time with constraints.

As expected, the constrained ASLC outperforms the standard scenario throughout most of the day. However, during peak traffic periods, the unconstrained ASLC demonstrates a greater reduction in mean travel time compared to the former. This is expected as the unconstrained system has more flexibility in aggressively adjusting speed limits to alleviate bottlenecks. Notably, toward the end of each peak period, when traffic demand begins to subside, the constrained ASLC surpasses the unconstrained ASLC in performance. This suggests that in

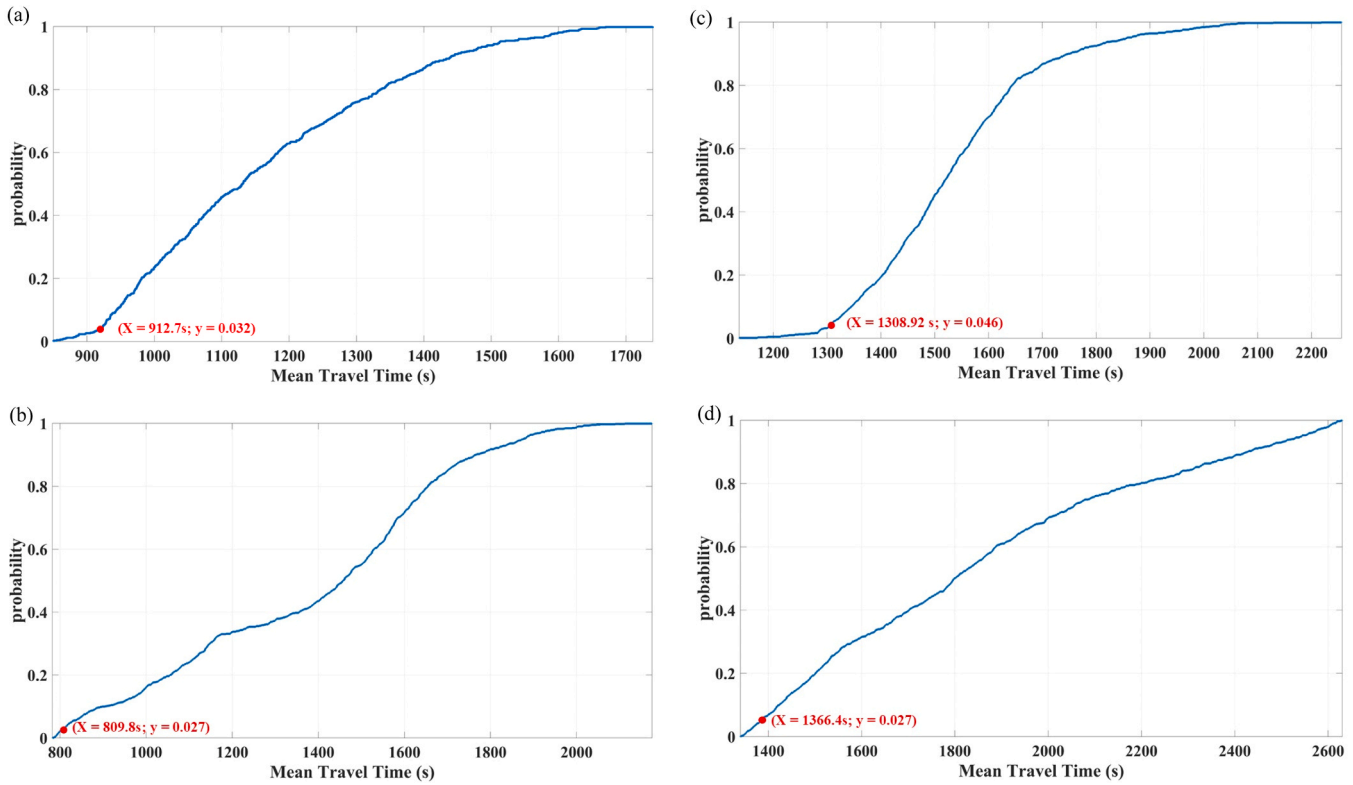

Fig. 15. Empirical cumulative density function.

