



MalLight: Influence-Aware Coordinated Traffic Signal Control for Traffic Signal Malfunctions

Qinchen Yang

Rutgers University

Piscataway, NJ, USA

qy129@rutgers.edu

Zejun Xie

Rutgers University

Piscataway, NJ, USA

zx180@scarletmail.rutgers.edu

Hua Wei

Arizona State University

Tempe, AZ, USA

hua.wei@asu.edu

Desheng Zhang

Rutgers University

Piscataway, NJ, USA

desheng@cs.rutgers.edu

Yu Yang

Lehigh University

Bethlehem, PA, USA

yuyang@lehigh.edu

ABSTRACT

Urban traffic is subject to disruptions that cause extended waiting time and safety issues at signalized intersections. While numerous studies have addressed the issue of intelligent traffic systems in the context of various disturbances, traffic signal malfunction, a common real-world occurrence with significant repercussions, has received comparatively limited attention. The primary objective of this research is to mitigate the adverse effects of traffic signal malfunction, such as traffic congestion and collision, by optimizing the control of neighboring functioning signals. To achieve this goal, this paper presents a novel traffic signal control framework (MalLight), which leverages an Influence-aware State Aggregation Module (ISAM) and an Influence-aware Reward Aggregation Module (IRAM) to achieve coordinated control of surrounding traffic signals. To the best of our knowledge, this study pioneers the application of a Reinforcement Learning(RL)-based approach to address the challenges posed by traffic signal malfunction. Empirical investigations conducted on real-world datasets substantiate the superior performance of our proposed methodology over conventional and deep learning-based alternatives in the presence of signal malfunction, with reduction of throughput alleviated by as much as 48.6%.

CCS CONCEPTS

- Computing methodologies → Artificial intelligence; • Applied computing → Transportation.

KEYWORDS

Deep reinforcement learning, traffic signal control, multi-agent system

ACM Reference Format:

Qinchen Yang, Zejun Xie, Hua Wei, Desheng Zhang, and Yu Yang. 2024. MalLight: Influence-Aware Coordinated Traffic Signal Control for Traffic Signal Malfunctions. In *Proceedings of the 33rd ACM International Conference*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CIKM '24, October 21–25, 2024, Boise, ID, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0436-9/24/10

<https://doi.org/10.1145/3627673.3679605>

on *Information and Knowledge Management (CIKM '24)*, October 21–25, 2024, Boise, ID, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3627673.3679605>

1 INTRODUCTION

Traffic signal malfunction refers to a situation when a traffic signal, also known as a traffic light, fails to function properly. One of the typical situations is that traffic lights are blackout caused by various reasons, such as flood [28, 52], hurricane [43], equipment aging [6], or human-made events, such as cyber-attacks [27, 16], traffic accidents and mismanagement [5]. When the traffic signal malfunction happens, it causes confusion and chaos on the road, leading to increased risk of accidents, traffic delays and traffic congestion. For example, we conducted a case study based on the traffic dataset from New York City (details in the Evaluation section), where traffic signals were set to be controlled by the FixedTime strategy [25], the most commonly used traffic signal control strategy. When the traffic signal at one intersection experiences a blackout, the average throughput of that intersection is reduced by nearly 30%. Moreover, the impacts also propagated to nearby intersections, such as intersections 1, 2, and 3-block away from the malfunctioning intersection, as shown in Figure 1. Similar observations are identified in other studies with increased travel time [15] and fatal accidents [37]. More importantly, these malfunctions occur more frequently than one might expect. For example, as shown in Figure 2, citizens from the City of Newark in New Jersey reported 94 complaints about traffic signal malfunctions at 67 different intersections between August 2020 and May 2022 on the Newark Connect platform [38] (i.e., a 311-like platform for citizens to report city issues), leading to three issues per month on average. Consequently, it is necessary to find a way that can mitigate the negative impacts (e.g., congestion) when traffic signal malfunctions occur.

The problem of traffic signal malfunctions has been studied in existing work [4, 44, 22]. In general, we can categorize these studies into two stages: detection and response. The malfunction detection mainly relies on road user report [4] and some recently developed intelligent signal systems to automatically detect signal malfunction [49, 44]. After detection, the ultimate response is to repair or replace the malfunctioning traffic signals. Unfortunately, due to limited resources (e.g., human resources and budgets), this response process can be time-consuming, with repairs sometimes taking months or even years to complete [50]. Therefore, immediate temporary responses become necessary. Current temporary

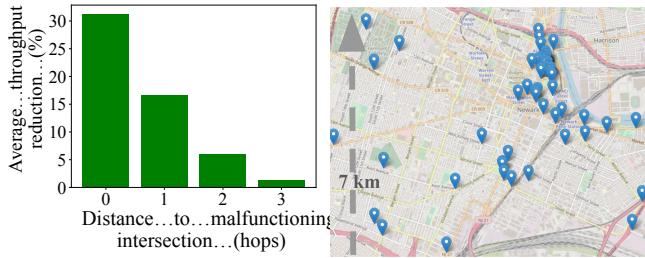


Figure 1: Throughput reduction at intersections around malfunctioning traffic signal

responses heavily rely on manual efforts, such as manual traffic guidance or Flash Mode [2]. However, abundant ground surveys have shown that Flash Mode is less effective [22], and manual traffic guidance also suffers from long time delays [33, 53]. Our work aims to provide a faster and more automated response after the malfunctions are detected and before significant human efforts get involved.

Our solution is inspired by the fundamental theory of traffic management from the perspective of demand (i.e., the number of vehicles that request to pass through an intersection in a unit of time) and capacity (i.e., the number of vehicles that an intersection can handle in the same unit of time) [1, 46]. Traffic congestion occurs when demand is higher than capacity. When a traffic signal malfunction occurs, i.e. the intersection changes from signalized to a blackout, capacity decreases [42], while demand remains relatively constant in the short term. This imbalanced relationship between demand and capacity motivates us to explore immediate responses to traffic signal malfunction: by carefully managing the networked traffic signals, especially those surrounding malfunctioning intersections, we can potentially control the traffic demand that enters the malfunctioning intersections in a desired manner, which leads to a more balanced demand and capacity.

However, coordinating intersections to control traffic in a desired manner is challenging, because (1) the impact from each intersection to malfunctioning ones is hard to measure. For example, one intersection has several routes to malfunctioning intersections. (2) Influence changes over time, such as during peak hours and non-peak hours. Recently, reinforcement learning (RL) has introduced a new possibility to coordinate traffic signal controls (TSC) [**presslight19**, 3, 14, 61, 71], where each agent controlling a single intersection can automatically learn to generate its actions to maximize its own benefits (e.g., the throughput of corresponding intersection). Since the reward function for each agent in these RL methods only depends on its own traffic condition, under a malfunction scenario, the agents of the well-functioning intersections are not aware of the specific needs of the malfunctioning intersections, even when its assistance benefits the whole system. Through experiments (Table 4, we observe that previous RL methods experience significant throughput reduction. For example, the IDQN method experiences a 42.8% reduction in throughput on the Hangzhou dataset, which is roughly 3.3 times higher than the 13% reduction observed with our model. Further, how the actions of each agent impact others is not well studied. To the best of our knowledge, none of the existing RL-based traffic signal controls can be applied in the malfunction setting.



