



CFLight: Enhancing Safety with Traffic Signal Control through Counterfactual Learning

Mingyuan Li
henryli_i@bupt.edu.cn
Lanzhou University;
Beijing University of Posts and
Telecommunications,
Lanzhou, China

Chunyu Liu
chunyuliu@bupt.edu.cn
Beijing University of Posts and
Telecommunications,
Beijing, China

Zhuojun Li
zhuojinli@bupt.edu.cn
Beijing University of Posts and
Telecommunications,
Beijing, China

Xiao Liu
liuxiao68@bupt.edu.cn
Beijing University of Posts and
Telecommunications,
Beijing, China

Guangsheng Yu
reimusaber@gmail.com
Independent Researcher
Sydney, Australia

Bo Du
bo.du@griffith.edu.au
Griffith University,
Brisbane, Australia

Jun Shen
jshen@uow.edu.au
University of Wollongong,
Wollongong, Australia

Qiang Wu*
wuqiang@lzu.edu.cn
Lanzhou University
Lanzhou, China

Abstract

Traffic accidents result in millions of injuries and fatalities globally, with a significant number occurring at intersections each year. Traffic Signal Control (TSC) is an effective strategy for enhancing safety at these urban junctures. Despite the growing popularity of Reinforcement Learning (RL) methods in optimizing TSC, these methods often prioritize driving efficiency over safety, thus failing to address the critical balance between these two aspects. Additionally, these methods usually need more interpretability. CounterFactual (CF) learning is a promising approach for various causal analysis fields. In this study, we introduce a novel framework to improve RL for safety aspects in TSC. This framework introduces a novel method based on CF learning to address the question: “What if, when an unsafe event occurs, we backtrack to perform alternative actions, and will this unsafe event still occur in the subsequent period?” To answer this question, we propose a new structure causal model to predict the result after executing different actions, and we propose a new CF module that integrates with additional “X” modules to promote safe RL practices. Our proposed algorithm, CFLight, achieves a near-zero collision control strategy for complex intersection safety. Extensive experiments on real-world and synthetic datasets demonstrate that CFLight significantly outperforms conventional and state-of-the-art safe RL methods in both safety and traffic efficiency. The data and code are available in the github¹.

*Qiang Wu is the corresponding author

¹<https://github.com/AdvancedAI-ComplexSystem/SmartCity/tree/main/CFLight>



This work is licensed under a Creative Commons Attribution 4.0 International License.
KDD '26, Jeju Island, Republic of Korea

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2258-5/2026/08

<https://doi.org/10.1145/3770854.3780308>

CCS Concepts

• **Computing methodologies** → **Computational control theory**; • **Applied computing** → **Transportation**.

Keywords

traffic signal control, reinforcement learning, counterfactual learning

ACM Reference Format:

Mingyuan Li, Chunyu Liu, Zhuojun Li, Xiao Liu, Guangsheng Yu, Bo Du, Jun Shen, and Qiang Wu. 2026. CFLight: Enhancing Safety with Traffic Signal Control through Counterfactual Learning. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1 (KDD '26)*, August 09–13, 2026, Jeju Island, Republic of Korea. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3770854.3780308>

Resource Availability:

The source code and artifacts of this paper are available at <https://doi.org/10.5281/zenodo.18041726>. The official version of the paper is accessible via <https://doi.org/10.1145/3770854.3780308>.

1 Introduction

Motivations. Road accidents have dire consequences, affecting victims, families, and societies worldwide. The World Health Organization reports an annual loss of 1.4 million lives and 20 to 50 million injuries due to road crashes in 2016 [25]. Intersections contribute significantly, causing about one-quarter of traffic fatalities and half of all injuries in the United States every year [6]. Meanwhile, traffic congestion remains an ever-escalating issue in contemporary urban environments, imposing substantial adverse impacts on various aspects of city life, including economic productivity and time consumption [28].

In recent years, Deep Reinforcement Learning (DRL) [33] has garnered increasing attention by adjusting traffic signal control (TSC)

to reduce traffic jams. By setting appropriate reward functions, such as minimizing vehicle travel time and maximizing traffic flow throughput. DRL continuously learns and optimizes the control strategies, adapting to changing traffic conditions in real time. Thus, DRL models make proper decisions regarding traffic light phases based on the current traffic flow state [3, 38]. Despite the many advances in RL model for TSC, most of the existing DRL models only consider traffic efficiency rather than safety. Now, the safe RL methods are proposed to ensure the agent’s learning and decisions align with safety constraints to avert hazardous outcomes [30]. It’s vital for real-world RL applications, safe RL holds the potential to tackle unresolved safety challenges in TSC. CounterFactual (CF) learning for safety is an emerging area of research in the field of AI and aims to improve the safety and robustness of AI systems by considering “what-if” scenarios during the training process [10]. Thus, CF learning provides some valuable tools for studying safety. Researchers often conduct CF policy evaluations before deploying any policy to the real world by using CF reasoning to reduce unsafe events.

The unsafe problems of intersections and the potential of CF learning for safe RL motivate us to develop a safe RL method for TSC to ensure intersection safety and efficiency at the same time.




Figure 1: The red vehicle, moving at high speed using permitted-only phase, collides with the opposing left-turning vehicle in permitted-only phase. We rewind to the state of pre-collision, and shift to protected-only phase to reduce collisions by avoiding the approach of hostile vehicles

Challenges. In the real-world, many intersections use a permitted-only phase, allowing both straight and left turns. While this improves efficiency, it increases the risk of collisions between straight-moving and left-turning vehicles. Alternatively, the protected-only phase, which separates left turns from straight movements, reduces collision risk but significantly decreases traffic efficiency. These current safe RL and TSC methods [5, 7] are still unsatisfactory in reducing the number of collisions and lack interpretability. Meanwhile, most safe RL methods are developed for specific domains or applications, lacking a universal framework for rapid research

and development of specific safe RL solutions. In the TSC field, few methods have been proposed to solve safety problems.

