

HumanLight: Incentivizing ridesharing via human-centric deep reinforcement learning in traffic signal control

Dimitris M. Vlachogiannis^{a,b,*}, Hua Wei^c, Scott Moura^{a,b}, Jane Macfarlane^{a,b}

^a Department of Civil and Environmental Engineering, University of California at Berkeley, Berkeley, 94720, CA, United States

^b Energy Technologies Area, Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, 94720, CA, United States

^c Department of Informatics, Ying Wu College of Computing, New Jersey Institute of Technology, Suite 5700 University Heights, Newark 07102, NJ, United States

ARTICLE INFO

Keywords:

Person-based traffic signal control
Decentralized adaptive control
Deep reinforcement learning
Ridesharing
Multimodal traffic environment

ABSTRACT

Single occupancy vehicles are the most attractive transportation alternative for many commuters, leading to increased traffic congestion and air pollution. Advancements in information technologies create opportunities for smart solutions that incentivize ridesharing and mode shift to higher occupancy vehicles (HOVs) to achieve the car lighter vision of cities. In this study, we present HumanLight, a novel decentralized adaptive traffic signal control algorithm designed to optimize people throughput at intersections. Our proposed controller is founded on reinforcement learning with the reward function embedding the transportation-inspired concept of pressure at the person-level. By rewarding HOV commuters with travel time savings for their efforts to merge into a single ride, HumanLight achieves equitable allocation of green times. Apart from adopting FRAP, a state-of-the-art (SOTA) base model, HumanLight introduces the concept of active vehicles, loosely defined as vehicles in proximity to the intersection within the action interval window. The proposed algorithm showcases significant headroom and scalability in different network configurations considering multimodal vehicle splits at various scenarios of HOV adoption. Improvements in person delays and queues range from 15% to over 55% compared to vehicle-level SOTA controllers. We quantify the impact of incorporating active vehicles in the formulation of our RL model for different network structures. HumanLight also enables regulation of the aggressiveness of the HOV prioritization. The impact of parameter setting on the generated phase profile is investigated as a key component of acyclic signal controllers affecting pedestrian waiting times. HumanLight's scalable, decentralized design can reshape the resolution of traffic management to be more human-centric and empower policies that incentivize ridesharing and public transit systems.

1. Introduction

Population growth along with urbanization have rendered today's transportation networks more saturated than ever before. Private vehicle ownership grows yearly, particularly since the COVID-19 outbreak (United States Department of Transportation, 2021). As a result, traffic congestion has become one of the key challenges metropolitan areas are facing. Apart from increased travel times and vehicle miles traveled, the associated environmental consequences have made transportation the number one contributor to the climate crisis in America today. In many urban areas, the inconveniences of unreliable transit are high for

* Corresponding author at: Department of Civil and Environmental Engineering, University of California at Berkeley, Berkeley, 94720, CA, United States.
E-mail addresses: d.vlachogiannis@berkeley.edu (D.M. Vlachogiannis), janemacfarlane@berkeley.edu (J. Macfarlane).

commuters (Malalgoda and Lim, 2019; Diab et al., 2015), which has resulted in diminishing public transit ridership, leaving private vehicles as the most attractive solution for commuters.

For policymakers, looking decades ahead, the reduction of trips performed in single occupancy vehicles (SOVs) is of paramount importance. With the COVID-19 pandemic coming to an end, incentivizing public transit use and shared modes of transportation is imperative. Ridesharing or pooling is defined as the act of two or more travelers sharing the same vehicle for a common trip. Shared rides will reduce vehicle miles traveled, energy use, and greenhouse gas emissions (Viegas et al., 2016). Apart from the system-level benefits, the users individually benefit from shared travel costs (Shaheen et al., 2016). Car pooling originally became popular in metropolitan regions with the establishment of High Occupancy Vehicle (HOV) lanes. People casually grouped in one vehicle to save on time, toll fares, and gas (Shaheen and Cohen, 2019). In the last decade, app-based pooling services that reserve, match, and process payments for rides on demand have been developed. By grouping passengers with similar origin and destination locations within a walking radius from the common route, transportation network companies in urban areas provide flexible and affordable ridesharing solutions (Lazarus et al., 2021). Compared to the ride-alone option, users of a pooled ride are quoted a discounted price for a usually longer estimated total travel time due to the added walking time and pick up of additional passengers.

Transit and purpose-driven shuttles (e.g. airport drop-off services) have been around for decades as fixed-guideway systems that cannot respond to dynamic passenger demand. On-demand transit-like services, typically comprised of vans, shuttles and buses, have also recently emerged and are commonly referred to as microtransit (Shaheen and Cohen, 2019; Bardaka et al., 2020). Particularly in the post-pandemic era which reshaped travel patterns to require greater flexibility, the on-demand services offered by microtransit with more direct routing, reduced transfers, and better coverage can significantly improve rider experience (Schank and Huang, 2022). As opposed to traditional transit, microtransit generates routes and stops in response to real-time demand. Furthermore, microtransit vehicles are guided by GPS with real-time traffic and can therefore adapt to on-the-ground conditions such as traffic jams and road closures. Microtransit solution experiments have been mostly successful and are envisioned as a key component in transportation networks by many forward-thinking transit agencies. For those reasons, this study incorporates microtransit in the formulation of HOV adoption scenarios to represent the multimodal composition of vehicle fleets in the upcoming decades.

The dynamic travel behavior and demand fluctuations across on-demand mobility options have led to a decline of carpooling to currently only 7.79% of the commute mode share in the United States for 2021 (United States Department of Transportation, 2021). The major progress in ride-hailing matching algorithms for mobility on demand is mostly achieving cheaper rides for commuters. However, Lazarus et al. (2021) have indicated that monetary benefits alone do not outweigh the loss in travel time poolers sacrifice to merge into a single ride. For pooling to be competitive, shared rides should be prioritized through smart mobility solutions that democratize travel times benefits. Currently, poolers and non-poolers share the travel time benefits of reduced vehicle counts in the transportation network thanks to poolers' shift towards HOVs. Instead, fairer and more socially equitable mechanisms would reward poolers for their efforts to merge rides by providing reduced travel times. HOV lanes are currently the only case where poolers receive travel time benefits. In urban environments, HOV lanes are mostly reserved for freeways and highways. According to Giuliano et al. (1990), the introduction of HOV lanes did not lead to a significant increase in ridesharing among the population of the examined route's commuters, but only in some specific periods. Results suggested that barriers to increased pooling and mode shift are formidable, requiring the development of additional carpooling incentivizing strategies based more on travel time savings. These findings are in alignment with studies suggesting that cost reductions of pooling solutions do not outweigh the potential delays in travel time (Cohen et al., 2021). Information and connected car technologies enabling real time data sharing and high performance computation now make the window of opportunity for the implementation of smart mobility solutions greater than ever.

In this study, we present HumanLight, a traffic management solution developed to operate with future vehicle to infrastructure (V2I) communication technologies. A novel decentralized adaptive signal control algorithm that optimizes people's throughput at intersections is formulated to support a future modal split among private vehicles, pooled rides and public transit. Our proposed solution sets the foundations for breaking free of today's dependence on cars by prioritizing people as opposed to vehicle movement. By devolving transportation policy to be more human-centric, ridesharing and public transit systems can be re-invigorated and attract the travel demand they truly merit in sustainable and multimodal urban environments. The reinforcement learning-based signal controller HumanLight is designed to reward, with more green times, vehicles carrying more people. We explore the headroom of such prioritization strategies at intersections at different rates of vehicle pooling adoption and demonstrate the higher travel time benefits for HOV adopters.

2. Literature review

2.1. Reinforcement learning in traffic signal control

Traffic controllers are typically classified into fixed-time, actuated, and adaptive controllers. Conventional traffic signal control (TSC) methods are mostly fixed-time and have been developed to heavily rely on pre-defined rules and assumptions on traffic conditions (Martinez et al., 2011; Roess et al., 2004). Traffic splits are typically derived to alleviate traffic congestion on uniform traffic flow distribution. Webster's method (Koonce and Rodegerdts, 2008) is one of the most widely-used methods in the field for a single intersection setting. It determines the optimum cycle length and phase split for a single intersection according to historical traffic data collected at different times during the day. Actuated signal controllers are more responsive to the traffic flows as they use real-time measurements from sensors. In most cases though, the timing plan parameters, such as maximum green and extension time, are optimized offline (Shabestary and Abdulhai, 2018).

Adaptive traffic signal controllers have been shown to outperform fixed-time and actuated controllers because they can adjust traffic phase splits, or skips traffic phases, according to dynamic and unpredictable traffic demand patterns. Reinforcement Learning (RL) based TSC is a model-free and self-learning adaptive strategy. By interacting with the environment, usually a microscopic simulator, the agent learns to adapt to the evolving real-time traffic conditions (Wei et al., 2018, 2019a; Chen et al., 2020). Recently, with the advent of deep learning and the use of deep neural networks to approximate key components of reinforcement learning, deep reinforcement learning (DRL) has enabled a continuous state space representation.

A key component of the design of a RL-based TSC system is the formulation of the algorithmic setting, comprising of the state and action space and the reward function. For the state representation, different quantitative descriptions of the environment have been proposed, including queue length (Wei et al., 2018; Bakker et al., 2010; Kuyer et al., 2008), waiting time (Van der Pol and Oliehoek, 2016; Chu et al., 2019), vehicle volume (El-Tantawy and Abdulhai, 2010; Wei et al., 2019a) approximations of delay (Arel et al., 2010), speeds (Casas, 2017; El-Tantawy and Abdulhai, 2010; Nishi et al., 2018) or phase setting (Aslani et al., 2017; Wei et al., 2019a, 2018; El-Tantawy et al., 2013; Chen et al., 2020). As for the design of the action space, two main representations are used in literature: (i) the algorithm selects traffic phases in an acyclic manner (Shabestary and Abdulhai, 2022; Wei et al., 2019a,b; Chen et al., 2020; Yang et al., 2019; Oroojlooy et al., 2020), (ii) the algorithm, following an ordered sequence of traffic phases, determines the traffic phase splits by regulating the duration of the current phase either by keeping or switching it (Mannion et al. (2016), Wei et al. (2018) and Aslani et al. (2017)). The reward, or objective function of the RL problem, has been modeled with several surrogate metrics as vehicle travel time can only be measured post trip completion. Metrics that have shown to provide effective formulations of the reward include queue length (Wei et al., 2019b; Van der Pol and Oliehoek, 2016; Kuyer et al., 2008; Aslani et al., 2017; Salkham and Cahill, 2010), waiting time (Bakker et al., 2010; Chu et al., 2019; Van der Pol and Oliehoek, 2016; Wei et al., 2018), speed (Wei et al., 2018; Van der Pol and Oliehoek, 2016; Casas, 2017), number of stops (Van der Pol and Oliehoek, 2016), throughput (Aslani et al., 2017; Wei et al., 2018; Salkham and Cahill, 2010) or even safety related features such as accident avoidance (Van der Pol and Oliehoek, 2016; Du et al., 2022).

In an effort to derive a more theoretically grounded reward function formulation and identify how much neighboring information is necessary in the state space representation, Wei et al. (2019a) drew inspiration from max pressure (MP), a state-of-the-art (SOTA) method in adaptive traffic signal control, as proposed in Varaiya (2013) and Lioris et al. (2016). The MP algorithm proved that by minimizing pressure of the intersection, defined as the difference between the total queue length on incoming approaches and outgoing approaches, the risk of over-saturation is reduced and the throughput of the whole road network is maximized. Wei et al. (2019a) tested using pressure as part of the long term reward function to achieve a formulation more theoretically justified without deploying combinations of heuristic metrics that also require weight tuning. PressLight is also the first RL model that automatically achieves coordination along arterials without prior knowledge. The framework was tested under uniform unidirectional traffic with the optimal solution being known to be the greenwave. PressLight learned to regulate consecutive traffic signals' switches with an offset equivalent to the expected vehicle travel time between intersections.