Figure 2: Locations of complaints about traffic signal malfunction

In this work, we design MalLight, an RL-based traffic signal control method to mitigate the impacts of traffic signal malfunctions, by modeling the spatially and temporally varied influence between well-functioning and malfunctioning intersections. Specifically, MalLight models the dynamic influences through the special design of the state module and the reward module in RL. To model the spatially and temporally varied influence, we adapt a graph diffusion convolution network to integrate the network-level state with the local state of an individual agent so that each RL agent is aware of its influence on other intersections while controlling its local traffic conditions. To guide each well-functioning intersection to be aware of the reward associated with malfunctioning intersections, we introduce the graph diffusion process. This process enables each agent to dynamically integrate rewards from multiple malfunctioning intersections while learning their policy. In summary, the main contributions of this work are as follows:

- This is the first work that utilizes adaptive traffic signal control to build a resilient and responsive traffic system in the presence of traffic signal malfunction.
- We design a novel RL model, MalLight, based on modeling influences among intersections under signal malfunction. We incorporate the graph diffusion process into the state and reward design to model the influences and guide the training of MalLight, making it resistant to signal malfunction.
- Experiments on real-world and synthetic datasets show that MalLight outperforms both traditional and state-of-the-art (SoTA) RL-based models, achieving a remarkable alleviation in throughput reduction by up to 48.6%.

2 PRELIMINARIES

In this section, we introduce the preliminaries of the traffic signal control problem [63, 61] and extend it to our scenario considering traffic signal malfunction. In our setting, each intersection is controlled by one agent, which has its own local observations of its intersection as the state. An agent controls the traffic signal in its own intersection by deciding the signal phase at the next time interval to optimize the local traffic, such as maximizing the throughput. We use centralized learning with a decentralized execution manner to reduce the computational complexity of traditional centralized methods [60]. Specifically, we formulate our problem as a Markov Decision Process(MDP) with the following components $< S, O, M, A, r, \pi, \gamma >$.

- **Set of malfunctioning signals M :** M is a set of intersections with malfunctioning traffic signals.

• **State space S and observation space O :** Assuming there are N intersections in our road network and N corresponding agents that control the traffic signals located at the intersections. All possible traffic states from N intersections form the state space S and each agent has access to partial state $s \in S$, which forms the agent's own observation $o \in O$. The state s of each intersection includes current phase p (one of the 8 total phases illustrated in Figure 4) and the number of vehicles in each lane (i.e., a 12-element vector with each element representing a number of vehicles in lanes illustrated in Figure 3). In MalLight, observations from all intersections in the road network are fed into a state module, which outputs a global state. The global state is combined with local state s to form observation o for each agent.

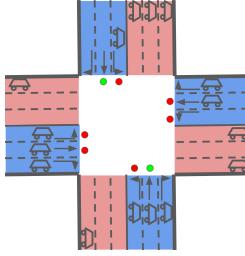


Figure 3: Intersection, incoming lanes(blue), outgoing lanes(red). Phase 3 is set.

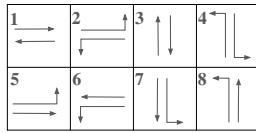


Figure 4: Eight phases (actions).

- **Action space A:** In the traffic signal control task, the i th agent generates an action a_i^t as the signal phase at time point t from its action space A_i for the following Δt period. Eight signal phases corresponding to eight actions are illustrated in Figure 4. The eight-phase setting is commonly used in real-world traffic signal control task [60]. We set Δt as 10 seconds based on the practice of the previous work [3]. Given the action a_i^t , the traffic signal chooses the corresponding phase p . For malfunctioning signal $m \in M$, a_m^t is set to be off denoted as p_{mal} .
- **Reward r :** Each RL agent obtains an immediate reward from the environment at time t by a rewarding function $S \times A_1 \times \dots \times A_N \rightarrow R$, where N is the number of intersections. We introduce the concept of *pressure* to measure the local reward for each intersection, which is the difference between the upstream and downstream queue length, indicating the inequivalence of the vehicle distribution [3]. More specifically, we denote the local reward of the i th intersection at time t as r_i^t , which is calculated as

$$r_i^t = -P_i = - \left| \sum_{l_{in} \in i} u_{i,l_{in}}^t - \sum_{l_{out} \in i} u_{i,l_{out}}^t \right| \quad (1)$$

where P_i represents the pressure of intersection i , $u_{i,l}^t$ is the queue length of lane l at intersection i at time t , l_{in} and l_{out} represent incoming and outgoing lane, respectively. For example, in Figure 3, the pressure of intersection is calculated as $|\#queueing\ cars\ in\ incoming\ lanes - \#queueing\ cars\ in\ outgoing\ lanes|$, which is $|3 + 3 + 1 + 2 - 1 - 1 - 3 - 2| = 2$. In MalLight, rewards from all intersections in the road network are fed into a reward module, which generates a global reward. Details of reward generation are described in Section 3.3. The global reward is then combined with the local reward r_i^t , to form the final reward r_i^t for each agent.

- **Policy set π and discount factor γ :** At time t , each agent chooses an action following a policy $\pi: S \rightarrow A$. This choice can also be expressed as action-value function Q^π , which indicates the expected quality of an action a_i^t if the i th agent takes action a_i^t based on policy

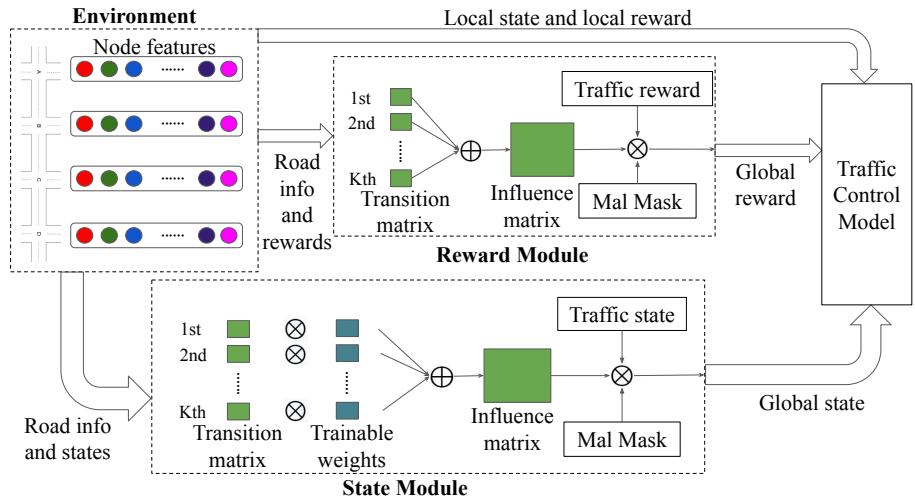


Figure 5: MalLight framework and modules.

π in state s_t^i . The formulation of the action-value function is

$$Q^\pi(s, a) = \mathbb{E} \left[\sum_{i=t}^{\infty} \gamma r_{t+i} | s_t^i = s, a_t^i = a \right] \quad (2)$$

Problem formulation and objective: In our problem, each intersection is controlled by one agent, and these agents share parameters from the same deep RL model. At each time step t , agent i (managing well-functioning traffic signals) obtains an observation o_i^t from the environment. Given traffic conditions and current signal phases at both local and other intersections, the objective of each agent is to determine the optimal action a_i^t that maximizes the reward r_i^t , formulated as:

$$\arg \max_{a_i^t \in A} r_i^t = -P_i - \sum_{j \in M} (w_j P_j) \quad (3)$$

Here, P_i represents the pressure at i th intersection, w_j is the weight of j th intersection with malfunctioning signals, for reward aggregation, and M is the set of intersections of malfunctioning signals.

3 METHOD

We introduce MalLight, a RL-based traffic signal control model built upon a general RL framework, incorporating new modules: an Influence-aware State Aggregation Module (ISAM) and an Influence-aware Reward Aggregation Module (IRAM). These modules enhance the model's resistance to signal malfunctions. In the following sections, we first introduce the overview of our model and then describe the detailed design of ISAM and IRAM.