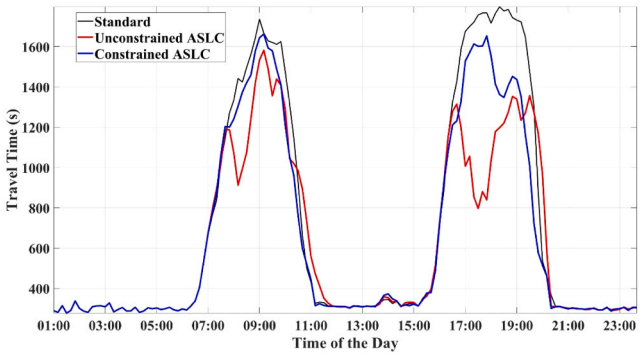

Fig. 16. Mean travel time in standard motorway (black), optimised by unconstrained ALC (red), and optimised by constrained ASLC (blue).

transitional phases, safer speed adaptations can result in improved efficiency. The results highlight the preliminary trade-offs between efficiency and safety, and support the inclusion of dynamic safety rules in practical deployment. Nonetheless, it is noted that further safety-specific research on the agent approach is expected to enhance the efficiency benefits.

5.4. Robustness and transferability of the ASLC

One important question is whether this methodology's outputs will be able to match the real operation of a network. To answer this question, it should be mentioned that there are three sources of output (travel time) in decision-making in this study. I) the real-world motorway, II) the high-fidelity simulator, and III) the IA (in this case ASLC). Simulators in the future are expected to become a full presentation of a real object's behaviour, such as traffic, and their outputs match the outputs from real data. Stakeholders currently use agent-based models to make decisions on main roads, such as in the case of the Dublin M50 motorway. Hence, the overarching vision of the proposed methodology is to demonstrate that if an IA's outputs can match those of a simulator, the IA is then an appropriate representation of reality. This aligns with the approach used for a Digital Twin [70], where the IA matches the operational dynamics of the real traffic network. This rationale uses the premise that any high-fidelity pairing of a physical system and its digital counterpart constitutes a high-fidelity simulator. Such simulators can be used offline to inform decision-making and hence allow a global understanding of the problem at hand. To do so, as shown in Fig. 17, there is a need to validate the ASLC's outputs by giving it an unseen dataset and to compare its outputs with the simulator.

As mentioned in Section 5.1, after training the estimator, an unseen dataset is used to validate the estimator's outputs. The results showed almost 97 % accuracy and a 0.98 correlation coefficient, which establishes the robustness of the estimator and ASLC. On the other hand, since the optimum number of the ANN's hidden layers and the number of neurons in the first and second hidden layers (i.e., two, six, and five, respectively) are less than the maximum allowed in the initial parameters (see Table 3), then increasing these and creating a larger ANN is not expected to produce a more accurate estimator. Moreover, as shown in Fig. 18, the reported minimum error was reached in iteration number 61

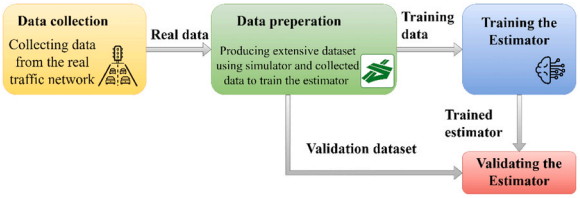

Fig. 17. Training and validation of the estimator.

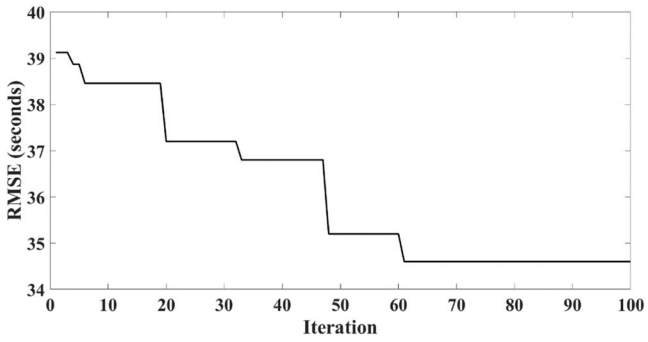

Fig. 18. Learning process of the estimator (test error in each iteration).

and the other 39 iterations after that did not increase the accuracy. Therefore, increasing the maximum number of iterations and the number of agents, hence increasing the computational cost is not expected to result in more accuracy.

5.5. Real-time assessment

As it was mentioned before, one important aspect of real-time optimality is to be able to find the optimum solution in so-called real-time. It is noted that the concept of real-time is relative to the system and its timescales, where some orders of seconds or minutes are still real-time, and where other decisions need to be achieved in milliseconds. Regardless, in all, having a low computational cost is critical. In the proposed framework, the decision maker that searches for and finds the optimum solution consists of an optimisation algorithm (as mentioned with SSA being applied) with a trained ANN (where the optimisation takes place). Hence, the computational cost of expensive traffic analyses based on high-fidelity data required to find accurate solutions is mitigated. In the current case study, the average time required to find the optimum is 10.18 s, which is less than 2 % of the time-interval considered (10 min). It is worth noting that the real-time performance of the proposed framework was evaluated on a personal laptop with the following specifications: 11th Gen Intel Core i7–1185G7 processor (base frequency 1.80 GHz, up to 4.80 GHz with Turbo Boost), 32 GB of RAM. As a result, it is assumed that in this timeframe no relevant changes occur that would change the input assumptions for decisions, and hence, the system can be assumed to be real-time. Moreover, it is noted that the methodology implemented is agnostic to the choice of estimator and decision-maker techniques. Thus, there is potential to further reduce the decision-making time with further research on different techniques. Furthermore, the application of alternative optimisation algorithms may also contribute to reducing this time. This process can be even shorter in the near future due to rapid technological development.