Despite directly interacting with a highly dynamic environment and learning to reach a long-term goal for traffic signal control, RL methods are susceptible to the curse of dimensionality (Rasheed et al., 2020; Yau et al., 2017) an issue whereby the state space becomes too large. This leads to higher computational costs during the exploration of the state-action pairs resulting in a longer learning time, as well as requiring a larger storage capacity to store the learned Q-values. MPLight (Chen et al., 2020) was recently designed as a decentralized framework to address those scalability issues. The proposed approach was applied for multi-intersection control, and specifically in a large-scale network of over 1000 intersections in New York City. Thanks to the pressure-based design and the parameter sharing capabilities in model learning, MPLight showcased strong performance and generalization ability. MPLight also adopts the FRAP architecture as its base model (Zheng et al., 2019). FRAP is a Deep Q-Learning method specifically developed for traffic signal control problems and designed based on the principles of phase competition and invariance. FRAP not only achieves superior performance with fast convergence, but is also resilient in handling complex intersection structures and multi-intersection environments.

The vast majority of SOTA studies showcase the potential of reinforcement learning in TSC in a vehicle-level optimization setting under the consideration of uni-modal traffic, neglecting the heterogeneous interests and complex interactions of multimodal traffic. Their vehicle-centric design provides optimal movement of vehicles disregarding their occupancy.

2.2. Transit signal priority and person-level traffic optimization

Human mobility and optimization of people throughput at intersections is essential for establishing efficient, sustainable and equitable transportation operations. Over the last decade, automatic passenger counts (APCs), the foundation of occupancy data, are rapidly growing in their adoption within transit fleets. APCs have already been established as standard practice by several transit agencies, as real-time occupancy is a necessary piece to the complete mobility picture. The American Public Transportation Association reported in 2020 that approximately 40% of all United States transit vehicles have APCs installed, with commuter buses exceeding 58%. This information is already helping agencies improve transit operations by reducing bus bunching and passenger pass-ups. Monitoring passenger counts is invaluable for TSC frameworks aiming to minimize person delays, as it facilitates the implementation of traffic management strategies capable of achieving efficient and reliable public transit.

A transit signal priority algorithm was proposed in Christofa and Skabardonis (2011) and Christofa et al. (2013), formulated as a mixed-integer non-linear program, to minimize total passenger delay while assigning priority to transit vehicles based on their passenger occupancy. The developed macroscopic mathematical model of delay at an intersection considers both regular and transit vehicles, and assumes under-saturated traffic conditions and fixed cycle lengths and phase sequences. The work was later extended to

the arterial level where intersection pairs were simultaneously optimized and schedule adherence was incorporated (Christofa et al., 2016). Another real-time TSC system, minimizing total person delays of cars and transit vehicles while ensuring a priority window for transit vehicles to address the issue of wasted priority due to bus arrival uncertainty, was proposed by Farid and Christofa (2017). Several other systems handling transit vehicle priority in bi-modal traffic environments have been developed (Hu et al., 2015; Yu et al., 2017; Zeng et al., 2015; Yu et al., 2018).

Only recently have traffic signal control algorithms founded on RL and accounting for person-level optimization appeared in literature (Rasheed et al., 2020). Using available information received from high-detailed traffic sensors, Shabestary and Abdulhai (2019) proposed a multimodal Deep RL based traffic signal controller that combines both regular traffic and public transit and minimizes the overall travelers' delay through the intersection. The controller builds on the authors previous work (Shabestary and Abdulhai, 2018). The position and average speed of vehicles are used as state inputs, the change of delay is defined as the reward function, and actions are set to flexibly choose the next phase in the pre-defined phase set. Long et al. (2022) also adopted a person-based reward function to propose an extended Dueling Double Deep Q-learning (DDQL) algorithm, eD3QNI, to improve bus operational efficiency and handle conflicting bus priority requests. Performance is evaluated by simulation for a single intersection with two traffic demands and random arrivals, schedule deviations, and occupancies of buses. The authors also performed an exploration around the penetration rate of connected buses, illustrating that it does not affect the convergence speed but it will affect performance. Another person-based approach also restricted in an isolated intersection was proposed by Wang et al. (2022). The reward function was built around passenger waiting time and queue length and optimized using the DDQL network. The authors even account for non-motorized traffic including pedestrians, and report a 6.3% decrease in waiting time compared to a baseline vehicle-level approach. To the best of our knowledge, existing literature still lacks a scalable RL-based multi-intersection traffic signal controller capable of handling multimodal traffic and generalizable to large scale networks.

2.3. Contributions and organization

The key contributions of this research work can be summarized as follows. HumanLight is the first human-centric RL-based decentralized adaptive traffic signal controller that is scalable to transportation networks of multiple intersections such as corridors and grids. HumanLight is designed to democratize urban traffic by allocating green times of the different phases in a socially equitable manner. By providing travel time savings to riders of HOVs, HumanLight can be a powerful tool for policymakers to incentivize a shift from privately owned motorized transportation towards shared and equitable mobility. Our proposed algorithm's effectiveness is evaluated in a variety of road network configurations and mode share scenarios. The different multimodal scenarios include diverse distributions of SOVs, carpools, microtransit and public transportation to quantify HumanLight's potential at different levels of HOV adoption.

We introduce some key methodological novelties that enable HumanLight to achieve robust performance in person-level optimization. Inspired by the concept of pressure (Varaiya, 2013; Wei et al., 2019a) quantifying the degree of disequilibrium between vehicle density on the incoming and outgoing lanes, HumanLight extends the idea to person pressure minimization to achieve optimal people throughput at intersections. We also introduce the concept of active vehicles, loosely defined as those in proximity to the intersection within the action interval window. Through systematic experiments, the utility of the active vehicle consideration during person-level optimization is shown to improve handling the high variance of vehicle occupancies in multimodal traffic environments.

In addition to enabling reduced passenger travel times and socially equitable green time allocation at signalized intersections, HumanLight enables policymakers and traffic engineers to control the aggressiveness in the prioritization of HOVs. We achieve this via a modification in the state embedding where vehicle occupancies are encoded. This way, the travel time benefits across vehicle types of different occupancies are parameterized for the system operator to adjust rather than being generated in a fixed approach.

Experiments are conducted to establish the most effective formulation of the algorithm (state space and reward function), as well as, parameter setting and tuning. Experiments include exposure to different traffic demand scenarios and network structures while the analytics extend further from traditional traffic metrics (travel times, delays and queues) to traffic signal design aspects such as average phase durations and number of changes. To examine the potential failure points of HumanLight in network settings, the study also includes analyses on maximum vehicle queues and stopping behaviors of the different vehicle types as a result of the person-level TSC policies.

The remainder of this paper is structured as follows. In Section 3, we introduce some preliminary definitions and formulate the decentralized multi-intersection traffic signal control problem. Section 4 presents the agent design and provides details on the deployed Deep Q-learning model. In Section 5, the experimental framework is demonstrated along with the methods of comparison and evaluation metrics. Section 6 discusses HumanLight's performance for different network configurations and demand profiles and presents analyses justifying the algorithmic formulation and the policymaking impacts of the proposed solution. Section 7 concludes the study and Section 8 suggests directions for future research.

3. Problem definition

3.1. Preliminaries

Definition 1 (Traffic Network). The traffic network is represented as a directed graph, with intersections modeled as nodes and road segments between intersections as edges. Each intersection may have both incoming and outgoing roads serving the upstream

and downstream traffic respectively. Each road may be comprised of multiple lanes. Road segment i of intersection I is denoted by r_i^I while lane l of r_i^I is denoted as $r_{(i,l)}^I$. We represent the set of incoming and outgoing lanes of intersection I as L_{in}^I and L_{out}^I respectively.

Definition 2 (Traffic Movement). Traffic movement is defined as the traffic traveling across an intersection from an incoming road to an outgoing road. We denote a traffic movement from r_i^I to r_j^I as (r_i^I, r_j^I) , in which $(r_i^I, r_j^I) = (r_{(i,l)}^I, r_{(j,m)}^I)$, $l \in r_i^I \subset L_{in}^I$, $m \in r_j^I \subset L_{out}^I$.

Definition 3 (Signal Phase). A traffic signal phase of intersection I , ϕ^I , is defined as a permissible combination traffic movements.

Definition 4 (Movement, Phase and Intersection Pressure). Vehicle pressure of a traffic movement is defined as the difference of vehicle density between the upstream lane $r_{(i,l)}^I$ and the downstream lane $r_{(j,m)}^I$:

$$p_v(r_{(i,l)}^I, r_{(j,m)}^I) = \frac{C_v(r_{(i,l)}^I)}{C_v^{max}(r_{(i,l)}^I)} - \frac{C_v(r_{(j,m)}^I)}{C_v^{max}(r_{(j,m)}^I)} \quad (1)$$

where $C_v(\cdot)$ is the vehicle count function, and $C_v^{max}(\cdot)$ the maximum permissible vehicle number for a road segment. If all lanes of the transportation network share the same maximum vehicle capacity, vehicle pressure can be simplified to capture the difference in vehicle counts between the upstream and downstream traffic.

The vehicle pressure of a traffic phase ϕ^I is defined as the sum of pressures over all traffic movements comprising the traffic phase:

$$p_v(\phi^I) = \sum_{(r_{(i,l)}^I, r_{(j,m)}^I) \in \phi^I} p_v(r_{(i,l)}^I, r_{(j,m)}^I) \quad (2)$$

The intersection vehicle pressure is defined as the sum of pressures over all traffic movements:

$$P_v^I = \sum_{(r_{(i,l)}^I, r_{(j,m)}^I) \in I} p_v(r_{(i,l)}^I, r_{(j,m)}^I) \quad (3)$$

Accordingly, we derive person intersection pressure (P_p^I) by substituting the vehicle count function $C_v(\cdot)$ of Eq. (1) with a person count function $C_p(\cdot)$:

$$P_p^I = \sum_{(r_{(i,l)}^I, r_{(j,m)}^I)} p_p(r_{(i,l)}^I, r_{(j,m)}^I) \quad (4)$$

with the maximum permissible person number, $C_p^{max}(\cdot)$, derived based on the vehicle type v which at capacity would maximize the total passenger count for the road segment:

$$C_p^{max}(\cdot) = \max_v (O_v \cdot C_v^{max}(\cdot)) \quad (5)$$

In real-world world scenarios, person counts can be generated via automated passenger counters (APCs).

Definition 5 (Active Vehicles (AV)). Literature has shown that vehicle-level traffic conditions can be adequately described by lane counts, including all vehicles in the incoming lanes (Casas, 2017; El-Tantawy and Abdulhai, 2010; Bakker et al., 2010; Wei et al., 2018), and queue lengths, including all vehicles stopped in queue waiting to traverse the intersection (Chen et al., 2020; El-Tantawy and Abdulhai, 2010; El-Tantawy et al., 2013; Wei et al., 2019a). In a person-level setting though, the distribution of people's locations along a segment may display significant variance, especially the more the world shifts to a multi-modal HOV setting.