Contributions. In this study, we design a CF framework to improve the safe RL model by answering the CF question: “What if, when an unsafe event occurs, we backtrack to perform alternative actions, will this unsafe event still occur in the subsequent period?”. Furthermore, we effectively achieve a balance between traffic efficiency and safety, we have specifically implemented this method in TSC scenarios (as shown in Figure 1) via leveraging this CF framework. Our contributions are summarized as follows:

- We propose the first instantiation of the causal learning for TSC by answering counterfactual questions to improve the safety of TSC. Specifically, we design structure casual models to predict the result after executing different actions and collect the CF result to enrich the data and improve performance.
- We develop a CF+“X” safe RL framework that integrates CF components with various “X” modules. Based on this framework, we introduce a novel algorithm, CFLight, which effectively addresses safety events and incorporates a new reward function to improve both safety and action interpretability.
- Theoretical analysis shows that CFLight, augmented with CF data, converges reliably. Experimental results demonstrate that CFLight significantly reduces collisions, achieving up to a 93.1% reduction in collision rates compared to 3DQN methods, while maintaining near-zero incident rates.

Furthermore, our “CF+X” framework serves as a generalized safe RL method that can be applied in various other domains.

2 Background

2.1 Counterfactual learning

CF learning, or counterfactual estimation or reasoning, is a fundamental concept utilized across diverse fields such as machine learning, causal inference, and decision-making. It has recently gained significant attention leverage unreal data instances [16, 21, 34]. The origin of CF learning traces back to the pioneering work in 1974 [31], where fundamental concepts such as potential outcomes and causal inference laid the groundwork for subsequent research. The twin network method [27] visually represents reasoning with two networks: one for the factual world and one for the counterfactual (imaginary) world. They share the same structure, but the intervened variables are removed in the counterfactual world. Researchers have successfully used this technique in the health-care sector to estimate the causal effects of medical treatments and interventions [29]. CF learning has also yielded significant breakthroughs in recommender systems, as seen in methods like CF matrix factorization [37]. Furthermore, RL and CF learning amalgamation have impressive performance in optimizing decision-making processes [20, 23, 24].

The scope applications of RL expands into safety-critical domains, ensuring the robustness and safety of RL agents [8]. In the early stages of addressing RL safety, primarily directed toward defending against and it reveals such attacks exploit learned policies, leading to unintended and potentially harmful behaviors within RL-based systems [32]. The event of agent safety within RL has

received significant attention, particularly in critical domains like healthcare and autonomous vehicles. A comprehensive framework has been proposed to address the agent safety problem in RL, ensuring agents prioritize safety to mitigate potential risks [8]. For TSC, rare applications are focused on safety rather than efficiency.

Structural Causal Models (SCMs) serve as a formal framework in causal inference and statistics to represent and analyze causal relationships among variables. In a SCM, the causal relationships among variables are explicitly depicted. The set of variables, denoted as $X_1, X_2, \dots, X_n$, is associated with structural equations:

$$X_i = f_i(Pa(X_i), U_i), \quad (1)$$

where: $Pa(X_i)$ signifies the parents of X_i in the causal graph, $f_i(\cdot)$ represents a deterministic function describing how X_i is generated based on its parents and U_i stands for the error term, and it represents the unobserved factors affect X_i . The CF Causal Effect (CCE) is formulated as the disparity between the potential outcome under an intervention and the observed outcome without the intervention: $CCE = \tilde{Y} - Y$ where Y signifies the observed outcome, and $\tilde{Y}$ represents the CF outcome. The core objective of CF learning lies in estimating the CF outcome $\tilde{Y}$ or the causal effect, utilizing observed data $\mathcal{D}$ in conjunction with the causal model.

2.2 Traffic Signal Control

Due to the continuous advancements in urban transportation, research on TSC has been advancing progressively since its inception. The fixed-time method [14] cannot make optimal signal choices in response to real-time traffic conditions. Self-Organizing Traffic Lights (SOTL) [4] decides phase maintenance or changing by using predefined rules and real-time traffic conditions. Given constraints of the traditional method, the RL TSC dynamically adjusts signal phases in real time based on current traffic, optimizing performance [9, 12, 26]. CoLight [35] employs a graph attention and significant optimization through neighboring intersection data. Advanced-CoLight [38] further advances TSC performance. SafeLight [5] integrates a safe component into RL, reducing collision rates.




Figure 2: Illustration of a four-way intersection.

Traffic network. A traffic network is composed of a series of intersections ($I_1, \dots, I_N$), with each intersection hosting a variety of

roads and each road further accommodating multiple lanes.

Traffic movement. Traffic movement is the vehicle motion within an intersection that involves lane changes upon entry and exit (as shown in Figure 2(a)).

Traffic signal phase. A Traffic Signal Phase (TSP) specifies vehicle movements at an intersection during a defined interval. It designates proceeding, yielding, or stopping lanes, each linked to a signal indication. Networks use 2, 4, or k phases, denoted as $a = 1, 2, \dots, k$. Figure 2(b)(c) shows 6 phases. To enhance traffic efficiency, most cities have established a set of regulations that focus on optimizing left-turn scenarios. Several left-turn phasing guidelines were built by Federal Signal Timing Manual [14]. These regulations encompass two primary modes:

- **Permitted-only:** allows both straight-ahead and left turns.
- **Protected-only:** permits only straight-ahead movements while prohibiting left turns.

Although the “permitted-only” traffic signal mode aims to improve traffic efficiency, it inadvertently increases the risk of collisions. This risk arises from faster-moving vehicles on a permissive signal road colliding with vehicles making left turns, posing a critical safety hazard.

To find a solution to this unsafe scenario, the “protected-only” signal mode has been designed to restrict left turns on the current road and eliminate conflicts with oncoming traffic. This aims to decrease collisions on risky roads, enhancing overall traffic safety.

3 Methods

In this section, we introduce the CF+“X” framework, which integrates the CF and X modules to enhance the safety of RL. We describe the creation of CF trajectories and present the safe RL algorithm CFLight for TSC by implementing the CF+“X” framework. Figure 3 shows the overview.

3.1 CF+“X” Framework

The CF+“X” framework comprises two primary modules: the CF module and the X module. The CF module employs a SCM model to collect CF trajectories, while the X module integrates a Safe component (e.g., collision penalty in reward) and an Efficient-RL component (e.g., throughput or vehicle waiting time optimization in reward).

3.1.1 CF Module. The CF module consists of CF trajectories and the CF objective, which together serve as a vital mechanism for improving RL safety by incorporating CF reasoning into the learning process. CF trajectories offer additional safety-related experiences from which both the TSC agent and the SCM model can learn. The CF objective guides the optimization process during training, encouraging safer decision-making in the TSC model.