6. Conclusion and future opportunities

To reduce traffic congestion and have more consistent traffic flow, a new IA called ASLC, comprising two different parts: an estimator and a decision-maker, is introduced in this study. The M50 motorway in Dublin is used as a case study of its implementation, and SUMO is used to train, validate, and test the ASLC's outputs. The scenarios that are used to train the ASLC were built in SUMO based on a real dataset. The main findings of this study are as follows:

- The hybridisation of an ANN and an SSA as the estimator can successfully predict mean travel time in 10-minute time-intervals with almost 97 % accuracy in different contexts of decision (i.e., traffic parameters) as output from a fusion of real-data and high-fidelity modelling. The ANN uses 12 inputs (5 DVs and 7 PS), demonstrating ASLC's capability as a suitable approximator for real motorway outputs.

- The findings indicate that ASLC achieves a greater reduction in travel time and waiting time for the slowest 5 % of vehicles compared to the average reduction observed across all vehicles. Therefore, it can be concluded that ASLC is able to address congestion by creating a more homogeneous and consistent traffic flow.
- Limiting the ASLC by safety constraints can reduce its performance in the time intervals where it can reduce the mean travel time the most.
- The computational time has a significant role in real-time VSL optimisation problems, which has a direct relationship with the number of search agents in SSA (see Table 3) and the performance function that drives the optimisation. Since the objective function is a pre-trained ANN, the computational time of the full optimisation with SSA is less than a second and it can be neglected when compared to the time segment's duration (i.e., 10 min), and, hence, in the present work it is assumed to provide real-time optimal decisions.

Nonetheless, there are some opportunities that can be explored in future studies. First, the ASLC was not able to reduce the travel time and waiting time in every time step. Hence, there is a need to further investigate techniques that may allow this. Moreover, since the IA needs to be trained before implementing it on different motorways, transfer learning and making the agent transferable present an opportunity. Preliminary results showed that there is limited dependence in terms of performance on the optimisation algorithm used. However, a more detailed analysis of the sensitivity to these may be of further interest in future work. Moreover, scaling the methodology to multiple segments or traffic networks would necessitate coordination mechanisms between agents and other traffic control systems, such as ramp metering, to ensure the smoothness of traffic conditions. Addressing this gap can bring this methodology one step closer to implementation.

The proposed agent uses a multi-fidelity framework. In recurrent scenarios, the low-fidelity model, i.e., ANN, predicts travel time, allowing the agent to optimise per-lane VSLs. In non-recurrent cases, such as accidents or lane closures, the agent remains inactive while data is collected and used to recalibrate the high-fidelity model, i.e., SUMO, and to update the ANN offline for future use. This offline process does not interfere with the agent's online operations (see Fig. 2). Nonetheless, it is noted that the agent is designed to handle non-recurrent scenarios as its knowledge base expands. However, this is beyond the scope of the current study, i.e. inference on near real-time optimality, and is intended for future research.

It is noted that the limitations and the feasibility of the implementation should be further explored in future work, where the goal of the present work is to show the large potential that real-time optimisation allows. Moreover, driver compliance still remains as one of the main challenges in the field of VSL. Enforcing optimum speed limits and researching approaches to ensure they receive more attention in future work. However, full compliance is not needed since the speed of each vehicle is dictated by the vehicles in front of it. Therefore, partial compliance is enough for the ASLC to work [68]. This is one of the main steps toward implementing VSL methodologies. In this study, 10-minute time segments are considered. Having other time-intervals with different durations can be considered, and the performance of ASLC on these should be addressed in the future. Finally, the ASLC can be used to optimise a route after an accident or any disruptive event. To conclude, it is also important to highlight that this research is part of a larger context for optimal intelligent transportation systems. The differences encountered in the two congestion peaks in the case study developed are an indicator of the need to further combine it with other approaches and interventions in the network.