The goal is to assure that only persons capable of traversing the intersection are accounted for in our state space and reward function calculation. Inspired by the work of Zhang et al. (2022), we introduce the concept of active vehicles (AV) to describe vehicles that are in range of the controlled intersection within the action interval window. As opposed to the predecessor study, vehicle range (L) is not calculated under maximum speed but by deploying the position equation of motion assuming constant acceleration and a maximum vehicle speed threshold, according to Eq. (6):

$$L_{(v,t_a)}(v_{(v,t_a)}, v_v^{max}, a_v^{max}) = v_v^{max} \Delta t - \frac{1}{2 a_v^{max}} (v_v^{max} - v_{(v,t_a)})^2 \quad (6)$$

where $v_{(v,t_a)}$ is vehicle's v speed at timestep t_a when action a is to be taken, v_v^{max} , a_v^{max} are the vehicle's maximum speed and acceleration, and Δt is the action interval.

For lanes upstream of the intersection, we compute the vehicles' maximum feasible projected location at the end of the upcoming action a . Our concept is also extended for outgoing lanes to properly inform the algorithm regarding the impact of the applied policies. For vehicles downstream of the intersection, we are interested on whether the controller's last selected action resulted in the vehicle crossing the intersection and joining the outgoing lane. In this case, the vehicle location at the moment of the




Fig. 1. (a) Standard four-way intersection with 8 incoming lanes. (b) Traffic movement illustration and pressure calculation. (c) Traffic phase illustration and pressure aggregation. The MaxPressure optimal phase is highlighted. We emphasize two key points: (i) Accounting for people movement returns a different optimal phase than the traditional vehicle-level pressure maximization strategy. (ii) The introduction of active vehicles has no impact in the vehicle-level approach (same phase is returned as optimal) but radically affects the person-level approach (the returned phases are not only different but share no common movements).

previously selected action can be either collected from past location data or approximated similarly to the incoming lane case. Eq. (7) summarizes the definition of the active vehicle set for intersection I at timestep t_a , $AV_{t_a}^I$:

$$AV_{t_a}^I = \left\{ v \in V \mid d_{(v,t_a)}^I \leq L_{(v,t_a)}, \forall v \in V \right\} \quad (7)$$

where $d_{(v,t_a)}^I$ is the distance of vehicle v from intersection I at timestep t_a , and V is the set of all vehicles.

In Section 6.4, we evaluate the efficacy of the AV consideration under different topological settings. Fig. 1 illustrates the defined concepts in a standard four-way intersection with 8 incoming and outgoing lanes. Traffic movements and phases are identified and the corresponding pressure values are calculated both at the vehicle and person-level, as well as, with or without the consideration of AV. Fig. 1 highlights that accounting for people movement returns a different optimal phase based on the MaxPressure algorithm than the traditional vehicle-level pressure maximization strategy. Another key observation is that the introduction of active vehicles has no impact in the vehicle-level approach (same phase is returned as optimal) but radically affects the person-level approach as the returned phases are not only different but share no common movements.

3.2. Formulation: Decentralized multi-intersection traffic signal control

Each traffic light regulated intersection in the network is controlled by a reinforcement learning agent based on the real-time people distribution along the incoming and outgoing lanes. Under the decentralized structure, agent steps are performed individually by each agent at every distinct action interval with duration Δt , while all agents of the network engage in the same learning process. Post training, the controllers will have learned to select the appropriate signal phases that maximize passenger throughput through the intersections.

The traffic signal control problem is formulated as a Markov Decision Process $\mathcal{M} = \langle S, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma \rangle$ (Sutton and Barto, 2018). S is a finite state space, $\mathcal{A}$ is a finite action space, $\mathcal{P}(s'|s, a) : S \times \mathcal{A} \times S \rightarrow [0, 1]$ is the state transition probability from state s to state s' determined by action a , and $\mathcal{R}(s, a) : S \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function defined as $\mathcal{R}(s, a) = \mathbb{E}[R_{t+1}|s_t = s, a_t = a]$. At timestep t , agent I aims to learn policy $\pi(a_t = a|s_t = s)$ returning the optimal action a given the state s to maximize the discounted reward:

$$J_t(\pi) = \sum_{m=0}^{\infty} \gamma^m R_{t+m+1} \quad (8)$$

The agent calibrates its estimates of the executed action's utility based on environmental feedback and will potentially adjust the rates of the actions leading up to the current action.

4. Methodology

4.1. Agent design

In this section, we describe the basic constructs of the reinforcement learning algorithm. The state and action spaces and the reward function are defined at the level of a single intersection I .

- State space: The state space includes the number of persons drawn from the active vehicles in each incoming lane and in each outgoing lane of the intersection, $C_p^{AV}(r_{(i,l)}^I)$ and $C_p^{AV}(r_{(i,m)}^I)$ respectively with leL_{in}^I and meL_{out}^I . The state space also includes the current traffic signal phase ϕ^I .
- Action space: The traffic phase ϕ_{t+1}^I for the next action interval. For most common transportation intersections with up to four entering and existing roads, the maximum compatible and non-conflicting phase combinations are eight for each isolated intersection (Fig. 1). As per common practice in RL research on traffic signal control (Wei et al., 2019b, 2018, 2019a,c; Chen et al., 2020), we adopt the acyclic phase paradigm which, as opposed to the cyclic paradigm, does not require a predetermined phase sequence to be imposed and for all available phases to appear at least once within a cycle. By enabling phase shortening and skipping, acyclic phasing schemes have demonstrated superior performance (Kanis et al., 2021) and are already implemented in urban environments (e.g. Amsterdam). However, our RL representation is still fully transferable to a cyclic phase modeling scheme.
- Reward function: For each individual agent, the reward is defined as the opposite of person intersection pressure as derived from Eq. (4):

$$R^I = -P_p^I \quad (9)$$

4.2. Base model: FRAP

The FRAP architecture (Zheng et al., 2019) is adopted as the base model for our traffic signal control system. FRAP is a Deep Q-Learning method specifically designed for traffic signal control problems. In vehicle-level traffic optimization, FRAP captures the competition relation between different traffic movements achieving superior performance and improved convergence. The authors also illustrated the model's transferability and adaptability to different traffic signal settings, roads structures and unbalanced traffic flows. The model is designed based on the principals of:

- Competition: Phases with higher traffic demand need to be prioritized.
- Invariance: Flipping or rotation of traffic flows along the intersection should not affect algorithmic performance.

The prediction of the Q-values is divided into three stages:

1. Phase demand modeling: Obtains a representation for the demand of each signal phase. The state features for each movement are extracted from the simulator and passed through two fully-connected layers. The outputs for the non-conflicting movements of every phase are added to generate the phase demand for green signal.
2. Phase pair representation: Establishes the phase pairs demand embeddings and applies convolutional layers with 1×1 filters to extract the phase pair representations, enabling the competing phases to interact with each other.
3. Phase pair competition: Predicts Q-values for each phase. The relative priorities of each phase are computed by multiplying the phase pair demand representation with a phase competition mask and applying an additional convolutional layer with 1×1 filter. The row sum of this pairwise competition matrix returns the array of Q-values, whose maximum value dictates the action to be selected.

All agents (traffic signal regulated intersections) of the evaluated network share the same FRAP model. The replay memory stores experiences (observations, actions and rewards) from all intersections which are used for the model's parameter updating.

Table 1
Simulation and model parameters.

Simulation			
Simulation time	1 h	CityFlow step length	1 s/step
Number of phases	8	Type of signal	Acyclic
Yellow and all red period	5 s	Action interval	10 s
Vehicle maximum speed	40 km/h	Vehicle maximum acceleration	2 m/s ²
Vehicle maximum deceleration	4.5 m/s ²	Vehicle minimum gap	2.5 m
Headway	2 s	SOV and carpool vehicle length	5 m
Microtransit (occupancy 10) vehicle length	7 m	Microtransit (occupancy 20) vehicle length	9.5 m
Microtransit (occupancy 30) vehicle length	12 m	Public transit vehicle length	15 m
FRAP model			
Discount factor γ	0.6	Episodes	200
Learning rate	0.001	Buffer size	10 000
Batch size	32	Greedy policy ϵ	1 $\rightarrow$ 0
Learning start	0	Optimizer	Adam
Number of layers (phase demand modeling)	2	Number of convolution layers w/ 1×1 filters	20




Fig. 2. (a) Adjusted driving commute mode share statistics drawn from the Bureau of Transportation Statistics for the year 2021. (b) Vehicle occupancy profiles for the HOV adoption scenarios (low, light, medium and high). Across each of the four scenarios, 15% of the single occupancy vehicle individuals shift to microtransit vehicles and public transit. Carpool ridership has been preserved at the same levels.

5. Experiment

All of our experiments are conducted in an open-source microscopic traffic simulator called CityFlow (Zhang et al., 2019; Mei et al., 2022), whose interface supports RL training for large-scale traffic signal control. Each green signal is followed by three-second yellow time during which -if possible- vehicles always choose to stop, and two-second all red time to prepare the signal phase transition. Right turning vehicles are considered to be allowed to turn during any phase but always yield priority to through moving vehicles in case of conflict. The action interval set to 10 s as per common practice in RL research on acyclic TSC (Wei et al., 2018, 2019b,a,c; Chen et al., 2020). Wei et al. (2018) also concluded that the action time interval has minimal influence on performance when in the range of 5 to 25 s. The adopted values for all simulation and model parameters are detailed in Table 1.

5.1. Mode share scenarios

Our experiments are set up by establishing a total people count served by the transportation network when it is operating almost at capacity. This is assessed through a critical movement analysis on the intersections' hourly expected flows (Urbanik et al., 2015). We maintain those people counts constant throughout the mode share scenarios investigating different levels of HOV adoption. Note that we assume no latent demand in the network due to the increased level of service in order to quantify the benefits of mode shift with person-based RL traffic signal controlled intersections.

Determining mode share percentages is a complex task and distributions may vary per location, especially after the COVID-19 pandemic. For the lowest HOV adoption scenario, the selected numbers aim to reflect the significant reduction transit ridership took globally during the pandemic. Even, in areas such as the Bay Area where transit was highly adopted, recent studies estimate that between 66% and 78% of Bay Area commuters drive alone when commuting (Cowan et al., 2021). As shown in Fig. 2a, commute mode data was drawn from the Bureau of Transportation Statistics for 2021 in California to extract the mode share for persons driving vehicles. People walking, working from home and biking were excluded. The percentage of the taxi, motorcycle, or other class was merged into the drive alone class. The carpooling class is split to two sub-classes of occupancy two and three with ratios

Table 2

Average flow per direction per lane in the single intersection experiment across all demand scenarios.

Demand scenario	North/South (veh/h/lane)		East/West (veh/h/lane)		v/c
	Through	Left	Through	Left	
Low HOV adoption	220	220	350	220	0.99
Light HOV adoption	194	194	307	194	0.87
Moderate HOV adoption	157	157	249	157	0.71
High HOV adoption	121	121	190	121	0.54




Fig. 3. Road network configurations: (a) Isolated four-way intersection where each road segment has one through and one left turning lane, (b) Six intersection corridor and (c) 4×4 grid of intersections with three lanes per road segment (one through, one left turning, and one right turning). Fixed public transit routelines are illustrated in green.

2/3 and 1/3 respectively. Fig. 2b summarizes the defined mode share scenarios. For our experimental setup, commute mode changes across scenarios with approximately 15% of the single occupancy vehicles individuals shifting to microtransit vehicles and public transit, while carpool ridership has been preserved at the same levels. The low HOV adoption scenario has transit set to its lowest levels of the last decade during allowing us to explore the potential of person-level traffic signal control optimization in a wide spectrum of HOV penetration environments.