3.1 Overview

To alleviate congestion in malfunctioning intersections through coordinated management with neighboring intersections (intersections can be multi-hop away), each agent should possess two key abilities. Firstly, each agent should have the ability to “care” about malfunctioning intersections. In previous works [61, 63, 3, 70], the reward for each RL agent was solely dependent on its own traffic conditions, which meant that each agent would only benefit from its own actions. In our setting, intercepting some vehicles at

a well-functioning intersection to prevent excessive traffic from entering malfunctioning intersections would lead to a reduction in the rewards earned by agents at well-functioning intersections. Therefore, to alleviate the traffic demand on malfunctioning intersections, each surrounding agent should be aware of the reward from these malfunctioning intersections. Secondly, each agent should have the ability to quantify how much influence it can exert on the malfunctioning intersection to determine the optimal actions. Different levels of influence require different policies.

The overall framework of MalLight is depicted in Figure 5. In line with the two aforementioned abilities, we introduce two modules, the Influence-aware State Aggregation module (ISAM), and the Influence-aware Reward Aggregation module (IRAM). In ISAM, we designed masked graph diffusion convolution and adapt it into our multi-agent scenario, where each agent needs to take nearby malfunctioning agents into consideration while being aware of its influence on these agents. In IRAM, we design a novel reward reshaping scheme based on diffusion process, which boost the performance of model when used in cooperation with masked diffusion convolution. Up to our knowledge, we are the first to adapt diffusion convolution into multi-agent RL scheme and the first to reshape reward by diffusion process.

To simplify the task, in MalLight, sensing devices for detecting the vehicle queue length of each lane, such as cameras or inductive loop detectors, are assumed to work well in all intersections including the malfunctioning ones, as traffic signals and sensing devices are generally managed by different systems [4, 26]. In addition, we assume that the system is aware of which intersections have malfunctioning traffic signals (e.g., through human reports or anomaly detection techniques [44, 49]), and the signal control agents can share information.

3.2 Influence-aware State Aggregation Module

The Influence-aware State Aggregation Module (ISAM) is designed to ensure that each agent's observation includes global states from malfunctioning intersections while taking into account its influence on these malfunctioning intersections. Higher influence corresponds to greater weight assigned to other intersections' states, and vice versa. ISAM takes road network static information and intersection-specific local states as input. For each agent, it generates an aggregated state that encompasses observations from malfunctioning intersections and their corresponding local states. This state information is subsequently fed into the RL model and serves as the agent's final environmental observation.

Inspired by the graph diffusion convolutional recurrent neural network (DCRNN) [29], which has achieved promising results in automatically modeling complex spatial dependencies on road networks, we design masked diffusion convolution and adapt it to multi-agent scheme, to capture influence among intersections, particularly the influence from well-functioning intersections to malfunctioning intersections. In the diffusion convolution operation, an influence vector is extracted for each agent. These vectors are used as weights to perform a weighted sum of states from other intersections. Additionally, before the weighted sum is calculated, a Malfunction Mask is applied to filter out data from well-functioning intersections. This helps agents at well-functioning intersections

focus on only themselves and the malfunctioning intersections. Details of ISAM are described below.

• **Diffusion Process.** Diffusion convolution [29] is a variation of the diffusion process [55] that incorporates trainable parameters into transition matrix of each layer, offering enhanced representativeness and flexibility in modeling influence. According to [55], the stationary distribution of the diffusion process can be expressed as a weighted combination of random walks on the graph:

$$\mathcal{P} = \sum_{k=1}^K \alpha(1-\alpha)^k (D_O^{-1}W)^k, \quad (4)$$

, where k represents the diffusion step, W is the edge weight matrix, D_O is the out-degree diagonal matrix of graph, $\alpha \in [0, 1]$ denotes the restart probability, and $D_O^{-1}W$ serves as the state transition matrix. This Markov process converges to the distribution $\mathcal{P} \in \mathbb{R}^{N \times N}$, where each i -th row $\mathcal{P}_{i,:}$ represents the likelihood of diffusion from node $v_i \in V$, where V represents the set of nodes in the graph, and N denotes the total number of nodes.

The edge weight matrix W is constructed based on the road distance between each pair of intersections. For normalization, we use a thresholded Gaussian kernel [48], defined as follows:

$$W_{i,j} = \exp\left(-\frac{dist(v_i, v_j)^2}{\sigma^2}\right) \quad (5)$$

In this formula, $W_{i,j}$ represents the edge weight between node v_i and node v_j , σ is the deviation of distances, and $dist(v_i, v_j)$ is 0 if nodes v_i and v_j are not connected by a road. Otherwise, it represents the distance between v_i and v_j .

• **Masked Diffusion Convolution.** Building upon the diffusion process described above, diffusion convolution introduces a trainable filter for each diffusion step k . Additionally, to ensure that agents do not concentrate their efforts on other functioning intersections, which do not require assistance, we apply the Malfunction Mask. The Malfunction Mask vector, denoted as $Mask \in \{0, 1\}^N$, where N represents the number of intersections in the road network. The element in the Malfunction Mask is set to 1 if the corresponding intersection is experiencing a malfunction, and 0 otherwise. The formulation of masked diffusion convolution is as follows:

$$S'_{:,p} = \left[\sum_{k=1}^K (\theta_k (D_O^{-1}W)^k) \odot Mask \right] S_{:,p}, p \in \{1, 2, 3, \dots, P\} \quad (6)$$

Here, \odot indicates Hadamard product, $S \in \mathbb{R}^{N \times P}$ represents the matrix of features in the graph (each agent's local state), P is the length of input features, $p \in [1, P]$ is an index of an element in input feature, θ_k represents the trainable filter for diffusion step k , and K is the number of diffusion steps. The masked diffusion convolution process involves random walks on the road network for each element of the input features. $S'_{i,:}$ indicates the information i -th agent needs considering its influence to other nodes(intersections). Masked diffusion convolution provides a more explainable method for making the model aware of the influence.

• **State Aggregation.** This step combines the local state of the target intersection with the global state S' originating from malfunctioning intersections. For the i -th agent operating at the i -th intersection, its corresponding global state is $S'_{i,:} \in \mathbb{R}^P$ and its local state is $S_{i,:} \in \mathbb{R}^P$. Finally, for the i -th agent, we obtain the state as:

$$S''_{i,:} = S'_{i,:} + S_{i,:}, i \in \{1, 2, 3, \dots, N\} \quad (7)$$

Here $S''_{i,:} \in \mathbb{R}^P$ contains all the information i -th agent should consider for making a decision and it serves as the input to the traffic control RL model.

3.3 Influence-aware Reward Aggregation Module

The Influence-aware Reward Aggregation Module (IRAM) is designed to ensure that each agent "cares" about the benefits of malfunctioning intersections and is "willing" to assist these intersections. IRAM takes road network static information and intersection rewards as input, and for each agent, generate aggregated reward containing rewards of malfunctioning intersections and corresponding local rewards. The reward information then serve as direction to optimize the RL model.

Reward Shaping [39] is a technique, where additional rewards are utilized to represent domain knowledge and guide the training of RL models using expert knowledge. In our context, we aggregate the rewards from malfunctioning intersections to each working agent, with different weights, representing influences between malfunctioning and working agents. The additional reward guides working agents in earning benefits for malfunctioning intersections. We employ a diffusion process without trainable parameters to perform a weighted sum of rewards. The formulation is as follows:

$$R' = \left[\sum_{k=1}^K (D_O^{-1} W)^k \odot \text{Mask} \right] R \quad (8)$$

where $R \in \mathbb{R}^N$ is the vector of rewards of all intersections, k is the diffusion step, $D_O^{-1} W$ is the state transition matrix. Similar to the state aggregation method described above, for i -th agent, we add the global reward R'_i and the local reward R_i to obtain the final reward, denoted as R''_i :

$$R''_i = R_i + R'_i, i \in \{1, 2, 3, \dots, N\} \quad (9)$$

3.4 Training

MalLight is updated using the Bellman Equation[13]:

$$Q(S''_i, a_i^t) = R''_i + \gamma \max Q(S''_i^{t+1}, a_i^{t+1}) \quad (10)$$

In this equation, R''_i represents the reward that i -th agent can obtain by taking action a_i^t based on the state S''_i , at decision time t . γ is discount factor. We employ the Mean Squared Error (MSE) as the loss function to measure the disparity between the current action-value estimation and the desired action-value estimation provided by the Bellman Equation. The RL model is optimized using the RMSprop[18] algorithm.