It is noted that the present framework depends on data for the training phase, and it is interesting to discuss where this data can be obtained in the traffic system. The current assumption in this work is that most of this data is, or will be, readily available in the near future, as there is already significant data being collected in traffic (tolls, car counters). Moreover, car users also provide important data that may in the future be used to further optimise traffic operation (cell data or connected cars). Such datasets would make all the required data for this implementation available. To conclude, it is important to remark that the purpose of the present work is to remark on how the ideas of IA can make a large impact in traffic management, in particular through the fusion of data, and offline and online training for real-time operation.

CRediT authorship contribution statement

Amirreza Kandiri: Software, Writing – original draft, Methodology, Visualization, Formal analysis, Validation, Data curation, Conceptualization. **Beatriz Martinez-Pastor:** Writing – review & editing. **Maria Nogal:** Supervision, Writing – review & editing. **Rui Teixeira:** Supervision, Conceptualization, Writing – review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. ANN

Neurons in the hidden layers receive weighted inputs from previous layers, sum them, and incorporate biases from the previous layer. This final value is then passed through an activation function. Output neurons follow a similar process but without using an activation function, while the input layer neurons simply import data into the ANN without applying any weights or biases [71]. ANNs are trained by optimising the value of the weights (w) and biases (b). Weights indicate the importance of a neuron's output on the input of a neuron in the next layer. Biases work as an offset or threshold, allowing neurons to produce an output even if all the weighted input values are zero. This adds a layer of adaptability, helping the network learn efficiently and make accurate predictions. ANNs learn different patterns and adapt to them using a constant value, i.e., bias. Fig. A.1 shows a schematic neuron in the hidden and output layers of an ANN.

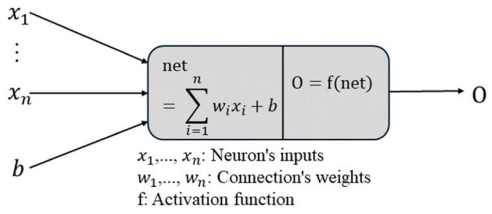

Fig. A.1. Neuron structure and its components adapted from [72]

Appendix B. Hybrid OANN technique to train the ANN

The various stages of the OANN approach are presented in Fig. B.1, and detailed explanations follow below. It should be noted that in Fig. B.1, SSA is used for the demonstration.

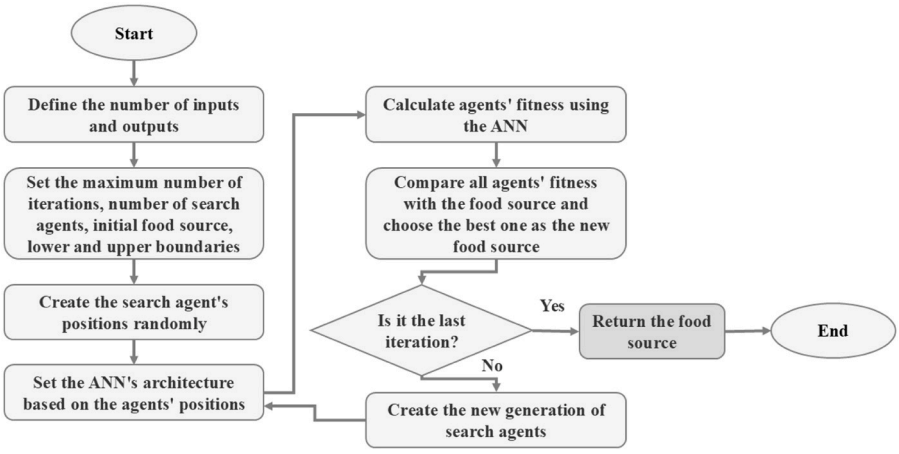

Fig. B.1. OANN's flowchart

The initialisation stage of OANN involves setting the algorithm's tuning parameters, which include maximum number of iterations, population size, initial food source (i.e., global optimum solution), upper and lower boundaries, and problem dimensions. Usually, the food source has an initial value that can be easily outperformed by the first generation of the population. Fig. B.2 illustrates how an individual of a population in an optimisation algorithm represents an ANN architecture, with two parts. In algorithms like the genetic algorithm, the representator is each individual's gene, and in swarm intelligence-based algorithms such as particle swarm optimisation, this representator is the individual's position. The first part is a binary code where 0 and 1 values of the *i*th hidden layer (h_i) indicate its inactivity or activity, respectively. The second part is related to the number of neurons in each layer, denoted by n_i. It should be noted that the maximum number of hidden layers is determined in the initialisation stage by upper and lower boundaries. For example, an individual's position of 1–0–1–2–4–6 represents an ANN architecture with two active hidden layers with 2 and 6 neurons in the first and third hidden layers, respectively, and an inactive hidden layer (the second hidden layer has a value of 0).