5.2. Infrastructure configurations

The set of synthetic experiments were designed to test the performance and efficiency of HumanLight at different road network configurations. These are outlined below and displayed in Fig. 3:

1. Single intersection
2. Corridor of six intersections (1×6) with buses operating only through the corridor.
3. Grid of intersections (4×4) with two fixed bus route-lines.

For the single four-way intersection experiment, each road segment has one through and one left turning lane while average hourly traffic demands are described in Table 2. We impose high left turning volumes equal to the through-moving traffic of the secondary direction. This allows us test the robustness of HumanLight on handling demand of different distributions across directions compared to the following configurations. For the corridor and grid configurations, each road segment has one through, one left turning and one right turning lane. The turning ratios at intersections are on average set as 10% left, 60% straight and 30% right as statistical analyses on real-world data sets have showcased (Chen et al., 2020; Vlachogiannis et al., 2023). The road networks have

Table 3
Demand configurations for corridor (1 × 6) and grid (4 × 4) experiments.

Demand scenario	Arrival rate (veh/h/road)
Low HOV adoption	700
Light HOV adoption	571
Moderate HOV adoption	469
High HOV adoption	372

segments of 300 m in length and connect at four-way intersections. The synthetic data fed into the simulator include bi-directional flows with turning traffic. To assess our traffic signal control models under various traffic demands the configurations of Table 3 are realized describing the different mode shift scenarios. The maximum speeds are set equally for all vehicle types since the imposed values for speed limits are not that high at 40 km/h.

5.3. Methods of comparison

In this section, we introduce classical TSC methods in the transportation field (Roess et al., 2004; Wei et al., 2019c; Martinez et al., 2011), and current RL-based methods. For the purposes of evaluating HumanLight's performance, all methods including parameter setting are tuned. The best performing method will be used as benchmark for the computation of the evaluation metrics.

- **Fixed Time Controller - Webster's Formula** (Koonce and Rodegerdts, 2008): Fixed-time controllers use a predetermined cycle and phase time plan. Webster's method calculates the cycle length and phase split for a single intersection setting. Under the assumption of uniform traffic flow during a certain period of time, using the closed-form solution of Eq. (10), Webster's formula derives the optimal cycle length C , minimizing the travel time of all vehicles passing the intersection.

$$C(V_c) = \frac{N * t_L}{1 - \frac{V_c}{3600/h * PHF * (v/c)}} \quad (10)$$

where N represents the number of phases, t_L the total loss time per phase (used to model the all-red time and the acceleration and deceleration of vehicles), parameter h the saturation headway time, PHF the peak hour factor (used to model traffic demand fluctuations within peak hours), and v/c is desired volume-to-capacity ratio. V_c expresses the sum of all critical lane volumes, with $V_c = \sum_{p_i}^N V_c^{p_i}$, where $V_c^{p_i}$ is the critical lane volume for phase p_i . The critical lane volume is determined from the approaching lane with the highest ratio of traffic flow to saturation flow during a phase.

As for the phase split, having established the cycle length, green times are calculated proportionally to the critical lane volumes served by each phase according to Eq. (11):

$$\frac{t_{p_i}}{t_{p_j}} = \frac{V_c^{p_i}}{V_c^{p_j}} \quad (11)$$

where t_{p_i} and t_{p_j} correspond to the phase duration for phases p_i and p_j respectively.

- **Self-Organizing Traffic Light (SOTL)** (Gershenson, 2004; Cools et al., 2013): SOTL is an actuated method that adaptively regulates traffic lights based on a hand-tuned threshold on the number of waiting vehicles. The controller switches the phase of a lane to green if the required minimum green phase duration is met for the current phase and provided that the number of vehicles on the lane exceeds the hand-tuned threshold. SOTL resembles a fully-actuated controller but instead of sending requests for green when a single vehicle is approaching, it does so when the number of vehicles exceeds the threshold.
- **MaxPressure** (Varaiya, 2013): Max Pressure sets the optimization objective as minimizing the vehicle pressure of phases for individual intersections, as defined in Section 3. The method greedily selects the phase with the maximum pressure, activates it and keeps the selected phase for a given period of time t_{min} . A sensitivity analysis on parameter t_{min} is carried out every time the algorithm is applied to identify the best performing value.
- **MPLight** (Chen et al., 2020): MPLight is a state-of-the-art deep reinforcement learning method enabling large scale road network control. Aside from parameter sharing, it accommodates a SOTA Q-network model structure (FRAP) as proposed in Zheng et al. (2019) and can incorporate the concept of vehicle pressure in the reward function as proposed in Wei et al. (2019a).

5.4. Evaluation metrics

Both person and vehicle based metrics will be used to evaluate the impact of the proposed traffic signal control strategy. These are:

1. Average Vehicle Travel Time (AVTT) & Average Person Travel Time (APTT)

$$AVTT = \sum_{v \in V} TT_v / |V| \quad (12)$$




Fig. 4. Convergence curves of HumanLight's reward function (top) and average person delay (bottom) throughout the learning episodes for all network structures across all HOV adoption scenarios: Episodic evolution of average person pressure is approximating average person delays. Average person queues display similar trends and are provided in Fig. A.1.

where TT_v denotes the travel time of vehicle v and $|V|$ is the cardinality of the vehicle set V , and

$$APTT = \sum_{v \in V} TT_v \cdot O_v / \sum_{v \in V} O_v \quad (13)$$

where O_v is the occupancy of vehicle v .

2. Vehicle Queue Length (VQL) & Person Queue Length (PQL)

$$VQL(t) = \left| \{v \in V | v_{(v,t)} \leq v_{stop}, \forall v \in V\} \right| \quad (14)$$

where v_{stop} is the speed lower bound for vehicles to be considered in motion and is set at 0.1 m/s, and

$$PQL(t) = \sum_{v \in V | v_{(v,t)} \leq v_{stop}} O_v \quad (15)$$

3. Average Vehicle Delay (AVD) & Average Person Delay (APD)

$$AVD = \sum_{v \in V} (TT_v - FTT_v) / |V| \quad (16)$$

where FTT_v is the free-flow travel time of vehicle v , derived accounting for the maximum speed allowed on each road segment along the vehicle's route, and

$$APD = \sum_{v \in V} (TT_v - FTT_v) \cdot O_v / \sum_{v \in V} O_v \quad (17)$$

6. Results

This section quantitatively highlights the key contributions of HumanLight. Firstly, the overall performance of the algorithm is discussed compared to SOTA controllers. Then, the socially equitable allocation of green times achieved by HumanLight is illustrated. A vehicle stopping behavior analysis as well as edge cases of vehicle queues are displayed to assure the stability of the framework. The impact of incorporating the concept of active vehicles in the formulation of our RL model is quantified through systematic experiments at different network structures. The capability of HumanLight to regulate the aggressiveness of HOV prioritization is showcased via a modification in the state embedding. Finally, the impact of the discount factor parameter γ on the generated phase profile is investigated.

The presented results in Sections 6.1–6.3 have been reproduced and validated across all network configurations considered (single intersection, corridor and grid). The algorithmic formulation as well as hyper-parameters have been set based on comprehensive




Fig. 5. Performance evaluation of HumanLight versus the SOTA vehicle-level optimization controllers in the 4×4 grid setting: the person-level benefits (a&b) of HumanLight increase as more people shift towards HOV alternatives. Thanks to the consideration of active vehicles, vehicle-level metrics (c&d) also improve. The corresponding results for the single intersection and corridor experiments are provided in [Tables A.1, A.2 and A.3](#).

experiments presented in Sections 6.4–6.6. We selectively include results from specific networks in this section, while adding the rest to the [Appendix](#) to avoid overwhelming readers with potentially similar patterns. As per common practice in RL research ([Long et al., 2022](#); [Henderson et al., 2018](#); [Vlachogiannis et al., 2020](#); [Dietterich et al., 1998](#)), multiple, in our study three, independent trials of learning are run for each scenario and the evaluation metrics (travel times, delays, queues) reported in the results are derived from averaging the last 20 episodes across all runs to smooth out episodic oscillations. The tables both in the main body and in the [Appendix](#) section provide the standard deviations of those values in parenthesis.

6.1. Convergence and overall performance

HumanLight displays strong performance across all considered road network configurations. The algorithm consistently converges after approximately 160 episodes of training across all HOV adoption scenarios ([Fig. 4](#) top). We observe that average person intersection pressure, the opposite of our reward function, follows a similar diminishing trend as average person delays ([Fig. 4](#) bottom) and queues ([Fig. A.1](#)), validating the suitability of person pressure for the task of optimizing people throughput at intersections.

[Fig. 5](#) summarizes the performance of the various controllers in a 4×4 grid setting. Not distinguishing between vehicle types when it comes to determining intersection priorities, vehicle-level optimization controllers render similar average average vehicle and person delays. We observe however, that vehicle delays are consistently lower than person delays in the corridor ([Table A.2](#)) and grid experiments ([Fig. 5](#), [Table A.3](#)). This is attributed to transit's routelines never including a right turn in those configurations ([Fig. 3](#)). With the most dense, occupancy-wise, vehicle type never making right turns which are independent from the current phase, person travel times are expected to be slightly higher. Instead, in the single intersection setting where no right turns exist, the trend does not appear ([Table A.1](#)).

By rewarding HOV riders with more green times, HumanLight achieves reduced person delays and queues. The higher the penetration of HOVs becomes, the improvements from the SOTA vehicle-level optimization controller in person delays and queues increase, even exceeding 55% in the high HOV adoption scenario. In [Fig. 5c&d](#), we observe that HumanLight achieves improvements even on vehicle-level metrics, attributed to the consideration of active vehicles in the RL problem formulation.

6.2. Socially equitable allocation of green times

Apart from superior performance compared to SOTA approaches, HumanLight achieves a fair allocation of green times to vehicles based on the number of passengers they are carrying. [Fig. 6](#) demonstrates the distributions of travel times for the different vehicle




Fig. 6. (a,c) Travel time distribution for different vehicle types over the last episode (200). (b,d) HumanLight prioritizes vehicles based on their occupancy as shown by the linear least-squares regression fitted on the box-plot means of all vehicle types over the last 20 episodes of each run.

types over the last episode of training (left) and the linear least-squares regression fitted on the box-plot means over the last 20 episodes of each run to smooth out the effect of fluctuations across episodes (right). All vehicles are considered in the single intersection setting, while only vehicles traversing all six intersections are accounted for in the corridor experiments to assure all vehicles share the same route length and thus have comparable travel time. In Fig. 6c, carpools and microtransit are purposefully not displayed due to data sparsity, as limited amount of some of those vehicles fulfilled the entire six intersection route especially in the lower HOV adoption scenarios.

Examining those distributions of travel times, the following key observations are derived. Firstly, the vehicle-level SOTA optimization controller (MPLight w/ FRAP & Pressure, shown in red) provides similar travel times across all vehicle types. Instead, HumanLight prioritizes vehicles based on their occupancy with the differences in travel time being eminent even at the low HOV adoption scenario (orange) with a strong coefficient of determination R^2 between the two variables. This trend being validated at the corridor level illustrates the scalability of the model to larger networks and paves the way for large-scale evaluations with a focus on route-level effects for public transit. The lower bound of the minimum travel time is only achievable in an unrealistic scenario where all vehicles arrive at the intersection from directions served by non conflicting maneuvers within the action interval window.