4 EXPERIMENTS

In this section, we conduct experiments to answer the following research questions:¹

- **RQ1:** Does MalLight outperform other SoTA methods on the malfunctioning signal scenario?
- **RQ2:** Does MalLight perform well on the normal traffic signal control scenario?
- **RQ3:** How the components in MalLight contribute to the performance?
- **RQ4:** Does MalLight decrease accidental risk?

¹Implementation can be found at: <https://anonymous.4open.science/r/MalLight-B3DC>

- **RQ5:** How the diffusion process distributes state and reward aggregation weights?
- **RQ6:** How do hyper-parameters affect the performance?
- **RQ7:** How does the number of malfunctioning intersections affect the performance?

4.1 Experiment Setting

Table 1: SUMO Environment Setting

Attribute	Value	Description
accel	2.0	max acceleration ability (m/s^2)
decel	4.5	max deceleration ability (m/s^2)
length	5.0	vehicle netto-length (m)
width	2.0	vehicle width (m)
maxSpeed	40	max velocity (km/h)
minGap	2.5	empty space after leader (m)
jmIgnoreFoeProb	0.05	the probability of vehicle ignoring foes

Environment setting: The experiments were conducted on a server with two Intel Xeon E5-2650 processors, 252GB of main memory, and two GeForce GTX 2080ti GPUs. We conducted experiments based on the LibSignal framework [35], which is a traffic signal simulation toolkit developed on top of simulator engines SUMO [30] and CityFlow [69, 7]. We used SUMO for our experiments due to its flexibility in simulating intersection malfunction. The SUMO settings are provided in Table 1. We generated vehicle trajectories by running the Dijkstra algorithm on each origin-destination pair in our datasets (introduced later). After feeding the traffic flow data into the simulator, each vehicle moves toward its destination following the predefined trajectory.

To better simulate traffic congestion and potential traffic accidents, we employ the Foelgnore strategy [14]. This mechanism enables each vehicle to ignore vehicles already within the intersection and potentially collide with them based on a predetermined probability. By simulating collisions at intersections, it mimics the congestion and chaos resulting from signal malfunctions.

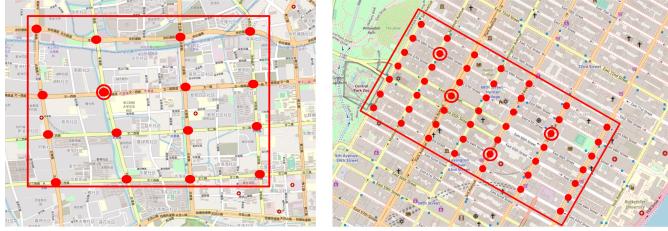
Model Setting: For the traffic control module, which makes decisions based on local and global information, we use a deep Q-learning model with a Multi-Layer Perceptron (MLP) model as the backbone. Details of MLP are shown in Table 2. The training algorithm is described in Section 3.4. Our model is trained for 200 episodes with a replay buffer size of 5000. For each episode, we run SUMO over the training dataset, with the signal phase determined by the current model. The model is updated with 10 iterations over all MDP samples in a replay buffer at the end of each episode. The learning rate is set to 0.001. The diffusion step K is 10. Discount factor γ is set to 0.95.

Table 2: Details of Backbone MLP

Layer	#Neurons
Input Layer	12
Hidden Layer 1	20
Hidden Layer 2	20
Output Layer	8

4.2 Datasets

We use two real-world and one synthetic traffic flow data in our experiments. The road networks for real-world datasets were imported from OpenStreetMap [41], as shown in Figure 6. The data



(a) Gudang Sub-district, Hangzhou, China (b) Upper East Side, Manhattan, New York, USA

Figure 6: Road networks for real-world datasets. Red dots mark intersections chosen for experiments. Red circles mark malfunctioning signals.

statistics are listed in Table 3. The road network of the synthetic Grid4*4 dataset is the same as that of Hangzhou dataset, and the only difference between them is the traffic flow pattern. In datasets, a vehicle is described as (o, d, t) , where o denotes the origin location, d denotes the destination location, t represents the departure time. Both origin (o) and destination (d) locations are within the road network.

- **Hangzhou Dataset** [57]: This is a widely used real-world dataset in prior works [61, 36, 65]. We selected 16 intersections in the Gudang Sub-district of Hangzhou and the signal at one of these intersections is set to be malfunctioning. Vehicle information was collected from roadside surveillance cameras. Each record in the camera data includes time, camera ID, and vehicle details. The camera ID indicates the vehicle’s location. Through the analysis of these records, we generated vehicle trajectories as they passed through the selected intersections.

- **New York Dataset** [40]: This is also a well-known real-world dataset used in prior works [3, 63, 61, 36]. We take 49 intersections in the Upper East Side of Manhattan as a demonstration. In this setting, the traffic signals on four intersections are set to be malfunctioning. Vehicle information was generated using open-source taxi trip data [40]. This taxi trip data includes the geo-locations of the origin and destination for each trip. We initially mapped these geo-locations to the nearest intersections and then retained the trips that fell within the selected area. Because taxi trips can be seen as a sample from real-world trip distributions, we empirically scaled up the number of trips by 4.5 to consider the impacts of personal and commercial vehicles. We made the simulation traffic volume roughly aligned with the real-world traffic volume indicated by New York Traffic Data Viewer [56].

- **Grid4*4 Dataset**: A synthetic dataset. The road network of Grid4*4 is the same as that of the Hangzhou dataset. We generate vehicles with a steady arrival rate and random origin and destination locations. Table 3 shows details of generated traffic flow.

Each dataset comprises 2 hours of traffic flow data, of which we utilize 1 hour for training purposes and the remaining hour for testing.

4.3 Evaluation Metrics

Following existing studies [60], we mainly use *throughput* to evaluate the performance of different models for traffic signal control. In addition, we simulate the number of accidents to evaluate the

Table 3: Data statistic of real-world traffic dataset

Dataset	#intersections	Arrival rate (#vehicles/300s)			
		Mean	Std	Max	Min
Hangzhou	16	1053.26	86.70	1352	512
New York	49	361.19	10.08	411	324
Grid4*4	16	1200	0	1200	1200

impacts of our model on reducing accidental risk. The detailed descriptions are as follows:

- **Network Throughput**: It represents the number of vehicles that have completed their trip in the road network during the one-hour testing period.
- **Intersection Throughput**: It represents the (average) number of vehicles that pass through malfunctioning intersection(s) during the one-hour testing period. We introduced this intersection-level metric to better compare the performance of models on malfunctioning intersections.
- **Number of Accidents**: It represents the number of traffic collisions simulated by SUMO in malfunctioning intersections. This is an additional metric. We do not calculate collisions at the road network level considering traffic collisions rarely occur in signalized intersections on SUMO.

4.4 Compared Methods

- **FixedTime** [25]: A policy with predefined fixed phase length and phase splits, which is commonly used in practice. The offset is randomly set in our experiments.
- **SOTL** [6]: A rule-based method based on the request from current phase and competing phases.
- **MaxPressure** [58]: A state-of-the-art (SoTA) optimization-based algorithm that minimizes pressure at each intersection. It reduces over-saturation by balancing queue lengths between each pair of neighboring intersections.
- **IDQN** [62]: A deep RL method in which information is not shared among agents. Each intersection is controlled by one agent, and these agents do not share model parameters.
- **CoLight** [61]: A deep RL method based on Graph Attention Network.
- **Advanced-CoLight** [70]: A SoTA RL algorithm template that incorporates advanced traffic state (ATS) into the RL model. It uses CoLight as a backbone network.
- **Advanced-PressLight** [70]: Similar to Advanced-CoLight, it uses PressLight [63] as the backbone network. PressLight is a deep RL method that coordinates traffic signals by minimizing the Pressure [58] at each intersection to maximize the throughput of the road network.
- **PRLight** [17]: A SoTA RL method that integrates a dynamic graph representation module and a traffic prediction module to effectively control traffic signals and mitigate action hysteresis in real-time traffic management.