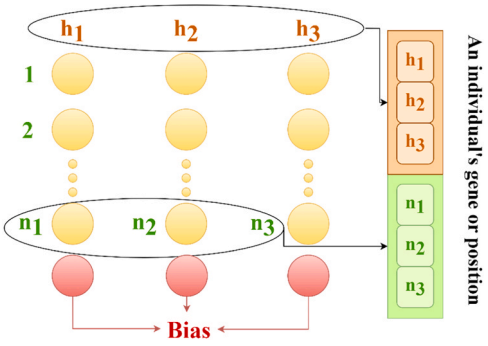

Fig. B.2. Slap position structure

Then, the initial positions of the population are randomly generated within the boundaries set in the previous stage. The objective of the OANN method is to minimise the error of the ANN. The objective function is defined as the average error of each individual, which is calculated using the k-fold cross-validation method. In this study, the root mean square error (RMSE) is applied as the statistical error indicator during the training, validation, and testing phases.

Appendix C. Salp Swarm Algorithm

The Salp Swarm Algorithm is a bio-inspired optimisation method that models the movement and foraging behaviour of salps in the ocean. Within this approach, potential solutions are represented by a group of salps arranged in a chain-like formation. The first salp in the chain, i.e., the leader, steers the swarm towards the target, i.e., food source, while the rest of the salps, i.e., followers, adjust their positions relative to the one ahead. This coordinated structure allows the algorithm to maintain a balance between exploring the search space and converging on promising solutions. SSA is particularly effective for solving complex, non-linear, and high-dimensional optimisation problems [34]. Fig. C.1. shows SSA's pseudo code.

```
Initialise the salp population
while (end condition is not satisfied)
  Calculate the fitness of each search agent (salp)
  F = the best search agent
  Update the position of the leader
  Update the position of the followers
  Amend the salps based on the upper and lower bounds of variables
end
return F
```

Fig. C.1. SSA's Pseudo code

Appendix D. ANN's weights and biases

$$\text{Input layer weights mathrix(IW)} = \begin{bmatrix} 0.001 & -0.04 & 0.03 & -0.05 & -0.13 & 1.73 & -2.20 & 0.12 & -0.13 & -0.11 & -0.03 & 0.10 \\ -0.036 & -0.20 & -0.18 & -0.18 & 3.05 & 5.47 & 0.86 & 0.63 & -0.02 & -1.25 & 0.42 & 1.05 \\ -0.03 & -0.09 & -0.003 & -0.11 & -0.24 & 1.79 & -1.6 & 0.08 & -0.07 & -0.23 & -0.13 & -1.17 \\ -0.02 & 0.01 & -0.06 & 0.07 & 0.13 & -3.33 & 0.49 & -0.14 & 0.10 & 0.01 & -0.40 & -0.41 \\ 0.02 & -0.02 & -0.12 & 0.003 & 0.09 & 0.03 & -0.46 & -0.60 & -0.71 & -0.08 & -0.34 & 7.12 \\ 0.001 & -0.02 & 0.04 & -0.04 & -0.10 & 0.60 & 0.64 & 0.13 & 0.09 & 0.11 & 0.21 & 0.18 \end{bmatrix}$$

$$\text{First hidden layer weights mathrix}(LW_1) = \begin{bmatrix} 2.66 & -0.0002 & -0.365 & 0.31 & 0.28 & 0.47 \\ -0.79 & -0.50 & -4.22 & -3.71 & 0.09 & -5.65 \\ 0.66 & 0.17 & -0.91 & 0.64 & 2.72 & 2.32 \\ 1.13 & -4.21 & 1.66 & -1.39 & 1.45 & 4.51 \\ 0.12 & -0.01 & -0.08 & 0.01 & -0.09 & -0.07 \end{bmatrix}$$

$$\text{Second hidden layer weights vector}(LW_2) = [-2.01 \quad -0.06 \quad 1.19 \quad 0.04 \quad -1.88]$$

$$\text{Input layer biases vector}(b_1) = [-0.05 \quad 2.32 \quad -0.92 \quad -0.10 \quad 3.61 \quad 0.93]$$

$$\text{First hidden layer biases vector}(b_2) = [-3.65 \quad 5.56 \quad -5.70 \quad -0.98 \quad -0.29]$$

$$\text{Second hidden layer biase}(b_3) = [-1.87]$$

Data availability

The data used was collected from a previous study.

References

[1] H. Chellapilla, et al., Bi-objective optimization models for mitigating traffic congestion in urban road networks, J. Traffic Transp. Eng. (Engl. Ed. 10 (1) (2023) 86–103.

[2] E. Walraven, M.T. Spaan, B. Bakker, Traffic flow optimization: a reinforcement learning approach, Eng. Appl. Artif. Intell. 52 (2016) 203–212.