6.3. Vehicle stopping behavior and queues

An analysis on vehicle stopping behavior is carried out in the 4×4 grid network. As mentioned in the queue definition, a vehicle is considered stopped or in queue during the duration its speed is less than or equal to 0.1 m/s. We present the per segment normalized number of stops for the various vehicle types to account for the different lengths of routes taken. In Fig. 7, both in terms of normalized per segment stopping frequency as well as average stop duration, HumanLight achieves reduced and shorter stops for HOVs. Instead, the vehicle-level approach (MPLight) generates similar stopping patterns across vehicle types. The peaks in the normalized stop frequency for transit can be explained from their routine. As shown in Fig. 3, bus-lines in the grid scenario never




Fig. 7. Vehicle stopping behavior: Both in terms of normalized per segment stopping frequency (a) as well as average stop duration (b), HumanLight achieves reduced and shorter stops for HOVs as opposed to the vehicle-level approach (MPLight) that generates similar stopping patterns for all vehicles.

take a right turn. In contrast, all other vehicles' routes may contain right turns that use dedicated right lanes, independent of the current phase and yielding priority only to through moving vehicles in cases of conflict.

A crucial aspect of implementing a human-centric approach, such as HumanLight, is the impact on vehicle-level metrics. Prioritizing HOVs can lead to network failure by generating queues in non-prioritized directions of traffic. Spill-back effects are caused when vehicle queues exceed road capacities, significantly deteriorating the operational efficiency of the signalized intersections. To quantify that effect on the 4×4 grid network, we evaluate the maximum queue length observed on the incoming lanes of each controlled intersection at every action interval. Fig. 8a compares the distributions of the averaged over time maximum queues generated from HumanLight versus MPLight, the state of the art vehicle-level approach. Thanks to the consideration of active vehicles in the formulation of the RL problem, HumanLight overall achieves similar or even better maximum vehicle queues while still prioritizing the maximization of people throughput. Only in the low and light HOV adoption scenarios, maximum vehicle queues degrade at some intersections compared to MPLight.

Fig. 8b isolates the average, over time, value of the maximum vehicle queues generated by HumanLight for each intersection in the light HOV scenario. In parentheses, the percentage change from MPLight is displayed. We highlight that MPLight is a SOTA vehicle-level controller, not currently in practice. The worst performing intersections experience the heaviest public transit load by serving conflicting and left turning bus routelines. To accommodate the arrival of buses, the controller dictates more frequent phase changes, overwhelming the intersection with stopped vehicles in the non-served traffic movements. Even though for some intersections the maximum queues observed degrade compared to MPLight, the absolute values (illustrated in bold) are still low and far lower from the segments' capacity at jam density. For higher HOV adoption scenarios, the active vehicles consideration in HumanLight renders even better comparative results because of the increased penetration of microtransit and buses (Fig. 8c).

6.4. Active vehicles: State space and reward formulation evaluation

This section evaluates the efficacy of active vehicles (AVs) in the state space and reward function formulation. Intuitively, AVs are conceptualized to provide a more accurate representation of the traffic dynamics. Accounting for vehicles and people not




Fig. 8. (a) Distributions of the maximum queue lengths observed on the incoming lanes for all intersections at every action interval on the 4×4 grid network as generated by MPLight and HumanLight. (b),(c) Maximum vehicle queues observed for the light and high HOV adoption scenarios respectively: In the light scenarios, maximum queues deteriorate on intersections serving the heaviest public transit load, while in the high scenario they improve throughout. Per intersection maximum queues for the other two scenarios are provided in Fig. A.2.

Table 4

Formulations explored around the efficacy of the active vehicles consideration in the state space and reward function.

Formulation	Description	State space embedding size	Person pressure computed from
All vehicles	Person counts drawn from all vehicles	$ L_{in}^I + L_{out}^I + 1$	All vehicles
Active vehicles	Person counts drawn from active vehicles	$ L_{in}^I + L_{out}^I + 1$	Active vehicles
All & active vehicle	Person counts drawn from both all and active vehicles	$2 \cdot L_{in}^I + L_{out}^I + 1$	Active vehicles

Table 5

Impact of active vehicle consideration on corridors of different segment lengths. Higher benefits are observed for longer road segments.

Formulation	Delay (s)				Queue			
	Vehicle	% Change	Person	% Change	Vehicle	% Change	Person	% Change
Corridor 1×3 (300 m segments)								
All vehicles	67.99 (2.40)	Baseline	49.98 (2.08)	Baseline	31.70 (1.06)	Baseline	45.01 (2.47)	Baseline
Active vehicles	61.10 (4.05)	-10.13%	45.07 (4.24)	-9.82%	28.46 (1.97)	-10.22%	40.12 (4.87)	-10.64%
All & active vehicle	61.49 (2.82)	-9.56%	45.25 (2.30)	-9.46%	28.64 (1.42)	-9.65%	40.18 (2.51)	-10.73%
Corridor 1×3 (800 m segments)								
All vehicles	116.10 (5.69)	Baseline	93.81 (6.98)	Baseline	55.23 (2.51)	Baseline	87.11 (13.64)	Baseline
Active vehicles	59.95 (3.27)	-48.36%	44.35 (2.76)	-52.72%	28.00 (1.69)	-49.30%	38.67 (3.25)	-55.61%
All & active vehicle	60.35 (2.23)	-48.02%	44.69 (1.89)	-52.36%	28.05 (1.13)	-49.21%	38.16 (2.02)	-56.19%

in range of the intersection within the time horizon of the action interval may potentially degrade the performance of control algorithms. Although vehicle counts across the entirety of lanes are extensively used in literature to create embeddings describing the traffic conditions of road segments, passenger loads can significantly fluctuate across vehicles, especially the more the world shifts towards a multimodal reality. These variations can make a substantial difference in the controller's decision making process when determining which phase needs to be prioritized on a person-level approach (Fig. 1). To demonstrate the motivation and the benefits of considering the active vehicles in our person-level approach, we explore three different formulations. In all three alternatives, the current traffic signal phase ϕ^I is included in the state space, while the way people counts in the incoming and outgoing lanes are encoded varies as described in Table 4. All alternatives are illustrated in a corridor of three intersections with the low HOV adoption scenario, while the demand is fixed nearly at capacity as described in Section 5.

As shown in Eqs. (6) and (7), a vehicle being considered as active depends on their relative location on the road segment, their current speed and maximum speed and acceleration. For the set values of vehicle features (Table 1), the active range for a vehicle leaving the intersection under zero speed is 80.25 m and under maximum speed is 111.11 m. These translate to approximately 1/3 of the typical road segment length (300 m) used in our experiments to be within the active vehicle range. To assess the sensitivity of AVs to the input parameters, we perform the three evaluation alternatives on an additional network structure with longer segments (800 m). We observe that all formulations converge after approximately 160 episodes (Fig. 9). When considering person counts from active vehicles versus all vehicles, in the 300 m segment corridor around 10% decreases in vehicle and person delays and queues are achieved (Table 5). The percentage improvements of those metrics almost reach or even exceed 50% in the 800 m segment corridor, while the actual values of delays and queues are similar to the 300 m segment corridor. Those results are in alignment with the intuitive expectation that the longer the road segments the higher the impact of the consideration of AVs would be. Table 5 illustrates how the consideration of both active and all vehicles in the state space achieves similar performance as




Fig. 9. Person delays (top) and queues (down) across the three representations of state space for the consideration of active vehicles. Two network structures are tested, one with segments of length 300 m (left) and one with segments of length 800 m (right).

the active vehicle representation alone without further benefits to justify the additional complexity in the embedding. The same experiments were conducted in the same 1×3 corridors under lighter traffic with an average intersection $v/c \approx 0.65$ (Table A.4). Results are consistent with Table 5. We do notice though that the improvements for the corridor with 800 m road segments under light traffic flatten out at approximately 20%, hinting that the benefits of AVs are higher for heavier traffic demand.

6.5. Regulating the aggressiveness of HOV prioritization

Apart from enabling reduced passenger travel times and socially equitable green time allocation at signalized intersections, HumanLight empowers policymakers and traffic engineers to regulate the aggressiveness in the prioritization of the HOV fleet. We achieve this via a modification in the state embedding where vehicle occupancies are encoded. This way, the travel time benefits across vehicle types of different occupancies is fully in the hands of the system operator rather than generated through a black-box approach.

This concept is illustrated in Fig. 10a, where five different encodings of person counts are displayed. The dark green representation ($e = 1$) depicts the person count encoding as defined in Section 3. With e decreasing, the impact of microtransit and public transit in the reward function compared to SOVs and carpools is diminishing. On the other extreme, the red representation ($e = 0$) completely cancels out the person count function, turning it into the vehicle counting function, treating all vehicles the same.

We evaluate how HumanLight can regulate the prioritization aggressiveness of vehicles of different occupancies in our high adoption mode share scenario. The high HOV adoption scenario was selected to provide higher number of samples (vehicles) in the microtransit and public transit categories, making the line fitting more robust by definition. Results were reproduced and remain consistent across all mode share scenarios. The evaluation is performed in the single intersection setting. That way every vehicle in our system follows a route of the same length, whether it includes a turn or not, thus making their travel times and delays comparable. Additionally, there is no bias originating from bus pre-specified routes. For example, in the corridor case where buses run along the corridor, vehicles riding from one end to the other without exiting are more likely to get a higher allocation of green times compared to turning vehicles.

Fig. 10b illustrates how the prioritization of HOVs intensifies with the increase of parameter e . A more detailed snapshot of those distributions on the last episode of training can be found in Fig. A.3. From evaluating the fitted lines for $e = 0.75$ and $e = 1.00$ in Fig. 10b, we observe that the incline of the travel time vs vehicle occupancy fitted line does not change much in absolute value. That is because HOVs cannot be further over-prioritized. Thus, the two scenarios return similar handling of inter-vehicles conflicts during arrivals at the intersection. As intuitively expected, the fitted line for $e = 0$ has an incline of 0.01, corresponding to a vehicle-level optimization approach. We need to stress that the travel time of 78 s corresponds to a delay of 24 s from the minimum free-flow




Fig. 10. (a) Five modified encodings of the person count function for the different vehicle types. (b) Linear least-squares regression fitted on means of travel times for all vehicle types.




Fig. 11. Evolution of average person delay and queue for the different values of discount factor γ during episodic training for the high traffic demand configuration.

travel time, which is approximately 5 s less than the vehicle delay of MPLight (Table A.1) in the same scenario. That improvement is a result of the active vehicle consideration. Overall, operators of HumanLight should take input from behavioral studies that evaluate the travel time elasticity of commuter mode choice. The aggressiveness in HOV prioritization should be scientifically informed to incentivize riders to mode shift. Over-penalizing SOVs' travel times must not be discouraging to commuters.

6.6. Impact of discount factor γ on phase profile

From the hyper-parameters tuned to optimize the performance of our reinforcement learning model, we conducted the most experiments on the discount factor γ . Apart from a significant impact on performance, the generated policies rendered different phase profiles. This section aims to advise system operators on the impact of prioritizing immediate versus long-term rewards. Four values are tested to capture the full spectrum of the effect, ranging from the fully myopic $\gamma = 0$ to $\gamma = 0.90$ that heavily weighs future rewards as described in Eq. (8). Similar to the active vehicle analysis (Section 6.4), we perform our evaluation on a corridor of three intersections with 300 m long segments to grasp any network level effects. We illustrate results under heavy traffic demand (average intersection v/c exceeding 0.92).