SCM Model. Assume that the next state S' and the CF result R are governed by the following SCM:

$$S' = \mathcal{M}_1(S, A, U_s), \quad R = \mathcal{M}_2(S, A, U_r), \quad (2)$$

where S and A denote the current state and action, respectively. The variable U represents a noise term, which may include factors such as weather conditions and transmission variability. This noise is an unobserved variable that is assumed to be independent of both S and A . The model $\mathcal{M}$ functions as a generative model, capable

of predicting unknown outcomes in real-world tasks. The SCM is trained using data stored in the experience buffer of the TSC agent and is used to collect CF trajectories.

CF Trajectories. We outline the process for collecting CF trajectories $(s, a_{cf}, s_{cf}, r_{cf})$. We begin by recording the states associated with all unsafe events prior to their occurrence, using a simulator denoted as *env*. Next, the simulation is rewound to the state immediately before each unsafe event. From this point, given alternative CF actions a_{cf} , the SCM model $\mathcal{M}$ is used to predict the corresponding CF reward r_{cf} and the resulting CF state s_{cf} . Finally, the outcomes before and after the rewind are compared in order to construct the CF objective.

CF Objective (Reward, Loss, and Q Value). We introduce the concept of a CF objective to enhance safety, defined as:

$$CF = f(R_{CF}, R_{AC}), \quad (3)$$

where R_{CF} represents the outcome resulting from the execution of the CF action a_{cf} , while R_{AC} denotes the outcome following the execution of the real world intended action. The function f is designed to measure the distance between R_{CF} and R_{AC} . Additionally, R_{CF} can be incorporated as a separate term in the overall optimization objective to explicitly account for CF considerations during learning.

3.1.2 X Module. The X Module consists of two primary components: the Safe component and the Efficient-RL component. The Safe component is designed to incorporate basic safety considerations into the RL model, while the Efficient-RL component represents the efficiency-driven part of the RL model. Different configurations of Safe and Efficient-RL components can be flexibly combined to form the integrated X Module.

Safe component. The Safe component focuses on embedding safety considerations into the design of the reward, loss, and Q functions within the RL framework. In the case of Safe Reward (*Safe_r*), safety concerns influence the reward function in a manner similar to reward shaping; when the safety model detects an unsafe action, a safety constraint is introduced to guide the learning process. For Safe Loss (*Safe_l*), the RL model is enhanced with an additional task by integrating the safe component into the loss function, enabling multi-task learning that optimizes multiple objectives simultaneously. In the case of Safe Value (*Safe_v*), the RL model is augmented by incorporating the safe component into the Q function, thereby establishing a multi-task learning framework that concurrently optimizes for both performance and safety.

Efficient-RL component. The Efficient-RL component utilizes advanced RL models aimed at optimizing efficiency. It can either maintain consistency with the original algorithm or adopt novel approaches to improve performance and adaptability.

3.2 CFLight

We propose CFLight, a method built upon the “CF+X” framework, designed to enhance both the efficiency and safety of TSC, as illustrated in Figure 3.

3.2.1 CF Module Implementation. SCM Model. A Bidirectional Conditional Generative Adversarial Network (BiCoGAN) [11] is used to approximate the SCM model. The core idea is to use a generator network to create fake samples that can deceive the

adversarial network. The generator network $\mathcal{M}$ is employed to estimate the next state S' and reward R , conditioned on (S, A, U_s, U_r) . Encoder is an inference mapping from S' and R to $(S, A, \hat{U}_s, \hat{U}_r)$.

$$S, A, \hat{U}_s, \hat{U}_r = E(S', R), \quad \hat{S}', \hat{R} = \mathcal{M}(S, A, U_s, U_r) \quad (4)$$

Using the data from the experience buffer, the SCM model is updated using the following formula:

$$\mathcal{L} = \min_{\mathcal{M}, E} \max_D V(D, \mathcal{M}, E) = \mathcal{L}_{DE} + \mathcal{L}_{DM} + \beta \mathcal{L}_{SR} + \lambda \mathcal{L}_{mono}, \quad (5)$$

where $\mathcal{L}_{DE} = \mathbb{E}[\log D(S, A, R, S', \hat{U}_s, \hat{U}_r)]$ and $\mathcal{L}_{DM} = \mathbb{E}[\log D(S, A, \hat{R}, \hat{S}', U_s, U_r)]$ are GAN-based loss functions used to optimize the discriminator, encoder, and generator models, helping to better distinguish between real and generated transitions. The loss $\mathcal{L}_{SR} = \|(S', R) - (\hat{S}', \hat{R})\|_2$ encourages the generator to produce states and rewards that more closely approximate the ground truth, using the L_2 norm distance. The monotonicity penalty, $\mathcal{L}_{mono} = \sum_l \sum_{i,j} \max(0, -W_{l,i,j})$, enforces the monotonicity of the generator $\mathcal{M}$, where l denotes the layer index and $W_{l,i,j}$ represents the weights of $\mathcal{M}$. The hyperparameters λ and β control the trade-off among the loss components.

THEOREM 1. Suppose R and S' satisfy the below causal model:

$$S' = \mathcal{M}_1(S, A, U_s), \quad R = \mathcal{M}_2(S, A, U_r), \quad (6)$$

where $U_s, U_r \perp\!\!\!\perp (S, A)$, $\mathcal{M}_1, \mathcal{M}_2$ is smooth and monotonic in U_s, U_r for fixed value of S, A . Suppose we have observed (S, A, R, S') , then for the CF action a_{cf} , the following CF outcome:

$$S'_{A=a_{cf}}, R_{A=a_{cf}} \mid S = s, A = a, S' = s', R = r, \quad (7)$$

is identifiable.

Proof: Given $S' = \mathcal{M}_1(S, A, U_s)$ and the observed values $S = s$, $A = a$, and $S' = s'$, and assuming that $\mathcal{M}_1$ is monotonic in U_s , we can invert the function to obtain $\hat{u}_s = \mathcal{M}_1^{-1}(s', s, a)$. Setting $A = a_{cf}$, and noting that $U_s \perp\!\!\!\perp (S, A)$, the noise term $\hat{u}_s$ remains unchanged. According to Theorem 1 in [21], we can then compute the CF next state as $S'_{A=a_{cf}} = \mathcal{M}_1(s, a_{cf}, \hat{u}_s)$. By the same reasoning, the CF reward is given by $R'_{A=a_{cf}} = \mathcal{M}_2(s, a_{cf}, \hat{u}_r)$. Therefore, the CF outcomes are identifiable.