4.5 Experiment Result

4.5.1 *Overall Performance.* Our experimental results using Hangzhou dataset and New York dataset are shown in Table 4. By analyzing the results, we have the following findings:

- **RQ1**: By analyzing the Reduction Ratio (RR) columns in Table 4,

Table 4: Results on Hangzhou, New York and Grid4*4 datasets

Methods	Intersection throughput				Road network throughput		
	No Mal ↑	Mal ↑	RR (%) ↓	#Acc ↓	No Mal ↑	Mal ↑	RR (%) ↓
Hangzhou							
FixedTime	538	390	27.5	18	2923	2653	9.2
SOTL	760	489	35.7	23	3674	3118	15.1
MaxPressure	592	405	31.5	27	3997	3777	5.5
IDQN	676	387	42.8	36	4554	4133	9.2
CoLight	702	457	34.9	31	4617	4219	8.6
Advanced-CoLight	740	526	28.9	27	4679	4295	8.2
Advanced-PressLight	752	561	25.3	25	4745	4370	7.9
PRLight	841	577	31.4	29	4779	4310	9.8
MalLight (ours)	822	715	13	11	4607	4416	4.2
New York							
FixedTime	540	365	32	31	2671	2239	16.1
SOTL	579	492	15	18	3172	2594	18.2
MaxPressure	630	570	9	13	3068	2852	7
IDQN	617	554	10.2	22	3223	2883	10.5
CoLight	627	561	10.5	18	3507	3119	11.1
Advanced-CoLight	625	562	10.1	16	3612	3154	12.6
Advanced-PressLight	642	542	15.5	17	3784	3276	13.4
PRLight	639	534	16.4	19	3695	3117	15.6
MalLight (ours)	633	592	6.5	8	3532	3309	6.3
Grid4*4							
FixedTime	579	381	34.2	25	3192	2582	19.1
SOTL	612	418	31.7	23	3463	2829	18.3
MaxPressure	645	454	29.6	21	3728	3053	18.1
IDQN	796	504	36.7	29	4339	3454	20.4
CoLight	832	589	29.2	19	4591	3778	17.7
Advanced-CoLight	854	624	26.9	17	4829	4076	15.6
Advanced-PressLight	883	679	23.1	17	4914	4236	13.8
PRLight	892	684	23.3	15	4963	4258	14.2
MalLight (ours)	877	767	12.5	9	4871	4515	7.3

No Mal denotes No Malfunction, Mal denotes Malfunction, RR denotes Reduction Ratio, #Acc denotes #Accidents

we can conclude that our model, MalLight, performs better in malfunctioning scenarios. When a malfunction occurs, the throughput reduction ratios, both at the intersection level and network level, are lower than those of other models. This suggests that our proposed state and reward modules genuinely enhance the resilience of the RL model to intersection malfunctions. By reducing the traffic demand in malfunctioning intersections, congestion is alleviated, and malfunctioning intersections experience a lower throughput reduction. Specifically, on Hangzhou dataset, compared with the second-best model Advanced-PressLight, our model alleviates throughput reduction by as much as 48.6%.

• **RQ2:** By analyzing the throughput when no malfunction, we conclude that MalLight performs well in normal (without malfunction) scenarios. The throughputs at both the intersection and network level when no malfunction are very close to the highest throughputs of other models. In addition, traditional Methods (FixedTime, SOTL, MaxPressure) perform poorly in both intersection and network levels. This is because traffic flow in our datasets is highly dynamic. These traditional methods rely heavily on human-defined assumptions of the environment, which may not be true under

some circumstances.

- An interesting finding is that the throughput reduction of the whole road network can be greatly higher than the throughput reduction of malfunctioning intersections, which means that congestion in malfunctioning intersections affects other well-functioning intersections. This occurs because when congestion arises in malfunctioning intersections, waiting vehicles can spill over into neighboring intersections, leading to congestion in those neighboring intersections and a subsequent reduction in their throughput. For example, in the case of SOTL under the New York dataset, the total throughput reduction of malfunctioning intersections is $(579 - 492) \times 4 = 348$, while the throughput reduction for the entire road network is 578, which is 66.1% higher than that of the malfunctioning intersections.

4.5.2 *Ablation Study.* We consider several variations of MalLight:

- **MalLight-S:** For state aggregation, instead of using diffusion convolution which contains trainable parameters, MalLight-S uses a diffusion process without any trainable parameter. This variant is to show the benefits of the flexible state aggregation in our model.
- **MalLight-R:** We remove reward aggregation in MalLight-R. This

means that each agent exclusively concentrates on its own reward and disregards the benefits of malfunctioning intersections.

- **MalLight-M** We remove Malfunction Mask in MalLight-M. Instead of having each agent focus on its own intersection and malfunctioning intersections, MalLight-M allows each agent to observe and care about the entire road network.

Table 5: Results of Ablation Study

Methods	Intersection Throughput Reduction (%) ↓		
	Hangzhou	New York	Grid4*4
MalLight*	13	6.5	12.5
MalLight-S	18.8	9.9	20.3
MalLight-R	30.9	16.1	27.4
MalLight-M	26.3	11.3	23.6

Table 5 shows the performance of variations of our model. We have the following findings (**RQ3**):

- The trainable parameters in diffusion convolution provide the model with greater adaptability, thereby enhancing its performance.
- The absence of reward aggregation, which makes the agent prioritize caring about malfunctioning intersections, leads to a significant decrease in performance.
- MalLight outperforms MalLight-M due to the presence of the Malfunction Mask, which filters out well-functioning intersections and enables each agent to concentrate on what we want it to pay attention to—malfunctioning intersections. Features and rewards coming from well-functioning intersections can be distracting for the agent, under which circumstances the agent may struggle to focus on malfunctioning intersections or even itself.

4.5.3 Accident Risk (RQ4). : We analyze the relationship between accident risk and throughput reduction. The number of accidents serves as an additional metric for our model. When an accident occurs at an intersection, the victims can block following vehicles, leading to congestion. By examining the column labeled #Acc, we can observe that a higher level of Reduction Ratio typically corresponds to a higher number of accidents.

4.5.4 Weight Distribution from Diffusion Process (RQ5). We select the upper-left intersection in Hangzhou road network as an example and calculate the influence weights generated by a diffusion process. Figure 7 demonstrates the influence (weights) from this node to other nodes at distances of 1 to 4 hops away. Because the road network is a bidirectional graph, with two edges connecting a pair of nodes and sharing the same initial weight (determined by road distance processed through a thresholded Gaussian kernel), this figure can also be viewed as representing the influence from this specific node to other nodes at varying distances (hops) in the graph. We can observe that as the distance (number of hops in the graph) increases, the influence between two nodes tends to decrease.