In terms of convergence, for the lowest γ value ($\gamma = 0.00$) we observe significant oscillations even after 200 episodes of training (Fig. 11). Performance-wise, the best performing value for hyper-parameter γ is 0.60 is minimizing delays and queues both in absolute value as well as standard deviation (Table 6).

When evaluating the impact of the discount factor γ to the profile of the generated phases (Table 7 and Fig. 12), we notice that the more short-sighted the controller is (smaller γ values), less phase changes are occurring translating to longer durations

Table 6
Impact of discount factor γ on performance metrics.

γ	Average delay		Average queue	
	Vehicle (s)	Person (s)	Vehicle (vehicles)	Person (persons)
0.00	92.59 (36.45)	75.59 (36.18)	43.02 (15.25)	61.12 (18.71)
0.30	70.52 (11.58)	53.41 (10.55)	33.07 (5.31)	48.15 (9.46)
0.60	61.10 (4.05)	45.07 (4.24)	28.46 (1.97)	40.12 (4.87)
0.90	65.40 (7.85)	51.00 (7.70)	31.37 (5.38)	44.25 (5.45)

Table 7
Sensitivity analysis on the number of phase changes per hour for different values of the discount factor γ on two different traffic demand levels at a synthetic corridor 1×3 of intersections.

Corridor 1×3 (300 m segments)		Discount factor γ			
		0.00	0.30	0.60	0.90
Average number of	$v/c \approx 0.92$	56.0 (12.0)	94.6 (16.3)	121.9 (10.2)	149.2 (9.8)
phase changes per hour	$v/c \approx 0.65$	144.3 (28.0)	151.00 (9.1)	151.33 (23.1)	175.00 (17.3)




Fig. 12. Phase profile for different values of γ under heavy traffic: A myopic controller, prone to postpone the yellow and all red period will result in less phase changes and higher durations.

of the imposed phases. The reason behind this pattern originates from the imposition of the 5 s. yellow and all red period before the initiation of a new (different from the current) phase. A myopic and greedier algorithm will learn to postpone that penalty period where no vehicles are allowed to traverse the intersection of control. Instead, when considering a longer time horizon in the optimization cost function, the controller will be more eager to endure the inevitable penalty period of a phase switch at the appropriate (as dictated by the cost function) point in time. As a benchmark, we note that a typical 8-phase fixed controller with

120 s signal length results in $3600/120 \cdot 8 = 240$ changes per hour. Therefore, the increased number of phase changes (Table 7) with the increase of γ should not pose any threats to pedestrian crossings waiting time but rather the opposite. Although in this paper HumanLight adopts the value of $\gamma = 0.60$ thanks to its superior performance and balanced phase profile, system operators may chose to sacrifice travel time gains for a desired phase profile. Although HumanLight is built on the acyclic phase scheme, it is fully compatible with a fixed-sequence (cyclic) phase configuration where the algorithm would tune phase extensions and shortenings.

For consistency, the evaluation is also performed under light traffic (average intersection $v/c \approx 0.65$). The results are provided in the Appendix (Table A.5 and Figs. A.4&A.5). Under light traffic, tuning γ seems to be less important as $\gamma = 0.60$ and $\gamma = 0.90$ perform equally well. Intuitively, lighter traffic implies the controller will require less consecutive action intervals to fully accommodate the traffic of each direction. This is validated when looking at the number of phase changes per hour (Table 7), where we observe that reduced vehicle density leads to increased phase changes. The phase duration profile varies even less than in the heavy traffic case, as long term implications are not as significant.

7. Conclusions

This study presents HumanLight, a decentralized adaptive signal control algorithm, which to the best of our knowledge is the first person-based approach founded on reinforcement learning in the context of network level control. Our traffic management solution requires fully developed V2I communication technology to optimize people's throughput at intersections. HumanLight is designed to equitably allocate green times at intersections. By rewarding HOV commuters with travel time savings for their efforts to merge into a single ride, HumanLight helps democratize urban traffic.

Considering a multimodal vehicle split among private vehicles, pooled rides, microtransit and public transit at various scenarios of HOV adoption, we showcase significant headroom compared to a state-of-the-art RL-based vehicle-level optimization controller. In a 4×4 intersection grid, improvements on person delays and queues exceed 25% for the lowest HOV adoption scenario which considers no microtransit vehicles and public transit adoption as low as in the COVID-19 pandemic era. For moderate and high HOV adoption, improvements on those same metrics reached over 40% and 55% respectively considering a $\sim 26\%$ and $\sim 41\%$ shift from single occupancy vehicles to HOV alternatives respectively. As travel time benefits increase along with HOV adoption, HumanLight can invigorate ridesharing and public transit systems to attract the travel demand they truly merit in sustainable and multimodal urban environments.

HumanLight formulates the reward function by extending the transportation theory inspired concept of pressure to the person level. To handle the highly variant occupancy profiles in multimodal urban environments, we introduce active vehicles as vehicles in proximity to the controlled intersections within the action interval window. Through systematic experiments on varying traffic flows and network structures, we showcase how generating the state space and computing rewards based on active vehicles, enables better decision making for optimizing people's throughput at intersections.

Apart from rendering reduced passenger travel times and socially equitable green time allocation, HumanLight allows policymakers and traffic engineers to regulate the aggressiveness in the prioritization of the HOV fleet. We illustrate how via a modification in the state embedding, the benefits for the different vehicle types can be tuned, ideally with input from behavioral studies on the travel time elasticity of mode choice. Furthermore, we explore the impact of the discount factor, which determines the importance of future rewards, on performance as well as the generated phase profile. System operators can be informed on the expected phase change and duration patterns which are critical aspects of acyclic signal controllers as they affect pedestrian waiting times. Finally, to assure that HOV prioritization does not lead to excessive queues in the non-prioritized directions of traffic, maximum vehicle queues from all incoming lanes are evaluated in the 4×4 grid setting. Across all scenarios, vehicle queues remain far lower than the segments' capacity at jam density, posing no threats to the operational efficiency of the signalized intersections.

8. Future directions

The field of human-centric RL-based traffic signal control is currently under-explored in literature, empowering us to present various directions to be considered for future research. Thanks to the decentralized architecture, HumanLight is a scalable traffic management solution. The algorithmic performance still needs to be evaluated on city-level networks with heterogeneous road and intersection structures. Advancements in high performance computing and simulation (Chan et al., 2018) have made large-scale algorithmic applications possible. As traffic volumes may significantly vary per location, more elaborate designs for coordination and cooperation among neighboring intersections or training multiple controllers may potentially yield better results.

From the demand side, open work contains an exploration at the metropolitan level considering advanced matching algorithms for pooling of trips and quantifying the gains in travel time for different distributions of individuals' walking radius flexibility to join a pooled ride, either at its origin or on route. For those different levels of adoption, it would interesting to assess the benefits of our human-centric controller in terms of emissions. Although results demonstrate potential as single occupancy vehicles are taken off the roads and HOVs are receiving priority, thus performing less stops, a follow up study quantifying emissions would be of great value. The sensitivity in the framework's effectiveness for varying penetration rates of vehicles with connected technology capabilities should also be evaluated. In cases where public transit occupancy data is not available, we recommend treating buses as full to ensure they remain the most prioritized vehicle type thus always providing commuters with riding incentives even early on routes.

From a planning perspective, the interaction of person-based traffic signal control with dedicated bus lanes needs to be investigated for policymakers to identify the optimal allocation of resources and whether these two HOV prioritization strategies

should be applied simultaneously. With HumanLight establishing near free-flow travel times for transit vehicles, the dedicated bus lanes can be freed and used as dedicated bike lanes or pedestrian safe-spaces, paving the way towards a more sustainable urban environment. As opposed to HOV lanes, our solution also achieves democratization of rides independent of the congestion levels, providing people with incentives to pool even at low traffic demand scenarios.

HumanLight allows the flexibility to balance even more complex priorities of multimodal traffic. Bike lanes and pseudo pedestrian lanes, modeling sidewalks, can be included in the codification of the transportation network if demands from these travel means were to be available. Apart from prioritizing bikes and pedestrians via higher weights, HumanLight can accommodate emergency vehicles and paratransit. By assigning ambulances and fire trucks extremely high occupancy values in cases of emergency, these vehicles would receive the highest level of prioritization. Although establishing transit adherence to schedule can be embedded in our reward functions, we believe that if the technology is adopted the bus line schedules should be adjusted accordingly to anticipate the near optimal travel time along the bus route resolving all potential issues of punctuality and bus bunching. With the transit network operating with higher efficiency, the schedule could also become denser offering riders an even higher level of service. In a similar manner, HumanLight can assist policymakers of prosperous urban societies to incentivize ridership of paratransit shuttles. Paratransit provides transportation to people whose disabilities prevent them from riding traditional fixed-route transit service. By accordingly adjusting the weights of paratransit vehicles in the space embedding, HumanLight can make paratransit more attractive for people with disabilities who face widespread lack of accessibility to built environments and roads. By prioritizing these vehicles, cities can foster participation and inclusion of all members of society.

Further research could explore the effectiveness of our method under cyclic traffic signal configurations. Although research has shown that acyclic phase transitions consistently yield better performance than cyclic phase transitions (Kanis et al., 2021), the fixed sequence provides with more reliability for pedestrian serving across all directions. Other approaches worth exploring include the incorporation of pedestrian waiting time in the reward function where elongated waiting times would be penalized, or the development of an acyclic controller with a minimum green time constraint for some traffic movements. Finally, the impact of the action interval on the benefits from the consideration of active vehicles needs to be investigated. A larger interval will increase the active range across all vehicles, and therefore more vehicles will be accounted for in the state representation.

Code and data availability

The code, data and relevant documentation are available on Github at <https://github.com/DimVlachogiannis/HumanLight.git>.

Funding

This work was sponsored by the U.S. Department of Energy (DOE) Vehicle Technologies Office (VTO) under the Big Data Solutions for Mobility Program, an initiative of the Energy Efficient Mobility Systems (EEMS) Program. The following DOE Office of Energy Efficiency and Renewable Energy (EERE) manager played important roles in establishing the project concept, advancing implementation, and providing ongoing guidance: Prasad Gupte. This scientific paper was additionally supported by the Onassis Foundation—Scholarship ID: F ZQ-010/2020–2021.

CRediT authorship contribution statement

Dimitris M. Vlachogiannis: Conceptualization, Methodology, Software, Data curation, Formal analysis, Validation, Writing – original draft, review & editing. **Hua Wei:** Methodology, Validation, Writing – review & editing, Supervision. **Scott Moura:** Writing – review & editing, Supervision. **Jane Macfarlane:** Conceptualization, Writing – review & editing, Supervision, Project administration.

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.

Data availability

A repository with the code, data, and documentation has been referenced in the manuscript.

Acknowledgments

We sincerely thank Professor Alex Skabardonis and Professor Daniel Rodriguez for their insightful comments and support.

Appendix

See [Figs. A.1–A.5](#) and [Tables A.1–A.5](#).




Fig. A.1. Convergence trend of average person queues with HumanLight throughout the learning episodes across all HOV adoption scenarios.

Table A.1

Metric evaluation across all scenarios in the single intersection set of experiments.