According to Theorem 1, we obtain the CF outcomes $s_{cf} = \mathcal{M}_1(s, a_{cf}, r_s)$ and $r_{cf} = \mathcal{M}_2(s, a_{cf}, r_r)$. This indicates that by changing to the CF action a_{cf} , we can generate the CF next state s_{cf} and the CF reward r_{cf} . Finally, we use the loss function defined in Equation 5 to optimize the SCM model $\mathcal{M}$.

CF Trajectories. Algorithm 1 shows the CF trajectories construction (CTC) step. CF thinking infers what would have happened if an agent had taken a different action. Given a sample (s, a_{cf}, u) , a CF action a_{cf} is randomly drawn from the action space and, together with state s and noise u , passed to the generator $\mathcal{M}$ to predict the CF reward r_{cf} and next state s_{cf} .

CF Objective. We design different instances of the function f in Equation 3 for reward, loss, and Q value, respectively, to define the CF Objective. Here, r_{ac} and r_{cf} represent the rewards obtained by executing the action and the CF action, respectively. These different CF objectives encourage the agent to learn safer strategies. To emphasize the difference between real-world and CF collisions,




Figure 3: The framework of CFLight consists of two modules: the X Module and the CF Module. During the TSC process, pre-collision states s are continuously collected. When a collision occurs, the system rewinds to the recorded pre-collision states and initiates the CF trajectory collection process. In the CF Trajectories Construction (CTC) step, different actions a_{cf} are evaluated by the SCM model M to predict potential collisions, the resulting next state s_{cf} , and to collect the CF trajectory $(s, a_{cf}, s_{cf}, r_{cf})$. Finally, the CF Objective is formulated, the CF trajectories are added to the replay buffer, and both the TSC agent and the SCM model are updated.

Algorithm 1: Counterfactual Trajectories Construction (CTC)

```

1 Function CTC(SCM models  $M_1$  and  $M_2$ , Pre-collision state  $s$ ,
   Real-world reward  $r_{ac}$ ):
2   trajectory = [ ];
3   for  $a_{cf} = 1$  to  $n$  do
4     Rewind pre-collision state  $s$  and initialize noise
        $u_s, u_r$ ;
5     Predict  $s_{cf} = M_1(s, a_{cf}, u_s)$ ,  $r_{cf} = M_2(s, a_{cf}, u_r)$ ;
6     Construct  $CF_r$  by Equation 8,  $CF_l$  by Equation 9 and
        $CF_q$  by Equation 10;
7     trajectory.append( $s, a_{cf}, s_{cf}, r_{cf}$ );
8   end
9   return trajectory;

```

we define the following reward-based CF objective:

$$CF_r = r_{cf} - r_{ac}, \quad (8)$$

where $r_{ac} = w_1 \cdot Safe_r + w_2 \cdot Eff_r$ captures the reward composed of the number of collisions ($Safe_r$) and the total vehicle waiting time (Eff_r), both observed after executing the pre-collision action a . In the computation of the CF reward, $r_{cf} = w_1 \cdot Safe_r^{cf} + w_2 \cdot Eff_r^{cf}$, various CF actions a_{cf} are explored to determine the number of CF collisions $Safe_r^{cf}$ and the efficiency measure Eff_r^{cf} .

This reward quantifies the causal effect of executing the CF action a_{cf} . Assuming that efficiency remains constant before and after the intervention, and that collisions decrease from 2 (in the real-world observation) to 0 (in the CF scenario), the resulting reward difference becomes $CF_r = 0 - (-2) = 2$. Such a CF trajectory

(s, a_{cf}, r_{cf}) with an improved reward of 2 can substantially enhance the model's learning performance for safety improvement, and conversely, a negative reward could indicate actions to avoid.

The distributional distance between the CF action and the action is formalized by:

$$CF_l = D_{KL}(A_{AC}(\pi_H(s)) || A_{CF}(s, a_{cf}; \theta)), \quad (9)$$

where a collision occurs upon executing a real-world action, and we obtain the unsafe advantage distribution A_{AC} generated by the policy $\pi_H(s)$, which takes the current traffic state as input and determines whether an action is unsafe. We capture collision-related information, trace it back to the CF scenario, and explore alternative safe actions to derive the safe advantage distribution $A_{CF} = Q(s, a_{cf}; \theta) - V(s)$, measuring the discrepancy via cross-entropy.

The occurrence of a collision in the CF world also introduces a novel objective aimed at risk mitigation. We explore the CF action a_{cf} and evaluate its corresponding value CF_q as:

$$CF_q = Q_{cf}(s, a_{cf}), \quad (10)$$

where Q_{cf} is a network used to assess the safety value of the CF action a_{cf} . This value contributes to the overall Q-value computation. The CF reward CF_r is used to update Q_{cf} based on the Bellman Equation 11.

3.2.2 CFLight. We introduce CFLight-R, a variant of the CFLight algorithm based on the "CF+X" framework. The CF Objective utilizes CF_r , the Safe module incorporates $Safe_r$, and the Efficient-RL module uses Eff_r , which represents the total vehicle waiting time. The components of CFLight-R can be succinctly represented as a combination: $CF_r + Safe_r + Eff_r$.

- **State.** The state S contains the vehicle speed and location.

Algorithm 2: CFLight Training (CFLight-R, CFLight-Loss, CFLight-Q)

Input: Input epoch times K , replay buffer D , D_{CF} , max steps $maxSteps$, pre train steps $preTrainSteps$, random epsilon eps , simulator env and SCM models $\mathcal{M}$.

```

1 for  $i = 1$  to  $K$  do
2    $s = env.getState()$ ;
3   Initialize efficient RL and safe component and get  $Q$ 
    network;
4    $totalSteps = 0$ ;
5   for  $j = 1$  to  $maxSteps$  do
6      $totalSteps++ = 1$ ;
7     if  $random\ number < eps$  then
8       Explore random action;
9     else
10      Choose action  $a$  based on  $Q$  network;
11    end
12    In real world, execute the action  $a$  by  $env$ , get  $Safe_r$ ,
    real world trajectory  $(s, a, r, s')$  and add it into
    replay buffer  $D$ ;
13    if  $totalSteps > preTrainSteps$  then
14      Sample trajectories from  $D$  and  $D_{CF}$ ;
15      Update  $Q$  by Equation 11(CFLight-R),
        Equation ??(CFLight-Loss) or Equation ??
        (CFLight-Q)
16    end
17  end
18  for  $i \% trainGap == 0$  do
19     $D_{CF}.clear()$ ;
20    Update  $\mathcal{M}$  by Equation 5;
21    if  $collision$  then
22      Go to CF world to collect CF trajectories
         $(s, a_{cf}, r_{cf}, s_{cf})$  according to CTC function;
23      Add CF trajectories into replay buffer  $D_{CF}$ ;
24    end
25  end
26 end

```

- **Action.** The action A is a combination of the **permitted-only** and **protected-only** phases.
- **Reward.** The reward r takes the value R_{AC} in the real world and CF_r in the CF scenario.