4.5.5 Parameter Sensitivity (RQ6). We set different K values and observe the throughput reduction metric. The parameter K in the diffusion process, as well as its counterpart in diffusion convolution, determines how many hops a signal can diffuse along edges. For example, if $K = 1$, the malfunctioning intersection will only consider directly connected intersections. The result is shown in Figure 8. When K is set too low, the observation scope of well-functioning

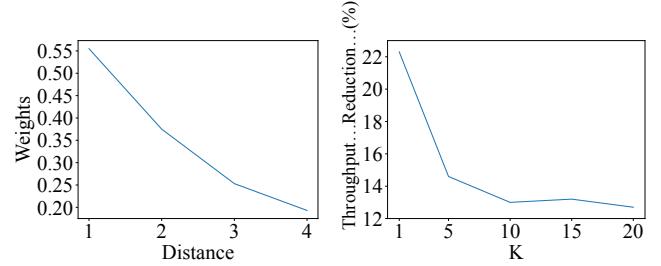


Figure 7: Impacts of intersection distances on their influence

intersections will be significantly limited, and the influence among intersections cannot be adequately modeled. It's worth noting that a higher value of K requires more computational resources. Based on the curve, we set $K = 10$ in experiments.

4.5.6 Performance on various number of malfunctioning intersections (RQ7). We conduct experiments using the MalLight framework on the New York dataset, varying the number of malfunctioning signals. For each experiment, our model is run five times, and scores are averaged. The result is shown in Figure 9. As the number of malfunctioning signals increases, the curve initially remains relatively flat before experiencing a sharp increase after approximately 15 signals. Our model operates under the assumption that nearby working intersections can provide assistance to malfunctioning ones. Consequently, if the percentage of malfunctioning intersections exceeds a certain threshold, the effectiveness of the model decreases significantly.

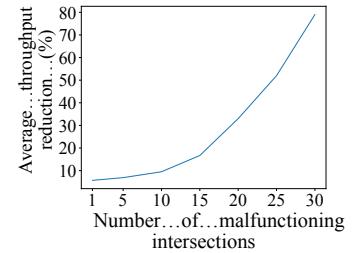


Figure 8: Impacts of diffusion steps K on throughput reduction ratio

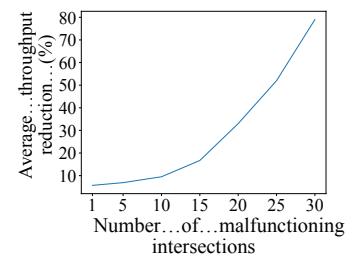


Figure 9: Performance on different #malfunctioning intersections

5 DISCUSSION

Lessons Learned: We summarize key lessons learned from work:

- We design our malfunction-resistant approach based on intersection cooperation, as manifested in the design of state and reward modules. In contrast to alternative methods that neglect the advantageous potential of well-functioning intersections in alleviating the plight of malfunctioning counterparts, our approach stands out by affording a notably reduced reduction in throughput.
- The integration of state and reward aggregation mechanisms, while taking into consideration the influences among agents, is of great importance in coordinating intersections and enhancing the model's resilience in the face of signal malfunction. In direct comparison to the Advanced-PressLight model which shares similar architecture but lacks the incorporation of these two pivotal modules, our approach achieves better performance in the presence of signal malfunction. In essence, by optimizing control of

surrounding intersections, the impact of signal malfunction can be alleviated, thus aligning with our central research idea.

Limitations: Constrained by the inherent limitations of the driver behavior model integrated within the SUMO traffic simulator, regrettably fails to authentically replicate the congestion resulting from signal malfunctions, we resorted to the use of traffic collisions as a surrogate for congestion induction. We have to acknowledge that not all instances of congestion are precipitated by traffic collisions. In the future, our objective is to enhance the fidelity of our simulation by refining the driver behavior model to more accurately emulate real-world congestion dynamics. We envision that the core concept of our research can be effectively transposed to a more realistic environment, which has a more comprehensive and precise representation of traffic congestion causality.

6 RELATED WORK

6.1 Traffic Signal Malfunction

The time series of traffic signal malfunctions and the subsequent actions are as follows: malfunction occurs, detection, and response. Methods designed to address malfunctioning traffic signals can be categorized into two groups: detection and response methods.

Detection methods: As stated in [4], malfunction detection typically relies on manual reports. Several resources [64] are available for reporting signal problems, including 24-hour signal repair hotlines, transportation emergency operations centers, and 911 calls. Manual detection, based on road user reports, often results in delays. Remote monitoring techniques may be employed in some areas when crew time and workload allow, but they are seldom effective due to the expensive human resources required.

Furthermore, some automatic detection systems have been developed. Purba et al [44] develop a self-diagnosis system to detect traffic signal malfunction immediately. Soh et al [49] design a monitoring system based on fuzzy technology for traffic signal malfunction detection. In our work, we assume malfunctioning signals are detected by the existing work.

Response methods: The response to put malfunctioning traffic signals back to fully functional relies heavily on manual repair or replacement actions taken by contractors or authorities [4]. In addition to manual response, some automatic response methods are also designed. Malfunction flash [2] is designed to prevent safety issues from traffic signal malfunction. When an error is detected, the signal is automatically placed into flash mode as a safety precaution (if the signal is still accessible). However, malfunction flash only serves as a warning to road users. However, some decisions, such as whether to cross or wait, and whether to stop before crossing, are left to drivers, based on their observation of traffic conditions. Hunter et al [22] studied the driver behaviors under malfunction flash and revealed that malfunction flash cannot eliminate the safety issue of traffic signal malfunctions, not to mention other impacts such as congestion. In addition, our method is applicable in situations where the traffic signal is in a blackout, where the traffic signal cannot even work as a flashing warning to drivers.

Some methods are designed for efficient human resources distribution. For example, Mathibela [34] designs a RUSBoost-based framework to predict the best distribution of human resources at critical intersections in the presence of malfunctioning traffic lights.

Our work is considered an immediate response before human resources get involved to prevent the heavy consequences resulting from traffic signal malfunction.

6.2 Traffic Signal Control

Traffic Signal Control is an important sub-domain of smart city [67, 72, 66, 21, 8]. Methods for controlling traffic signals can be divided into classic optimization-based methods [58, 6, 51] and RL-based methods [62, 63, 70, 61, 17]. Among these techniques, RL is most popular these years because of its ability to learn directly from complex conditions without assumptions about the environment [54, 31, 32, 59, 24, 47, 23]. Thus, RL-based methods and Deep Learning based methods usually outperform classic optimization-based methods [19, 20, 68, 12]. However, current RL-based methods exhibit suboptimal performance under scenarios involving signal malfunctions due to the isolation of reward mechanisms. Specifically, each agent operates with a focus exclusively on optimizing the traffic flow at its own intersection, neglecting the potential benefits of adjacent intersections. While this approach is effective in standard conditions, where each intersection is regulated by a functioning agent, it falters in the event of a signal malfunction. In such cases, the affected intersection lacks the capability to leverage support from neighboring intersections, leading to increased congestion.

Recently, a number of works have focused on traffic disruptions, such as missing data [36], incidents [14, 45], or weather changes [9, 10, 11]. Traffic signal malfunction is different from these disruptions in that: Firstly, these disruptions affect the RL model's input, while signal malfunction disables the RL agent from interacting with the environment, i.e. affects the output. Secondly, these disruptions test a single model's robustness, while signal malfunction tests the coordination ability of the whole system. Because a malfunctioning signal loses the ability to control traffic flow, it is naturally in need of help from other well-functioning signals, which makes cooperation an important factor in this context. In conclusion, traffic signal malfunction is a unique scenario that has not been considered in previous traffic disruption and robustness works.

7 CONCLUSION

In this paper, we study the problem of intelligent traffic signal control in the presence of traffic signal malfunction. We design MalLight, an traffic signal control model founded on RL principles. In the MalLight framework, we introduce two pioneering components, namely the Influence-aware State Aggregation Module and the Influence-aware Reward Aggregation Module, which are seamlessly integrated into the RL architecture. We subsequently conduct a comprehensive array of empirical investigations on two real-world datasets. Our empirical findings demonstrate that MalLight surpasses the performance of extant methods in both normal conditions and scenarios characterized by signal malfunctions. Notably, our results reveal a substantial improvement in the reduction of intersection throughput, with a remarkable mitigation rate of up to 48.6%.