Mode share scenario	Controller	Average vehicle delay (s)	Average person delay (s)	Average vehicle queue (vehicles)	Average person queue (persons)
Low HOV adoption	Fixed (Websters)	235.76	231.86	98.09	105.75
	SOTL	93.38	95.11	46.53	52.23
	MaxPressure	62.56	62.36	31.09	33.90
	MPLight with FRAP + Pressure	58.75 (2.05)	59.55 (2.29)	29.20 (1.14)	32.47 (1.34)
	HumanLight	61.25 (3.21)	58.95 (2.29)	30.56 (1.78)	32.11 (1.41)
	% Improvement from SOTA	-4.26%	1.01%	-4.66%	1.11%
Little HOV adoption	Fixed (Websters)	169.65	174.23	70.20	90.20
	SOTL	68.63	69.70	29.40	38.25
	MaxPressure	63.36	64.01	31.15	39.98
	MPLight with FRAP + Pressure	47.03 (2.63)	46.68 (2.41)	20.16 (1.78)	25.54 (1.39)
	HumanLight	47.95 (3.02)	40.53 (2.56)	20.54 (1.63)	21.77 (1.27)
	% Improvement from SOTA	-1.96%	13.17%	-1.88%	14.76%
Moderate HOV adoption	Fixed (Websters)	109.07	116.51	40.36	67.14
	SOTL	54.86	52.86	18.52	27.75
	MaxPressure	56.71	52.83	23.74	34.20
	MPLight with FRAP + Pressure	38.48 (0.93)	40.10 (2.37)	13.05 (0.34)	21.33 (1.41)
	HumanLight	35.36 (1.20)	25.44 (1.03)	11.82 (0.45)	12.63 (0.53)
	% Improvement from SOTA	8.11%	36.56%	9.43%	40.79%
High HOV adoption	Fixed (Websters)	64.64	61.43	18.17	35.59
	SOTL	39.62	36.50	9.81	18.67
	MaxPressure	53.05	57.76	17.51	39.59
	MPLight with FRAP + Pressure	29.04 (0.95)	28.89 (1.58)	7.23 (0.26)	14.88 (1.01)
	HumanLight	29.31 (1.23)	16.56 (0.77)	7.45 (0.38)	7.96 (0.49)
	% Improvement from SOTA	-0.93%	42.68%	-3.04%	46.51%

Table A.2

Metric evaluation across all scenarios in the 1 × 6 corridor set of experiments.

Mode share scenario	Controller	Average vehicle delay (s)	Average person delay (s)	Average vehicle queue (vehicles)	Average person queue (persons)
Low HOV adoption	Fixed (Websters)	352.49	368.53	84.89	100.16
	SOTL	314.97	337.22	101.53	119.69
	MaxPressure	68.83	73.84	26.23	30.73
	MPLight with FRAP + Pressure	70.22 (4.18)	73.41 (4.59)	25.73 (1.92)	27.82 (1.87)
	HumanLight	54.78 (3.45)	53.79 (3.62)	19.81 (1.20)	21.24 (1.85)
	% Improvement from SOTA	20.41%	26.73%	23.01%	23.65%
Little HOV adoption	Fixed (Websters)	237.16	268.08	63.84	92.93
	SOTL	120.18	153.51	38.70	63.32
	MaxPressure	66.96	76.77	24.78	36.23
	MPLight with FRAP + Pressure	48.23 (1.77)	52.43 (1.88)	15.67 (0.62)	21.35 (1.43)
	HumanLight	41.16 (1.80)	37.37 (1.82)	12.95 (0.66)	14.52 (0.79)
	% Improvement from SOTA	14.66%	28.72%	17.36%	31.99%

(continued on next page)

Table A.2 (continued).

Mode share scenario	Controller	Average vehicle delay (s)	Average person delay (s)	Average vehicle queue (vehicles)	Average person queue (persons)
Moderate HOV adoption	Fixed (Websters)	116.06	151.72	32.17	67.73
	SOTL	109.07	116.51	40.36	67.14
	MaxPressure	59.95	70.59	19.07	35.01
	MPLight with FRAP + Pressure	37.11 (0.75)	42.28 (1.34)	9.75 (0.24)	16.28 (1.02)
	HumanLight	28.77 (1.41)	22.92 (1.24)	7.18 (0.45)	8.35 (0.54)
	% Improvement from SOTA	22.47%	45.79%	26.36%	48.71%
High HOV adoption	Fixed (Websters)	76.92	117.54	17.60	54.80
	SOTL	67.04	100.72	13.69	41.44
	MaxPressure	55.31	66.04	14.41	35.00
	MPLight with FRAP + Pressure	28.49 (0.98)	31.13 (1.24)	5.70 (0.26)	12.75 (1.55)
	HumanLight	23.71 (2.11)	16.30 (1.26)	4.61 (0.51)	5.64 (0.68)
	% Improvement from SOTA	16.78%	47.64%	19.12%	55.76%

Table A.3Metric evaluation across all scenarios in the 4×4 grid traffic network configuration.

Mode share scenario	Controller	Average vehicle delay (s)	Average person delay (s)	Average vehicle queue (vehicles)	Average person queue (persons)
Low HOV adoption	Fixed (Websters)	500.86	512.51	55.95	62.98
	SOTL	446.3	462.9	69.51	79.49
	MaxPressure	159.39	162.97	30.74	34.49
	MPLight with FRAP + Pressure	143.81 (5.62)	146.29 (5.63)	25.25 (0.78)	27.82 (1.36)
	HumanLight	110.40 (8.23)	107.00 (8.20)	19.40 (1.63)	20.50 (4.10)
	% Improvement from SOTA	23.23%	26.86%	23.17%	26.31%
Little HOV adoption	Fixed (Websters)	410.14	427.24	49.51	66.24
	SOTL	265.09	280.01	40.33	54.34
	MaxPressure	139.92	141.57	24.94	32.24
	MPLight with FRAP + Pressure	100.94 (2.14)	103.23 (2.38)	15.83 (0.31)	20.14 (0.98)
	HumanLight	85.20 (3.08)	74.80 (2.68)	13.10 (0.54)	14.40 (0.92)
	% Improvement from SOTA	15.59%	27.54%	17.25%	28.50%
Moderate HOV adoption	Fixed (Websters)	239.82	265.94	31.42	54.33
	SOTL	179.23	197.73	22.5	38.39
	MaxPressure	122.33	128.66	18.99	30.92
	MPLight with FRAP + Pressure	76.01 (0.91)	77.86 (1.48)	9.81 (0.13)	15.54 (0.70)
	HumanLight	59.80 (2.60)	45.30 (2.03)	7.50 (0.37)	8.40 (0.55)
	% Improvement from SOTA	21.33%	41.82%	23.55%	45.95%
High HOV adoption	Fixed (Websters)	169.11	193.41	18.26	41.88
	SOTL	136.35	154.19	13.41	30.12
	MaxPressure	111.66	113.09	14.28	28.86
	MPLight with FRAP + Pressure	57.44 (1.11)	59.78 (1.45)	5.76 (0.13)	11.58 (0.88)
	HumanLight	40.50 (2.84)	25.70 (1.80)	3.90 (0.31)	4.50 (0.38)
	% Improvement from SOTA	29.49%	57.01%	32.29%	61.14%

Table A.4

Impact of active vehicle consideration on experimental corridors under light traffic demand.

	Delay (s)				Queue			
	Vehicle	% Change	Person	% Change	Vehicle	% Change	Person	% Change
Scenario name	Corridor 1 $\times$ 3 (300 m segments)							
All vehicles	43.27 (1.38)	Baseline	29.31 (0.86)	Baseline	15.46 (0.55)	Baseline	17.36 (0.58)	Baseline
Active vehicles	35.95 (2.09)	-16.92%	24.90 (1.41)	-15.05%	12.57 (0.88)	-18.69%	15.00 (1.18)	-13.59%
Scenario name	Corridor 1 $\times$ 3 (800 m segments)							
All vehicles	44.17 (1.72)	Baseline	29.88 (1.27)	Baseline	15.88 (0.69)	Baseline	17.80 (0.84)	Baseline
Active vehicles	35.69 (1.86)	-19.20%	24.56 (1.40)	-17.80%	12.48 (0.75)	-21.41%	14.07 (0.93)	-20.96%




Fig. A.2. Maximum vehicle queue per lane comparison between MPLight and HumanLight: Deteriorations are observed at the low HOV adoption scenario on intersections serving the heaviest public transit load (left). Improvements throughout the grid are reported in the moderate HOV adoption scenario (right).




Fig. A.3. For episode 200, the penalization of SOVs in terms of travel time is becoming less intense with the increase of parameter e .

Table A.5

Impact of discount factor γ on performance metrics under light traffic demand.

Gamma	Average delay		Average queue	
	Vehicle (s)	Person (s)	Vehicle (vehicles)	Person (persons)
0.00	50.58 (60.80)	39.52 (56.21)	17.19 (16.95)	22.88 (23.88)
0.30	37.57 (1.80)	26.32 (1.69)	13.28 (0.72)	15.73 (1.17)
0.60	35.95 (2.09)	24.40 (1.41)	12.57 (0.88)	14.70 (1.18)
0.90	34.01 (0.81)	24.10 (0.59)	11.76 (0.30)	14.52 (0.64)




Fig. A.4. Evolution of average person delay and queue for the different values of discount factor γ during episodic training for the light traffic demand configuration.




Fig. A.5. Phase profile for different values of γ under light traffic.

References

- Arel, I., Liu, C., Urbanik, T., Kohls, A.G., 2010. Reinforcement learning-based multi-agent system for network traffic signal control. *IET Intell. Transp. Syst.* 4 (2), 128–135.
- Aslani, M., Mesgari, M.S., Wiering, M., 2017. Adaptive traffic signal control with actor-critic methods in a real-world traffic network with different traffic disruption events. *Transp. Res. C* 85, 732–752.
- Bakker, B., Whiteson, S., Kester, L., Groen, F.C., 2010. *Traffic Light Control by Multiagent Reinforcement Learning Systems*. Springer.
- Bardaka, E., Hajibabai, L., Singh, M.P., 2020. Reimagining ride sharing: Efficient, equitable, sustainable public microtransit. *IEEE Internet Comput.* 24 (5), 38–44.
- Casas, N., 2017. Deep deterministic policy gradient for urban traffic light control. *arXiv preprint arXiv:1703.09035*.
- Chan, C., Wang, B., Bachan, J., Macfarlane, J., 2018. Mobiliti: Scalable transportation simulation using high-performance parallel computing. In: 2018 21st International Conference on Intelligent Transportation Systems. ITSC, IEEE, pp. 634–641.
- Chen, C., Wei, H., Xu, N., Zheng, G., Yang, M., Xiong, Y., Xu, K., Li, Z., 2020. Toward a thousand lights: Decentralized deep reinforcement learning for large-scale traffic signal control. In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 34, pp. 3414–3421.
- Christofa, E., Ampountolas, K., Skabardonis, A., 2016. Arterial traffic signal optimization: A person-based approach. *Transp. Res. C* 66, 27–47.
- Christofa, E., Papamichail, I., Skabardonis, A., 2013. Person-based traffic responsive signal control optimization. *IEEE Trans. Intell. Transp. Syst.* 14 (3), 1278–1289.
- Christofa, E., Skabardonis, A., 2011. Traffic signal optimization with application of transit signal priority to an isolated intersection. *Transp. Res. Rec.* 2259 (1), 192–201.
- Chu, T., Wang, J., Codecà, L., Li, Z., 2019. Multi-agent deep reinforcement learning for large-scale traffic signal control. *IEEE Trans. Intell. Transp. Syst.* 21 (3), 1086–1095.
- Cohen, M.C., Fitzer, M.-D., Ratzon, A., Sasson, R., 2021. Incentivizing commuters to carpool: A large field experiment with waze. *Manuf. Serv. Oper. Manage.*
- Cools, S.-B., Gershenson, C., D'Hooghe, B., 2013. Self-organizing traffic lights: A realistic simulation. In: *Advances in Applied Self-Organizing Systems*. Springer, pp. 45–55.
- Cowan, G., Monteleone, I., Bellisario, J., 2021. The Megaregional Case for a New Transbay Rail Crossing. Tech. Rep., Bay Area Council Economic Institute, San Francisco, CA, p. 56, URL http://www.bayareaeconomy.org/wp-content/uploads/2021/01/FINAL_NTRCmegaregion_1.27.2021.pdf.
- Diab, E.I., Badami, M.G., El-Geneidy, A.M., 2015. Bus transit service reliability and improvement strategies: Integrating the perspectives of passengers and transit agencies in North America. *Transp. Rev.* 35 (3), 292–328.