Finally, we update the network through 3DQN and Q is updated according to the Bellman function:

$$Q(s, a) \leftarrow Q(s, a) + \alpha \left[R + \gamma \max_{a' \in \mathcal{A}} Q'(s', a') - Q(s, a) \right] \quad (11)$$

Using CF data augmentation, the following lemma ensures that the Q -function will achieve the optimal value [21]:

LEMMA 2. *Given the transition dynamics, Q -learning on the counterfactually augmented data set converges with probability one to the optimal value function Q^* , as long as the state and action spaces are finite, and the learning rate α_t satisfies $\sum_t \alpha_t = \infty$ and $\sum_t \alpha_t^2 < \infty$.*

Algorithm 2 shows a specific implementation based on CF module CF_r , X module ($Safe_r$ as safe component; Eff_r as efficient-RL component). The specific CFLight easily accommodates different CF and X modules.

4 Experiments

In this section, we concentrate on non-periodic acyclic phase switching and extensively experiment with synthetic and real-world datasets. These experiments are conducted in the microscopic traffic simulator, SUMO [15]. We set the collision rate to 40% under time-variant traffic flow.

4.1 Datasets

We experiment with a synthetic intersection [19] and a real-world intersection [22].

Synthetic: The synthetic intersection comprises four methods and twelve one-way vehicular movements with synthetic traffic flow adopted into SUMO.

Real-world: The real-world dataset is derived from the Cologne, Germany intersection, incorporating real traffic patterns into the SUMO environment. This intersection features 8 approaching lanes. In both cases, the SUMO simulation environment assumes accurate vehicle detection, consistent deceleration rates (modifiable as needed), and permitted left turns within the timing plan.

4.2 Compared Methods

In this study, we compared four types of RL methods for intersection control. These methods include:

Traditional method: Fixed-time [14].

RL methods: 3DQN [19], IPPO, Syn-R [13], Syn-Q [7], CPO [1] and SafeLight-Loss [5].

Our methods: CFLight-1 corresponds to $CF_r + Safe_r + RL_{3DQN}$. CFLight-2 corresponds to $CF_l + Safe_l + RL_{3DQN}$. CFLight-3 corresponds to $CF_q + Safe_v + RL_{3DQN}$.

All methods share the same action space, which includes both protected-only and permitted-only phases. We evaluate these TSC methods from the perspectives of efficiency and safety. The efficiency measure includes Average Delay (average vehicle waiting time) and the number of vehicles throughput within a period of time, while the safety measure is the sum of collisions over time.


4.3 Results

We average the results from the last 10 rounds across two benchmark datasets, with key outcomes presented in Table 1. On the synthetic dataset, CFLight-Loss reduces collision rates by 93.1% compared to 3DQN, outperforming current state-of-the-art (SOTA) methods. On the real-world dataset, the three CFLight variants achieve an average 32.6% reduction in collisions over non-CF methods, with CFLight-R emerging as the top-performing solution based on collision metrics.

Figure 4 illustrates the training curves of the model after incorporating the CF module. It can be observed that the training curve for collisions with the CF module shows an improvement compared to the one without the CF module, demonstrating the effectiveness of our approach.

Table 1: Performance of all methods on synthetic and real-world datasets. Green colors show the percentage improvement when adding the CF module. Red colors indicate the percentage of degradation.

Synthetic — Acyclic Action										
	Safelight-Act	CPO	Syn-R	CFLight-R (Improved)	Safelight-Loss	CFLight-Loss (Improved)	Syn-Q	CFLight-Q (Improved)	3DQN	Fixed-time
Average Delay (s)	4.73 ±1.04	22.16±3.23	9.25±1.03	8.67±1.07 (+6.3%)	30.29±1.28	29.89±2.15(+1.3%)	8.54±0.14	8.87±1.11(-3.8%)	4.24±0.13	31.22±1.34
Throughput	1775.3 ±52.35	1760.2±43.31	1784.8±50.03	1802.8±80.25 (+1.0%)	1106.2±30.05	1107.2±30.05(+0.01%)	1753±40.09	1781.3±40.07(+1.59%)	1785.6±0.04	1782±0.05
Collision Count	15.9 ±6.21	13.2±3.53	3.5±1.16	2.5±1.05(+28.6%)	1.8±0.15	1.3±0.16 (+27.7%)	5.7±4.33	4.6±2.12 (+19.2%)	19±0.15	7.6±0.16
Real-world — Cologne, Acyclic Action										
Average Delay (s)	23.32 ±3.12	44.06±8.24	12.59±3.14	20.27±3.13 (-61.1%)	66.86±1.19	47.74±3.56 (+40.0%)	16.20±5.18	22.28±6.16(-37.5%)	5.54±1.12	67.25±3.25
Throughput	2013 ±0.04	2014.6±0.04	2014±0.07	2014±0.05	1733.5±5.03	2014.7±0.04 (+16.2%)	2014.4±0.05	2014.8±0.04(+0.01%)	2014±0.05	2015±0.04
Collision Count	13.2 ±6.32	41.3±8.56	20.3±6.21	5.3±4.56 (+73.8%)	13.4±5.34	13.3±4.23 (+0.01%)	13.06±2.51	9.9±3.34(+24.1%)	40.07±4.13	35.1±3.32

**Figure 4: The figures show the training curve for safety and efficiency for using CFLight in the Cologne dataset.**

4.4 Safety and Efficiency Trade-off


In this section, we illustrate the influence of adjusting weights on both safety and efficiency in synthetic dataset. Leveraging the CFLight-1 algorithm, we conduct 800 training rounds and establish the benchmark value by averaging the results from the last 10 rounds, as detailed in Figure 2. The outcome reveals the emergence of a Pareto optimal surface. Importantly, a harmonious equilibrium is attained when the weights assigned to safety and efficiency closely align. These insights are derived from testing on a synthetic dataset using CFLight-R.