[3] M. Zhang, et al., Urban travel time and residential location choice: the impacts of traffic congestion, Sustain. Cities Soc. 99 (2023) 104975.

[4] H. Dezani, et al., Optimizing urban traffic flow using genetic algorithm with petri net analysis as fitness function, Neurocomputing 124 (2014) 162–167.

[5] Mead, J.L.. R-TIS: real time travel information system for multi-modal real time information for Southampton Airport, the Olympic Park and the London Transport Museum. Stevenage: IET.

[6] A. Kumar, et al., A meta-heuristic-based energy efficient route modeling for EVs integrating start/stop and recapturing energy effect, Sustain. Cities Soc. 91 (2023) 104420.

[7] H. Majid, C. Lu, H. Karim, An integrated approach for dynamic traffic routing and ramp metering using sliding mode control, J. Traffic Transp. Eng. (Engl. Ed. 5 (2) (2018) 116–128.

[8] Y. Han, et al., A physics-informed reinforcement learning-based strategy for local and coordinated ramp metering, Transp. Res. Part C Emerg. Technol. 137 (2022) 103584.

[9] Y. Wu, et al., Double-layer ramp-metering model for incident congestion on expressway, J. Traffic Transp. Eng. (Engl. Ed. 1 (2) (2014) 129–137.

[10] W. Lu, et al., TD3LVSL: a lane-level variable speed limit approach based on twin delayed deep deterministic policy gradient in a connected automated vehicle environment, Transp. Res. Part C. Emerg. Technol. 153 (2023) 104221.

[11] Y. Han, et al., A new reinforcement learning-based variable speed limit control approach to improve traffic efficiency against freeway jam waves, Transp. Res. Part C. Emerg. Technol. 144 (2022) 103900.

[12] J. Feng, et al., Multi-Lane Differential Variable Speed Limit Control via Deep Neural Networks Optimized by an Adaptive Evolutionary Strategy, Sensors 23 (10) (2023) 4659.

[13] J. Cao, et al., Exploring the impact of a coordinated variable speed limit control on congestion distribution in freeway, J. Traffic Transp. Eng. (Engl. Ed. 2 (3) (2015) 167–178.

[14] I. Karafyllis, M. Papageorgiou, Feedback control of scalar conservation laws with application to density control in freeways by means of variable speed limits, Automatica 105 (2019) 228–236.

[15] P. Goatin, S. Göttlich, O. Kolb, Speed limit and ramp meter control for traffic flow networks, Eng. Optim. 48 (7) (2016) 1121–1144.

[16] L. Maggi, S. Sacone, S. Siri, Freeway traffic control considering capacity drop phenomena: comparison of different MPC schemes. 2015 IEEE 18th International Conference on Intelligent Transportation Systems, IEEE, 2015.

[17] J.R.D. Frejo, B. De Schutter, Logic-based traffic flow control for ramp metering and variable speed limits—Part 1: Controller, IEEE Trans. Intell. Transp. Syst. 22 (5) (2020) 2647–2657.

[18] A. Hegyi, B. De Schutter, H. Hellendoorn, Model predictive control for optimal coordination of ramp metering and variable speed limits, Transp. Res. Part C Emerg. Technol. 13 (3) (2005) 185–209.

[19] R.C. Carlson, et al., Optimal mainstream traffic flow control of large-scale motorway networks, Transp. Res. Part C Emerg. Technol. 18 (2) (2010) 193–212.

[20] R. Li, et al., Traffic control optimization strategy based on license plate recognition data, J. Traffic Transp. Eng. (Engl. Ed. 10 (1) (2023) 45–57.

[21] A. Messmer, M. Papageorgiou, METANET: A macroscopic simulation program for motorway networks, Traffic Eng. Control 31 (9) (1990).

[22] C. Daganzo, The cell transmission model Part I: a simple dynamic representation of higheay traffic, PATH Rep. 3 (1993) 93–0409.

[23] C.F. Daganzo, The cell transmission model: A dynamic representation of highway traffic consistent with the hydrodynamic theory, Transp. Res. Part B Methodol. 28 (4) (1994) 269–287.

[24] Z. Li, et al., Reinforcement learning-based variable speed limit control strategy to reduce traffic congestion at freeway recurrent bottlenecks, IEEE Trans. Intell. Transp. Syst. 18 (11) (2017) 3204–3217.

[25] F. Zhu, S.V. Ukkusuri, Accounting for dynamic speed limit control in a stochastic traffic environment: A reinforcement learning approach, Transp. Res. Part C Emerg. Technol. 41 (2014) 30–47.