- Dietterich, T.G., et al., 1998. The MAXQ method for hierarchical reinforcement learning. In: ICML. Vol. 98, pp. 118–126.
- Du, W., Ye, J., Gu, J., Li, J., Wei, H., Wang, G., 2022. SafeLight: A reinforcement learning method toward collision-free traffic signal control. arXiv preprint arXiv:2211.10871.
- El-Tantawy, S., Abdulhai, B., 2010. An agent-based learning towards decentralized and coordinated traffic signal control. In: 13th International IEEE Conference on Intelligent Transportation Systems. IEEE, pp. 665–670.
- El-Tantawy, S., Abdulhai, B., Abdelgawad, H., 2013. Multiagent reinforcement learning for integrated network of adaptive traffic signal controllers (MARLIN-ATSC): methodology and large-scale application on downtown toronto. IEEE Trans. Intell. Transp. Syst. 14 (3), 1140–1150.
- Farid, Y.Z., Christofa, E., 2017. Real-Time Signal Control with Transit Priority Window: A Person-Based Approach. Tech. Rep., Transportation Research Board 96th Annual Meeting.
- Gershenson, C., 2004. Self-organizing traffic lights. arXiv preprint nlin/0411066.
- Giuliano, G., Levine, D.W., Teal, R.F., 1990. Impact of high occupancy vehicle lanes on carpooling behavior. Transportation 17 (2), 159–177.
- Henderson, P., Islam, R., Bachman, P., Pineau, J., Precup, D., Meger, D., 2018. Deep reinforcement learning that matters. In: Proceedings of the AAAI Conference on Artificial Intelligence. Vol. 32, pp. 3207–3214.
- Hu, J., Park, B.B., Lee, Y.-J., 2015. Coordinated transit signal priority supporting transit progression under connected vehicle technology. Transp. Res. C 55, 393–408.
- Kanis, S., Samson, L., Bloembergen, D., Bakker, T., 2021. Back to basics: Deep reinforcement learning in traffic signal control. arXiv preprint arXiv:2109.07180.
- Koonce, P., Rodegerdts, L., 2008. Traffic Signal Timing Manual. Tech. Rep., Federal Highway Administration, United States.
- Kuyer, L., Whiteson, S., Bakker, B., Vlassis, N., 2008. Multiagent reinforcement learning for urban traffic control using coordination graphs. In: Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2008, Antwerp, Belgium, September 15–19, 2008, Proceedings, Part I 19. Springer, pp. 656–671.
- Lazarus, J.R., Caicedo, J.D., Bayen, A.M., Shaheen, S.A., 2021. To pool or not to pool? Understanding opportunities, challenges, and equity considerations to expanding the market for pooling. Transp. Res. A 148, 199–222.
- Lioris, J., Kurzhanskiy, A., Varaiya, P., 2016. Adaptive max pressure control of network of signalized intersections. IFAC-PapersOnLine 49 (22), 19–24.
- Long, M., Zou, X., Zhou, Y., Chung, E., 2022. Deep reinforcement learning for transit signal priority in a connected environment. Transp. Res. C 142, 103814.
- Malalgoda, N., Lim, S.H., 2019. Do transportation network companies reduce public transit use in the US? Transp. Res. A 130, 351–372.
- Mannion, P., Duggan, J., Howley, E., 2016. An experimental review of reinforcement learning algorithms for adaptive traffic signal control. Auton. Road Transp. Support Syst. 47–66.
- Martinez, F.J., Toh, C.K., Cano, J.-C., Calafate, C.T., Manzoni, P., 2011. A survey and comparative study of simulators for vehicular ad hoc networks (VANETs). Wirel. Commun. Mob. Comput. 11 (7), 813–828.
- Mei, H., Lei, X., Da, L., Shi, B., Wei, H., 2022. LibSignal: An open library for traffic signal control. arXiv preprint arXiv:2211.10649.
- Nishi, T., Otaki, K., Hayakawa, K., Yoshimura, T., 2018. Traffic signal control based on reinforcement learning with graph convolutional neural nets. In: 2018 21st International Conference on Intelligent Transportation Systems. ITSC, IEEE, pp. 877–883.
- Oroojlooy, A., Nazari, M., Hajinezhad, D., Silva, J., 2020. Attendlight: Universal attention-based reinforcement learning model for traffic signal control. Adv. Neural Inf. Process. Syst. 33, 4079–4090.
- Van der Pol, E., Olthoek, F.A., 2016. Coordinated deep reinforcement learners for traffic light control. In: Proceedings of Learning, Inference and Control of Multi-Agent Systems (at NIPS 2016). Vol. 8, pp. 21–38.
- Rasheed, F., Yau, K.-L.A., Noor, R.M., Wu, C., Low, Y.-C., 2020. Deep reinforcement learning for traffic signal control: A review. IEEE Access 8, 208016–208044.
- Roess, R.P., Prassas, E.S., McShane, W.R., 2004. Traffic Engineering. Pearson/Prentice Hall.
- Salkham, A.a., Cahill, V., 2010. Soile: A decentralized approach to optimization of fluctuating urban traffic using reinforcement learning. In: 13th International IEEE Conference on Intelligent Transportation Systems. IEEE, pp. 531–538.
- Schank, J., Huang, E., 2022. Microtransit: A Good Idea Just Got Even Better. San Jose State University.
- Shabestary, S.M.A., Abdulhai, B., 2018. Deep learning vs. discrete reinforcement learning for adaptive traffic signal control. In: 2018 21st International Conference on Intelligent Transportation Systems. ITSC, IEEE, pp. 286–293.
- Shabestary, S.M.A., Abdulhai, B., 2022. Adaptive traffic signal control with deep reinforcement learning and high dimensional sensory inputs: Case study and comprehensive sensitivity analyses. IEEE Trans. Intell. Transp. Syst. 23 (11), 20021–20035.
- Shabestary, S.M.A., Abdulhai, B., 2019. Multimodal intelligent deep (mind) traffic signal controller. In: 2019 IEEE Intelligent Transportation Systems Conference. ITSC, IEEE, pp. 4532–4539.
- Shaheen, S.A., Chan, N.D., Gaynor, T., 2016. Casual carpooling in the san Francisco Bay Area: Understanding user characteristics, behaviors, and motivations. Transp. Policy 51, 165–173.
- Shaheen, S., Cohen, A., 2019. Shared ride services in North America: definitions, impacts, and the future of pooling. Transp. Rev. 39 (4), 427–442.
- Sutton, R.S., Barto, A.G., 2018. Reinforcement Learning: An Introduction. MIT Press.
- United States Department of Transportation, 2021. Commute mode: Bureau of transportation statistics. URL <https://www.bts.gov/browse-statistical-products-and-data/state-transportation-statistics/commute-mode>.
- Urbanik, T., Tanaka, A., Lozner, B., Lindstrom, E., Lee, K., Quayle, S., Beaird, S., Tsoi, S., Ryus, P., Gettman, D., et al., 2015. Signal Timing Manual, vol. 1, Transportation Research Board Washington, DC.
- Varaiya, P., 2013. Max pressure control of a network of signalized intersections. Transp. Res. C 36, 177–195.
- Viegas, J., Martinez, L., Crist, P., Masterson, S., 2016. Shared mobility: Innovation for liveable cities. In: International Transport Forum's Corporate Partnership Board. pp. 1–56.
- Vlachogiannis, D.M., Moura, S., Macfarlane, J., 2023. Intersense: An xgboost model for traffic regulator identification at intersections through crowdsourced GPS data. Transp. Res. C 151, 104112. <http://dx.doi.org/10.1016/j.trc.2023.104112>, URL <https://www.sciencedirect.com/science/article/pii/S0968090X23001018>.
- Vlachogiannis, D.M., Vlahogianni, E.I., Golias, J., 2020. A reinforcement learning model for personalized driving policies identification. Int. J. Transp. Sci. Technol. 9 (4), 299–308.
- Wang, L., Ma, Z., Dong, C., Wang, H., 2022. Human-centric multimodal deep (HMD) traffic signal control. IET Intell. Transp. Syst.
- Wei, H., Chen, C., Zheng, G., Wu, K., Gayah, V., Xu, K., Li, Z., 2019a. Presslight: Learning max pressure control to coordinate traffic signals in arterial network. In: Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 1290–1298.
- Wei, H., Xu, N., Zhang, H., Zheng, G., Zang, X., Chen, C., Zhang, W., Zhu, Y., Xu, K., Li, Z., 2019b. Colight: Learning network-level cooperation for traffic signal control. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management. pp. 1913–1922.
- Wei, H., Zheng, G., Gayah, V., Li, Z., 2019c. A survey on traffic signal control methods. arXiv preprint arXiv:1904.08117.
- Wei, H., Zheng, G., Yao, H., Li, Z., 2018. Intellilight: A reinforcement learning approach for intelligent traffic light control. In: Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 2496–2505.
- Yang, S., Yang, B., Wong, H.-S., Kang, Z., 2019. Cooperative traffic signal control using multi-step return and off-policy asynchronous advantage actor-critic graph algorithm. Knowl.-Based Syst. 183, 104855.
- Yau, K.-L.A., Qadir, J., Khoo, H.L., Ling, M.H., Komisarczuk, P., 2017. A survey on reinforcement learning models and algorithms for traffic signal control. ACM Comput. Surv. 50 (3), 1–38.

- Yu, Z., Gayah, V.V., Christofa, E., 2017. Person-based optimization of signal timing: Accounting for flexible cycle lengths and uncertain transit vehicle arrival times. *Transp. Res. Rec.* 2620 (1), 31–42.
- Yu, Z., Gayah, V.V., Christofa, E., 2018. Implementing phase rotation in a person-based signal timing optimization framework. In: 2018 21st International Conference on Intelligent Transportation Systems. ITSC, IEEE, pp. 20–25.
- Zeng, X., Sun, X., Zhang, Y., Quadrioglio, L., 2015. Person-based adaptive priority signal control with connected-vehicle information. *Transp. Res. Rec.* 2487 (1), 78–87.
- Zhang, H., Feng, S., Liu, C., Ding, Y., Zhu, Y., Zhou, Z., Zhang, W., Yu, Y., Jin, H., Li, Z., 2019. Cityflow: A multi-agent reinforcement learning environment for large scale city traffic scenario. In: The World Wide Web Conference. pp. 3620–3624.
- Zhang, L., Wu, Q., Shen, J., Lü, L., Du, B., Wu, J., 2022. Expression might be enough: representing pressure and demand for reinforcement learning based traffic signal control. In: International Conference on Machine Learning. PMLR, pp. 26645–26654.
- Zheng, G., Xiong, Y., Zang, X., Feng, J., Wei, H., Zhang, H., Li, Y., Xu, K., Li, Z., 2019. Learning phase competition for traffic signal control. In: Proceedings of the 28th ACM International Conference on Information and Knowledge Management. pp. 1963–1972.