4.5 Ablation Analysis

We conduct ablation experiments on real datasets involving CF_r , CF_r+3DQN , $3DQN$, $Safe_r$ and CF_r+Safe_r configurations to assess the efficacy of the CF module. Notably, employing only the CF reward CF_r leads to a 45.5% reduction in collision incidents compared to $3DQN$, but with a huge increase in waiting time because it works without optimization in waiting time. By using CF_r+RL_{3DQN} method, collision rates decrease by 16.3%, and waiting times improve significantly compared to the standalone CF module CF_r . Compared to the $Safe_r$ method, $CF+Safe_r$ module the collisions reduce about 36.2% which shows the validity of our CF module.

4.6 Case Study

In this section, we present a visualization of the next state and reward predicted by the SCM before and after a collision event.

**Figure 5: In the real world collision case, execute the permitted-only ($a = 0$) phase for state, next state, and reward. The third subgraph shows safety (top) and efficiency (bottom).**

The traffic signal phase prior to the collision is set to permitted-only, and we evaluate the state and reward variations when the SCM predicts outcomes for a protected-only phase. The state is modeled as a 60×60 matrix, with the x - and y -axes representing vehicle coordinates and the heatmap values indicating normalized vehicle speeds. The CF reward is computed as $CF_r = r_{cf} - r_{ac}$. In the third subplot on the right side of the figure, the upper half of the reward plot reflects safety (higher values improve safety) and the lower half represents efficiency. The results reveal that adopting the protected-only phase, as predicted by the SCM, significantly increases the safety reward, highlighting its effectiveness in improving the learning of the CF trajectory. Moreover, the CF next states




Figure 6: In the CF world collision case, execute the protected-only ($a_{cf} = 3$) phase for state state, next state and reward CF_r predict by SCM $\mathcal{M}$.




Figure 7: The left image shows the safety-efficiency trade-off, while the right compares protected-only and permitted-only phases based on average waiting time and collision counts in the Cologne dataset.

closely resemble the real-world next states in structure, validating the robustness and accuracy of our proposed method.

5 Discussion

(1) **How does the CF module’s SCM with GAN approximation operate and what are its benefits?** In CFLight, to realistically simulate the aftermath of traffic collisions, we configure the SUMO simulator to introduce a delay after a collision vehicles involved are briefly held before being rerouted to their destinations. This setup leads to temporary congestion and increased waiting time, aligning with real-world traffic scenarios. Moreover, CFLight introduces a SCM and employs a BiCoGAN to generate counterfactual trajectories under alternative actions, enabling counterfactual data augmentation. This technique creates diverse and realistic hypothetical scenarios, enhancing data diversity and model generalization. As a result, it improves sample efficiency and overall performance without requiring additional real-world data, and comes with theoretical convergence guarantees (Lemma 2).

(2) **Why not just use protected-only phases?** We do experiments using only protected-only in the real-world Cologne (1 intersection) dataset as shown in Figure 7, and we see that passage efficiency is greatly reduced. The number of collisions is greatly increased by using only permitted-only phase. In the real world, we observe that there are only two lanes on each road and each lane contains two directions, and the protected-only phase only allows passage in one direction, which may be the reason for the reduced efficiency. The original intention of our CFLight design is to flexibly switch between these two phases, ensuring both efficiency and safety in

Table 2: CFLight performance on 100 intersections of synthetic dataset.

Synthetic 100 Intersections Acyclic Action			
	CFLight-R	Syn-R	Fixed-time
Average Delay (s)	3098.90	3162.37	32708.37
Throughput	6831	6852	7065
Collision Count	47	55	699

Table 3: Multiple intersection experiments in Cologne3 and Cologne8.

Real-world Cologne3 Acyclic Action			
	CFLight-R	IPPO	Fixed-time
Average Delay(s)	15.21±2.16	7.37±3.13	65.9±5.32
Collisions	14.07±3.14	30.07±3.65	22±4.13
Real-world Cologne8 Acyclic Action			
Average Delay(s)	12.34±4.15	5.01±5.16	110.33±5.34
Collisions	32.31±5.34	64.31±6.67	47±6.78

traffic.

(3) **Are there any other scenario experiments for our proposed method?** We conduct lane-changing experiments in an autonomous driving setting to evaluate the CF+X framework. In this task, a vehicle must change across two lanes while assessing collision risk. The state includes the speed of nearby vehicles, the action indicates whether the vehicle changes lanes, and the reward reflects whether a collision occurs. We train an RL agent using CF+X with CF_r for 100 rounds and compare it to a baseline without CF+X. As shown in Figure 8, the left graph shows that CF+X reduces the average collision rate to 0.09, compared to 0.34 without the framework, representing a 73.5% reduction in collisions.




Figure 8: The collision result of using our proposed CF+X and without using CF+X under the lane change scenario. Our proposed method improved 73.5% performance in terms of average collisions on 100 rounds

(4) **Can our method be applied to multiple intersections?** We used the TSC benchmark [2] datasets with multiple intersections (3, 8 and 100 intersections) and conduct experiments with IPPO as the backbone, CFLight uses CF_r as the CF objective. Our results, as shown in Table 2 and Table 3, indicate that CFLight effectively enhances performance, demonstrating the effectiveness of our method.

(5) **What are the strengths and limitations of the various CFLight extension methods? In which scenarios are these methods most appropriately applied?** CFLight-Loss leverages SCM-derived action advantages and jointly optimizes safety constraints and efficiency in a unified loss function, demonstrating flexibility but higher sensitivity to hyperparameters. CFLight-Q enhances interpretability by directly optimizing Q-values under CF scenarios, yielding near-zero collision rates, albeit at the cost of increased complexity and reliance on Q-network convergence. Each variant offers distinct strengths and trade-offs, making them suitable for different deployment priorities. We also create a table comparing the advantages and disadvantages of CFLight against other safe methods.

(6) **How do you identify which action/state in the trajectory was causal?** To identify causal actions/states, we compare outcomes by backtracking to pre-collision states and executing safe actions. If a collision is avoided, it reveals causal relationships. Our CTC generates tuples $(s, a_{cf}, r_{cf}, s_{cf})$ for capturing these causal relationships.