8 ACKNOWLEDGMENTS

The work was supported in part by the National Science Foundation under Grant No. 2047822, 1952096, 2318697 and 2421839. We thank all the reviewers for their insightful feedback to improve this paper.

REFERENCES

[1] Md Aftabuzzaman. "Measuring traffic congestion-a critical review". In: *30th Australasian transport research forum*. ETM GROUP London, UK, 2007, pp. 1–16.

[2] B Benioff, FC Dock, and C Carson. *A STUDY OF CLEARANCE INTERVALS, FLASHING OPERATION, AND LEFT-TURN PHASING AT TRAFFIC SIGNALS. VOLUME 2. CLEARANCE INTERVALS*. Tech. rep, 1980.

[3] Chacha Chen et al. "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. 04. 2020, pp. 3414–3421.

[4] Wenling Chen, Larrie Henley, and Jeff Price. "Assessment of traffic signal maintenance and operations needs at Virginia department of transportation". In: *Transportation research record* 2128.1 (2009), pp. 11–19.

[5] Md Mamun Chowdhury. "Traffic congestion and mismanagement in dhaka city". In: *Planned Decentralization: Aspired Development, World Town Planning Day* (2013).

[6] Seung-Bae Cools, Carlos Gershenson, and Bart D'Hooghe. "Self-organizing traffic lights: A realistic simulation". In: *Advances in applied self-organizing systems* (2013), pp. 45–55.

[7] Longchao Da et al. "CityFlowER: An Efficient and Realistic Traffic Simulator with Embedded Machine Learning Models". In: *arXiv preprint arXiv:2402.06127* (2024).

[8] Longchao Da et al. "Open-ti: Open traffic intelligence with augmented language model". In: *International Journal of Machine Learning and Cybernetics* (2024), pp. 1–26.

[9] Longchao Da et al. "Prompt to transfer: Sim-to-real Transfer for Traffic Signal Control with Prompt Learning". In: *In Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI'24)*. 2024.

[10] Longchao Da et al. "Sim2Real Transfer for Traffic Signal Control". In: *2023 IEEE 19th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2023, pp. 1–2.

[11] Longchao Da et al. "Uncertainty-aware Grounded Action Transformation towards Sim-to-Real Transfer for Traffic Signal Control". In: *In Proceedings of 62nd IEEE Conference on Decision and Control (CDC 2023)*. 2023.

[12] Yi Ding et al. "P2-Loc: A Person-2-Person Indoor Localization System in On-Demand Delivery". In: *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 6.1 (2022), 9:1–9:24. doi: 10.1145/3517238, URL: <https://doi.org/10.1145/3517238>.

[13] Zihan Ding et al. "Introduction to reinforcement learning". In: *Deep reinforcement learning: fundamentals, research and applications* (2020), pp. 47–123.

[14] Wenlu Du et al. "Safelight: A reinforcement learning method toward collision-free traffic signal control". In: *Proceedings of the AAAI Conference on Artificial Intelligence*. Vol. 37. 12. 2023, pp. 14801–14810.

[15] Barry C Ezell et al. "Cyber risk to transportation, industrial control systems, and traffic signal controllers". In: *Environment Systems and Decisions* 33 (2013), pp. 508–516.

[16] Branden Ghena et al. "Green Lights Forever: Analyzing the Security of Traffic Infrastructure". In: *WOOT 14* (2014), pp. 7–7.

[17] Xiao Han et al. "Mitigating Action Hysteresis in Traffic Signal Control with Traffic Predictive Reinforcement Learning". In: *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 2023, pp. 673–684.

[18] Geoffrey Hinton, Nitish Srivastava, and Kevin Swersky. "Neural networks for machine learning lecture 6a overview of mini-batch gradient descent". In: *Cited on 14.8* (2012), p. 2.

[19] Zhiqing Hong et al. "CoMiner: nationwide behavior-driven unsupervised spatial coordinate mining from uncertain delivery events". In: *Proceedings of the 30th International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2022, Seattle, Washington, November 1–4, 2022*. Ed. by Matthias Renz and Mohamed Sarwat. ACM, 2022, 10:1–10:10. doi: 10.1145/3557915.3560944. URL: <https://doi.org/10.1145/3557915.3560944>.

[20] Zhiqing Hong et al. "FastAddr: real-time abnormal address detection via contrastive augmentation for location-based services". In: *Proceedings of the 30th International Conference on Advances in Geographic Information Systems, SIGSPATIAL 2022, Seattle, Washington, November 1–4, 2022*. Ed. by Matthias Renz and Mohamed Sarwat. ACM, 2022, 64:1–64:10. doi: 10.1145/3557915.3560999. URL: <https://doi.org/10.1145/3557915.3560999>.

[21] Zhiqing Hong et al. "SmallMap: Low-cost Community Road Map Sensing with Uncertain Delivery Behavior". In: *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.* 8.2 (2024), 50:1–50:26. doi: 10.1145/3659596. URL: <https://doi.org/10.1145/3659596>.

[22] Michael Hunter et al. "Mode of flashing for malfunctioning traffic signals". In: *Journal of transportation engineering* 137.7 (2011), pp. 438–444.

[23] Haoyuan Jiang et al. "GuideLight: Industrial Solution" Guidance for More Practical Traffic Signal Control Agents". In: *arXiv preprint arXiv:2407.10811* (2024).

[24] Haoyuan Jiang et al. "X-Light: Cross-City Traffic Signal Control Using Transformer on Transformer as Meta Multi-Agent Reinforcement Learner". In: *arXiv preprint arXiv:2404.12090* (2024).

[25] Peter Koonce and Lee Rodegerds. *Traffic signal timing manual*. Tech. rep. United States. Federal Highway Administration, 2008.

[26] P Kuppusamy et al. "Design of smart traffic signal system using internet of things and genetic algorithm". In: *Advances in Big Data and Cloud Computing*. Springer, 2018, pp. 395–403.

[27] JooChan Lee, JangHoon Kim, and JungTaek Seo. "Cyber attack scenarios on smart city and their ripple effects". In: *2019 international conference on platform technology and service (PlatCon)*. IEEE, 2019, pp. 1–5.

[28] Serge Lhomme et al. "Analyzing resilience of urban networks: a preliminary step towards more flood resilient cities". In: *Natural hazards and earth system sciences* 13.2 (2013), pp. 221–230.

[29] Yaguang Li et al. "Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting". In: *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net, 2018. URL: <https://openreview.net/forum?id=SjHxGWAZ>.

[30] Pablo Alvarez Lopez et al. "Microscopic Traffic Simulation using SUMO". In: *The 21st IEEE International Conference on Intelligent Transportation Systems*. IEEE, 2018. URL: <https://elib.dlr.de/124092/>.

[31] Wenjun Lyu et al. "REDE: Exploring Relay Transportation for Efficient Last-mile Delivery". In: *39th IEEE International Conference on Data Engineering, ICDE 2023, Anaheim, CA, USA, April 3-7, 2023*. IEEE, 2023, pp. 3003–3016. doi: 10.1109/ICDE55515.2023.00230. URL: <https://doi.org/10.1109/ICDE55515.2023.00230>.

[32] Wenjun Lyu et al. "Towards Fair Workload Assessment via Homogeneous Order Grouping in Last-mile Delivery". In: *Proceedings of the 31st ACM International Conference on Information & Knowledge Management, Atlanta, GA, USA, October 17-21, 2022*. Ed. by Mohammad Al Hasan and Li Xiong. ACM, 2022, pp. 3361–3370. doi: 10.1145/3511808.3557132. URL: <https://doi.org/10.1145/3511808.3557132>.

[33] Hashim MN Al-Madani. "Dynamic vehicular delay comparison between a police-controlled roundabout and a traffic signal". In: *Transportation Research Part A: Policy and Practice* 37.8 (2003), pp. 681–688.