[26] T.D. Kulkarni, et al., Hierarchical deep reinforcement learning: Integrating temporal abstraction and intrinsic motivation, Adv. Neural Inf. Process. Syst. (2016) **29**.

[27] R.S. Sutton, Introduction: The challenge of reinforcement learning. in Reinforcement learning, Springer, 1992, pp. 1–3.

[28] M. Gregurić, K. Kušić, E. Ivanjko, Impact of deep reinforcement learning on variable speed limit strategies in connected vehicles environments, Eng. Appl. Artif. Intell. 112 (2022) 104850.

[29] Y. Wu, H. Tan, B. Ran, arXiv preprint, arXiv:1810.10952,, Differ. Var. Speed Limits Control Free. Recurr. Bottle via Deep Reinf. Learn. (2018).

[30] C. Hickert, S. Li, C. Wu, Cooperation for scalable supervision of autonomy in mixed traffic, IEEE Trans. Robot. (2023).

[31] J. Guo, S. Cheng, Y. Liu, Merging and diverging impact on mixed traffic of regular and autonomous vehicles, IEEE Trans. Intell. Transp. Syst. 22 (3) (2020) 1639–1649.

[32] J. Kennedy, R. Eberhart, Particle swarm optimization. in Proceedings of ICNN'95-international conference on neural networks, IEEE, 1995.

[33] S. Mirjalili, S.M. Mirjalili, A. Lewis, Grey wolf optimizer, Adv. Eng. Softw. 69 (2014) 46–61.