(7) **Why is it necessary to assume that $\mathcal{M}$ is monotonic in order to predict CF in an SCM?** The monotonicity assumption on $\mathcal{M}_1$ and $\mathcal{M}_2$ with respect to the exogenous variables U_s, U_r is essential for counterfactual prediction. Specifically, given observed data (S, A, S', R) and the structural equations:

$$S' = \mathcal{M}_1(S, A, U_s), \quad R = \mathcal{M}_2(S, A, U_r), \quad (12)$$

identifying the hidden variables U_s and U_r from the observations requires inverting $\mathcal{M}_1$ and $\mathcal{M}_2$. If $\mathcal{M}_1$ and $\mathcal{M}_2$ are monotonic in U_s and U_r respectively (for fixed S, A), then these mappings are invertible and the values of U_s, U_r are uniquely determined:

$$U_s = \mathcal{M}_1^{-1}(S, A, S'), \quad U_r = \mathcal{M}_2^{-1}(S, A, R). \quad (13)$$

This uniqueness is critical for generating consistent counterfactual outcomes under a hypothetical action a_{cf} , as it ensures that the same latent variables can be reused to simulate the counterfactual outcome:

$$S'_{cf} = \mathcal{M}_1(S, a_{cf}, U_s), \quad R_{cf} = \mathcal{M}_2(S, a_{cf}, U_r). \quad (14)$$

Without monotonicity, the inversion may be non-unique or ill-defined, making counterfactual inference ambiguous.

(8) **How are CF actions selected to ensure safety for CFLight-Loss?** CF actions are selected by evaluating their advantage using the learned SCM and Q-function. Specifically: We monitor the real-world or simulated trajectory and identify actions leading to safety violations. For each detected unsafe state-action pair (s, a) , we generate a set of alternative candidate actions $\{a_{cf}^{(1)}, a_{cf}^{(2)}, \dots\}$ and simulate their outcomes using the SCM:

$$s'_{cf} = \mathcal{M}_1(s, a_{cf}, u_s), \quad r_{cf} = \mathcal{M}_2(s, a_{cf}, u_r) \quad (15)$$

For each candidate action, we compute its CF advantage:

$$A_{CF}(s, a_{cf}) = Q(s, a_{cf}) - V(s) \quad (16)$$

We choose the CF action a_{cf} with the highest advantage, under the constraint that it does not trigger unsafe outcomes as predicted by the SCM (i.e., r_{cf} indicates safety). This process balances *efficiency* and *safety*, enabling the agent to improve policy performance while avoiding unsafe behaviors.




Figure 9: Performance under noise of CFLight-R.

(9) **How sensitive is the method to inaccurate SCM predictions under varying traffic conditions or noise?** Figure 9 shows that, CFLight-R exhibits moderate sensitivity to inaccurate SCM predictions under varying traffic conditions, with noise modeled as $N(0, scale)$ and states normalized to the $[0, 1]$ range. The Collisions Comparison shows that higher noise (0, 1) results in slightly more collisions and variability compared to lower noise (0, 0.5). Similarly, the Wait Average Comparison reveals a higher initial wait time and fluctuations with noise scale 1, compared to the smoother decline with scale 0.5, with the impact of noise diminishing over epochs. In the future, we will explore more control methods to improve the robustness of CFLight like [17, 18, 36].

6 Conclusion

In this study, we introduced a flexible “CF+X” safe reinforcement learning framework, integrating counterfactual learning with the “X” module consisting of adjustable safety and efficient RL components. Within this framework, we developed a novel algorithm called CFLight, addressing CF and safety concerns while smoothly integrating with existing RL methods for TSC. Our method was evaluated in the context of safe and efficient traffic management at intersections, showcasing an impressive 93.1% reduction in collision rates compared to DQN while concurrently enhancing traffic efficiency. This work holds promise for broader applications in safe reinforcement learning, with potential implications for autonomous systems and robotics, offering a robust solution to enhance safety in complex decision-making scenarios.

Acknowledgments

This work is supported by Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (Grant No. JYB2025XDXM910), National Natural Science Foundation of China (Grant Nos. 42595590, 42595593), Talent Scientific Fund of Lanzhou University (Grant No. 561120208), and Supercomputing Center of Lanzhou University. G. Yu contributed as an independent researcher in a personal capacity. This work was completed before Australia’s 2025 international collaboration policy updates. G. Yu received no funding or material support for this research, and no institutional resources, facilities, or work hours from his affiliation were utilized. This contribution is outside his institutional employment and adheres to all intellectual property regulations.