[34] Bonolo Mathibela. "Another broken traffic light? Reducing traffic congestion in resource constrained environments". In: *2017 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2017, pp. 1305–1310.

[35] Hao Mei et al. "LibSignal: An Open Library for Traffic Signal Control". In: *arXiv preprint arXiv:2211.10649* (2022).

[36] Hao Mei et al. "Reinforcement Learning Approaches for Traffic Signal Control under Missing Data". In: *arXiv preprint arXiv:2304.10722* (2023).

[37] National Center for Statistics and Analysis. *Traffic Safety Facts 2020: A Compilation of Motor Vehicle Crash Data*. Report DOT HS 813 375. National Highway Traffic Safety Administration, Oct. 2022.

[38] Newark Connect. <https://www.newarknj.gov/card/newark-connect>. Accessed: 2023-06-30.

[39] Andrew Y Ng, Daishi Harada, and Stuart Russell. "Policy invariance under reward transformations: Theory and application to reward shaping". In: *Icml*. Vol. 99. Citeseer, 1999, pp. 278–287.

[40] NYC-Taxi-Trips. <https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page>.

[41] OpenStreetMap. <https://www.openstreetmap.org>.

[42] Basil Paul and H Pitale. "Comparison of capacity at signalised and unsignalised intersection". In: *IJSTE-Int J Sci Technol Eng* 3 (2017), pp. 108–111.

[43] Maria Pregnolato et al. "Assessing urban strategies for reducing the impacts of extreme weather on infrastructure networks". In: *Royal Society open science* 3.5 (2016), p. 160023.

[44] Aleksander Purba, Rahayu Sulistyorini, and Ageng Sadnowo. "Developing Monitoring System of Traffic Signal Using Microcontroller Device by SMS of GSM Network". In: *Prosiding Semnas SINTA FT UNILA Vol. 1 Tahun 2018* 1.1 (2018), pp. 273–277.

[45] Filipe Rodrigues and Carlos Lima Azevedo. "Towards robust deep reinforcement learning for traffic signal control: Demand surges, incidents and sensor failures". In: *2019 IEEE intelligent transportation systems conference (ITSC)*. IEEE, 2019, pp. 3559–3566.

[46] Sandra Rosenbloom. "Peak-period traffic congestion: A state-of-the-art analysis and evaluation of effective solutions". In: *Transportation* 7.2 (1978), pp. 167–191.

[47] Jingjing Ruan et al. "CoSLight: Co-optimizing Collaborator Selection and Decision-making to Enhance Traffic Signal Control". In: *arXiv preprint arXiv:2405.17152* (2024).

[48] David I Shuman et al. "The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains". In: *IEEE signal processing magazine* 30.3 (2013), pp. 83–98.

[49] Azura Che Soh, Asnor Juraiza Ishak, and Mohd Hanif Zaini. "Smart monitoring fault detection system for malfunction traffic light operation". In: *2013 IEEE 8th Conference on Industrial Electronics and Applications (ICIEA)*. IEEE, 2013, pp. 549–554.

[50] Some Chicago traffic lights still broken months, even a year, after reporting to 311. <https://abc7chicago.com/chicago-traffic-light-broken-311-cdot/12300361/>. Accessed: 2023-06-30.

[51] Aleksandar Stevanovic, Cameron Kergaye, and Peter T Martin. "Scoot and scats: A closer look into their operations". In: *88th Annual Meeting of the Transportation Research Board, Washington DC*. 2009.

[52] Pablo Suarez et al. "Impacts of flooding and climate change on urban transportation: A systemwide performance assessment of the Boston Metro Area". In: *Transportation Research Part D: transport and environment* 10.3 (2005), pp. 231–244.

[53] Henk Taale, PHG van Bekkum, and MLD van Rij. "Evaluation of traffic management by the traffic police". In: (2004).

[54] Heng Tan et al. "Joint Rebalancing and Charging for Shared Electric Micro-mobility Vehicles with Energy-informed Demand". In: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM 2023, Birmingham, United Kingdom, October 21–25, 2023*. Ed. by Ingo Frommholz et al. ACM, 2023, pp. 2392–2401. doi: 10.1145/3583780.3614942. URL: <https://doi.org/10.1145/3583780.3614942>.

[55] Shang-Hua Teng et al. "Scalable algorithms for data and network analysis". In: *Foundations and Trends® in Theoretical Computer Science* 12.1–2 (2016), pp. 1–274.

[56] *Traffic Data Viewer*. <https://www.dot.ny.gov/tdv/>.

[57] *Traffic Signal Control Datasets*. <https://traffic-signal-control.github.io/index.html#open-datasets/>.

[58] Pravin Varaiya. "Max pressure control of a network of signalized intersections". In: *Transportation Research Part C: Emerging Technologies* 36 (2013), pp. 177–195.

[59] Dimitris M Vlahogiannis et al. "HumanLight: Incentivizing ridesharing via human-centric deep reinforcement learning in traffic signal control". In: *Transportation Research Part C: Emerging Technologies* (2023).

[60] Hua Wei et al. "A survey on traffic signal control methods". In: *arXiv preprint arXiv:1904.08117* (2019).

[61] Hua Wei et al. "Colight: Learning network-level cooperation for traffic signal control". In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 2019, pp. 1913–1922.

[62] Hua Wei et al. "Intellilight: A reinforcement learning approach for intelligent traffic light control". In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2018, pp. 2496–2505.

[63] Hua Wei et al. "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*. 2019, pp. 1290–1298.

[64] *what should you do if a traffic is malfunctioning*. <https://elteccorp.com/news/other/what-should-you-do-if-a-traffic-light-is-malfunctioning/>.

[65] Libing Wu et al. "DynSTGAT: Dynamic spatial-temporal graph attention network for traffic signal control". In: *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*. 2021, pp. 2150–2159.

[66] Hua Yan et al. "Robust Route Planning under Uncertain Pickup Requests for Last-mile Delivery". In: *Proceedings of the ACM on Web Conference 2024, WWW 2024, Singapore, May 13–24, 2024*. Ed. by Tat-Seng Chua et al. ACM, 2024, pp. 3022–3030. doi: 10.1145/3589334.3645595. URL: <https://doi.org/10.1145/3589334.3645595>.

[67] Guang Yang et al. "CARPG: Cross-City Knowledge Transfer for Traffic Accident Prediction via Attentive Region-Level Parameter Generation". In: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM 2023, Birmingham, United Kingdom, October 21–25, 2023*. Ed. by Ingo Frommholz et al. ACM, 2023, pp. 2939–2948. doi: 10.1145/3583780.3614802. URL: <https://doi.org/10.1145/3583780.3614802>.

[68] Qinchen Yang et al. "Mapping Plastic Mulched Farmland for High Resolution Images of Unmanned Aerial Vehicle Using Deep Semantic Segmentation". In: *Remote. Sens.* 11.17 (2019), p. 2008. doi: 10.3390/RS11172008. URL: <https://doi.org/10.3390/rs11172008>.

[69] Huichu Zhang et al. "Cityflow: A multi-agent reinforcement learning environment for large scale city traffic scenario". In: *The world wide web conference*. 2019, pp. 3620–3624.

[70] Liang Zhang et al. "Expression might be enough: Representing pressure and demand for reinforcement learning based traffic signal control". In: *International Conference on Machine Learning*. PMLR, 2022, pp. 26645–26654.

[71] Guanjie Zheng et al. "Learning phase competition for traffic signal control". In: *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 2019, pp. 1963–1972.

[72] Shuxin Zhong et al. "RLIFE: Remaining Lifespan Prediction for E-scooters". In: *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, CIKM 2023, Birmingham, United Kingdom, October 21–25, 2023*. Ed. by Ingo Frommholz et al. ACM, 2023, pp. 3544–3553. doi: 10.1145/3583780.3615037. URL: <https://doi.org/10.1145/3583780.3615037>.