[34] S. Mirjalili, et al., Salp Swarm Algorithm: A bio-inspired optimizer for engineering design problems, Adv. Eng. Softw. 114 (2017) 163–191.
[35] A. Lad, P. Kanaujia, Y. Solanki, Computer Vision enabled Adaptive Speed Limit Control for Vehicle Safety. 2021 International Conference on Artificial Intelligence and Machine Vision (AIMV), IEEE, 2021.
[36] B. Zhai, et al., Adaptive Control Strategy of Variable Speed Limit on Freeway Segments under Fog Conditions, J. Transp. Eng. Part A Syst. 149 (10) (2023) 04023097.
[37] F. Vrbanić, et al., Reinforcement learning-based dynamic zone positions for mixed traffic flow variable speed limit control with congestion detection, Machines 11 (12) (2023) 1058.
[38] M. Gregurić, F. Vrbanić, E. Ivanjko, Impact of federated deep learning on vehicle-based speed control in mixed traffic flows, J. Parallel Distrib. Comput. 186 (2024) 104812.
[39] R. Teixeira, et al., Metamodel-based metaheuristics in optimal responsive adaptation and recovery of traffic networks, Sustain. Resilient Infrastruct. 7 (6) (2022) 756–774.
[40] B.M. Angadi, M.S. Kakkasageri, S.S. Manvi, Computational intelligence techniques for localization and clustering in wireless sensor networks. in Recent Trends in Computational Intelligence Enabled Research, Elsevier, 2021, pp. 23–40.
[41] R. Teixeira, et al., The role of multi-fidelity modelling in adaptation and recovery of engineering systems, Acta Polytech. CTU Proc. 36 (2022) 224–230.
[42] B. Peherstorfer, K. Willcox, M. Gunzburger, Survey of multifidelity methods in uncertainty propagation, inference, and optimization, Siam Rev. 60 (3) (2018) 550–591.
[43] X. Li, C. Wang, H. Shi, A travel time prediction method: Bayesian reasoning state-space neural network. in The 2nd International Conference on Information Science and Engineering, IEEE, 2010.
[44] N. Wisitpongphan, W. Jitsakul, D. Jieamumporn, Travel time prediction using multi-layer feed forward artificial neural network. 2012 Fourth International Conference on Computational Intelligence, Communication Systems and Networks, IEEE, 2012.
[45] P. Gao, et al., Travel time prediction with immune genetic algorithm and support vector regression. 2016 12th World Congress on Intelligent Control and Automation (WCICA), IEEE, 2016.
[46] B. Yu, et al., Prediction of bus travel time using random forests based on near neighbors, Comput. -Aided Civ. Infrastruct. Eng. 33 (4) (2018) 333–350.
[47] Y. Zhang, A. Haghani, A gradient boosting method to improve travel time prediction, Transp. Res. Part C Emerg. Technol. 58 (2015) 308–324.
[48] S. Oh, et al., Short-term travel-time prediction on highway: a review of the data-driven approach, Transp. Rev. 35 (1) (2015) 4–32.
[49] M. Bai, et al., Travel-time prediction methods: a review, *December 10–12, 2018, Proceedings*, in: in Smart Computing and Communication: Third International Conference, SmartCom 2018, 3, Springer, *Tokyo, Japan*, 2018.
[50] A. Kandiri, F. Sartipi, M. Kioumarsi, Predicting compressive strength of concrete containing recycled aggregate using modified ANN with different optimization algorithms, Appl. Sci. 11 (2) (2021) 485.
[51] K. Hosseini, et al., E-bike to the future: Scalability, emission-saving, and eco-efficiency assessment of shared electric mobility hubs, Transp. Res. Part D Transp. Environ. 133 (2024) 104275.
[52] [cited 2022 18 Aug]; Available from: ⟨https://www.rte.ie/news/business/2022/0818/1316570-m50-toll-income-rose-13-last-year-as-traffic-recovered/⟩.
[53] R. Corbally, L. Yang, A. Malekjafarian, Predicting the duration of motorway incidents using machine learning, Eur. Transp. Res. Rev. 16 (1) (2024) 14.
[54] A. Kandiri, et al., Travel time prediction for an intelligent transportation system based on a data-driven feature selection method considering temporal correlation, Transp. Eng. 18 (2024) 100272.
[55] M. Rogers, S. Darcy, Traffic-flow impact of toll booths on M50 motorway, Dublin. in Proceedings of the Institution of Civil Engineers-Transport, Thomas Telford Ltd, 2007.
[56] D. Laoide-Kemp, M. O'Mahony, Dealing with latency effects in travel time prediction on motorways, Transp. Eng. 2 (2020) 100009.
[57] de Paor, C., et al., The Role of Motorway Traffic Flow Optimisation Indicators in Enhancing Motorway Operation Services in the Irish Road Network.
[58] M. Guériau, I. Dusparic, Quantifying the impact of connected and autonomous vehicles on traffic efficiency and safety in mixed traffic. 2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC), IEEE, 2020.
[59] P.A. Lopez, et al., Microscopic traffic simulation using sumo. 2018 21st international conference on intelligent transportation systems (ITSC), IEEE, 2018.
[60] J. Erdmann, SUMO's lane-changing model, *May 15-16, 2014*. in Modeling Mobility with Open Data: 2nd SUMO Conference 2014, Springer, *Berlin, Germany*, 2015.
[61] P.-Y. Ting, et al., Freeway travel time prediction using deep hybrid model–taking Sun Yat-Sen freeway as an example, IEEE Trans. Veh. Technol. 69 (8) (2020) 8257–8266.
[62] X. Miao, et al., Examining the impact of different periodic functions on short-term freeway travel time prediction approaches, J. Adv. Transp. 2020 (1) (2020) 3463287.
[63] L. Li, et al., Short-term highway traffic flow prediction based on a hybrid strategy considering temporal–spatial information, J. Adv. Transp. 50 (8) (2016) 2029–2040.
[64] A. Petrella, M. Miozzo, P. Dini, Mobile traffic prediction at the edge through distributed and deep transfer learning, IEEE Access (2024).
[65] Y. Tian, et al., LSTM-based traffic flow prediction with missing data, Neurocomputing 318 (2018) 297–305.
[66] K.Y. Chan, et al., Neural-network-based models for short-term traffic flow forecasting using a hybrid exponential smoothing and Levenberg–Marquardt algorithm, IEEE Trans. Intell. Transp. Syst. 13 (2) (2011) 644–654.
[67] D.D. Hema, K.A. Kumar, Levenberg–marquardt–lstm based efficient rear-end crash risk prediction system optimization, Int. J. Intell. Transp. Syst. Res. 20 (1) (2022) 132–141.
[68] F. Alasiri, Y. Zhang, P.A. Ioannou, Per-lane variable speed limit and lane change control for congestion management at bottlenecks, IEEE Trans. Intell. Transp. Syst. 24 (12) (2023) 13713–13728.
[69] T. Yuan, et al., Evaluation of integrated variable speed limit and lane change control for highway traffic flow, IFACPap. 54 (2) (2021) 107–113.
[70] K. Kušić, R. Schumann, E. Ivanjko, A digital twin in transportation: Real-time synergy of traffic data streams and simulation for virtualizing motorway dynamics, Adv. Eng. Inform. 55 (2023) 101858.
[71] A. Kandiri, et al., Modified Artificial neural networks and support vector regression to predict lateral pressure exerted by fresh concrete on formwork, Int. J. Concr. Struct. Mater. 16 (1) (2022) 64.
[72] A. Kandiri, E.M. Golafshani, A. Behnood, Estimation of the compressive strength of concretes containing ground granulated blast furnace slag using hybridized multi-objective ANN and salp swarm algorithm, Constr. Build. Mater. 248 (2020) 118676.