References

- [1] Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. 2017. Constrained policy optimization. In *Proceedings of the International Conference on Machine Learning (ICML)*.
- [2] James Ault and Guni Sharon. 2021. Reinforcement learning benchmarks for traffic signal control. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 1)*.
- [3] Chacha Chen, Hua Wei, Nan Xu, Guanjie Zheng, Ming Yang, Yuanhao Xiong, Kai Xu, and Zhenhui Li. 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. 3414–3421.
- [4] Seung-Bae Cools, Carlos Gershenson, and Bart D'Hooghe. 2013. Self-organizing traffic lights: A realistic simulation. In *Advances in applied self-organizing systems*. Springer, 45–55.
- [5] Wenlu Du, Junyi Ye, Jingyi Gu, Jing Li, Hua Wei, and Guiling Wang. 2023. Safelight: A reinforcement learning method toward collision-free traffic signal control. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 14801–14810.
- [6] FHWA. 2023. Intersection Safety. <https://highways.dot.gov/safety/intersection-safety/about> Accessed on: August 1, 2023.
- [7] Yaobang Gong, Mohamed Abdel-Aty, Jinghui Yuan, and Qing Cai. 2020. Multi-objective reinforcement learning approach for improving safety at intersections with adaptive traffic signal control. *Accident Analysis & Prevention* 144 (2020), 105655.
- [8] Shangding Gu, Long Yang, Yali Du, Guang Chen, Florian Walter, Jun Wang, Yaodong Yang, and Alois Knoll. 2023. A Review of Safe Reinforcement Learning: Methods, Theory and Applications. [arXiv:2205.10330 \[cs.AI\]](https://arxiv.org/abs/2205.10330)
- [9] Mengyu Guo, Pin Wang, Ching-Yao Chan, and Sid Askary. 2019. A reinforcement learning approach for intelligent traffic signal control at urban intersections. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 4242–4247.
- [10] Yimin He, Stephanie C. Payne, Xiang Yao, and Rachel Smallman. 2020. Improving workplace safety by thinking about what might have been: A first look at the role of counterfactual thinking. *Journal of Safety Research* 72 (2020), 153–164. doi:10.1016/j.jsr.2019.12.010
- [11] Ayush Jaiswal, Wael AbdAlmageed, Yue Wu, and Premkumar Natarajan. 2019. Bidirectional conditional generative adversarial networks. In *Computer Vision—ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III* 14. Springer, 216–232.
- [12] Junchen Jin, Xiaoliang Ma, and Iisakki Kosonen. 2017. An intelligent control system for traffic lights with simulation-based evaluation. *Control engineering practice* 58 (2017), 24–33.
- [13] Mohamed A Khamis and Walid Gomaa. 2014. Adaptive multi-objective reinforcement learning with hybrid exploration for traffic signal control based on cooperative multi-agent framework. *Engineering Applications of Artificial Intelligence* 29 (2014), 134–151.
- [14] P Koonce. 2008. Traffic signal timing manual. technical report. Kittelson & Associates, Inc., Portland, OR, USA, Tech. Rep. FHWA-HOP-08-024 (2008).
- [15] Daniel Krajzewicz, Jakob Erdmann, Michael Behrisch, and Laura Bieker. 2012. Recent development and applications of SUMO-Simulation of Urban MOBility. *International journal on advances in systems and measurements* 5, 3&4 (2012).
- [16] Jiahui Li, Kun Kuang, Baoxiang Wang, Furui Liu, Long Chen, Fei Wu, and Jun Xiao. 2021. Shapley counterfactual credits for multi-agent reinforcement learning. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 934–942.
- [17] Mingyuan Li, Jiahao Wang, Bo Du, Jun Shen, and Qiang Wu. 2025. FuzzyLight: A Robust Two-Stage Fuzzy Approach for Traffic Signal Control Works in Real Cities. [arXiv preprint arXiv:2501.15820](https://arxiv.org/abs/2501.15820) (2025).
- [18] Mingyuan Li, Jiahao Wang, Guangsheng Yu, Xu Wang, Qianrun Chen, Wei Ni, Lixiang Li, and Haipeng Peng. [n. d.]. RobustLight: Improving Robustness via Diffusion Reinforcement Learning for Traffic Signal Control. In *Forty-second International Conference on Machine Learning*.
- [19] Xiaoyuan Liang, Xunsheng Du, Guiling Wang, and Zhu Han. 2019. A deep reinforcement learning network for traffic light cycle control. *IEEE Transactions on Vehicular Technology* 68, 2 (2019), 1243–1253.
- [20] Tongtong Liu, Joe McCalmon, Thai Le, Md Asifur Rahman, Dongwon Lee, and Sarra Alqahtani. 2023. A novel policy-graph approach with natural language and counterfactual abstractions for explaining reinforcement learning agents. *Autonomous Agents and Multi-Agent Systems* 37, 2 (2023), 34.
- [21] Chaochao Lu, Biwei Huang, Ke Wang, José Miguel Hernández-Lobato, Kun Zhang, and Bernhard Schölkopf. 2020. Sample-efficient reinforcement learning via counterfactual-based data augmentation. [arXiv preprint arXiv:2012.09092](https://arxiv.org/abs/2012.09092) (2020).
- [22] Hao Mei, Xiaoliang Lei, Longchao Da, Bin Shi, and Hua Wei. 2022. LibSignal: An Open Library for Traffic Signal Control. [arXiv preprint arXiv:2211.10649](https://arxiv.org/abs/2211.10649) (2022).
- [23] Thomas Mesnard, Théophane Weber, Fabio Viola, Shantanu Thakoor, Alaa Saade, Anna Harutyunyan, Will Dabney, Tom Stepleton, Nicolas Heess, Arthur Guez, et al. 2020. Counterfactual credit assignment in model-free reinforcement learning. [arXiv preprint arXiv:2011.09464](https://arxiv.org/abs/2011.09464) (2020).
- [24] Yusuke Narita, Kyohei Okumura, Akihiro Shimizu, and Kohei Yata. 2023. Counterfactual Learning with General Data-Generating Policies. *Proceedings of the AAAI Conference on Artificial Intelligence* 37, 8 (2023), 9286–9293.
- [25] World Health Organization et al. 2018. Global status report on road safety 2018: Summary (No. WHO/NMH/NVI/18.20). *World Health Organization* (2018).
- [26] Afshin Oroojlooy, Mohammadreza Nazari, Davood Hajimezhad, and Jorge Silva. 2020. Attendlight: Universal attention-based reinforcement learning model for traffic signal control. *Advances in Neural Information Processing Systems* 33 (2020), 4079–4090.
- [27] Judea Pearl. 2009. *Causality*. Cambridge university press.
- [28] Bob Pishue. 2021. 2021 inrix global traffic scorecard. In *INRIX (December 2021)*.
- [29] Mattia C. F. Proserpi, Yi Guo, Matthew Sperrin, James S. Koopman, Jae S. Min, Xing He, Shannan N. Rich, Mo Wang, Iain E. Buchan, and Jiang Bian. 2020. Causal inference and counterfactual prediction in machine learning for actionable healthcare. *Nat. Mach. Intell.* 2, 7 (2020), 369–375. doi:10.1038/s42256-020-0197-y
- [30] Faizan Rasheed, Kok-Lim Alvin Yau, and Yeh-Ching Low. 2020. Deep reinforcement learning for traffic signal control under disturbances: A case study on Sunway city, Malaysia. *Future Generation Computer Systems* 109 (2020), 431–445.
- [31] Donald B. Rubin. 1974. Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies. *Journal of Educational Psychology* 66, 5 (1974), 688–701.
- [32] Jianwen Sun, Tianwei Zhang, Xiaofei Xie, Lei Ma, Yan Zheng, Kangjie Chen, and Yang Liu. 2020. Stealthy and Efficient Adversarial Attacks against Deep Reinforcement Learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 34 (2020), 5883–5891.
- [33] Elise Van der Pol and Frans A Olthoek. 2016. Coordinated deep reinforcement learners for traffic light control. *Proceedings of Learning, Inference and Control of Multi-Agent Systems (at NIPS 2016)* (2016).
- [34] Sahil Verma, Varich Boonsanong, Minh Hoang, Keegan E. Hines, John P. Dickerson, and Chirag Shah. 2022. Counterfactual Explanations and Algorithmic Recourses for Machine Learning: A Review. [arXiv:2010.10596 \[cs.LG\]](https://arxiv.org/abs/2010.10596)
- [35] Hua Wei, Nan Xu, Huichu Zhang, Guanjie Zheng, Xinshi Zang, Chacha Chen, Weinan Zhang, Yanmin Zhu, Kai Xu, and Zhenhui Jessie Li. 2019. CoLight: Learning Network-level Cooperation for Traffic Signal Control. *Proceedings of the 28th ACM International Conference on Information and Knowledge Management* (2019).
- [36] Qiang Wu, Mingyuan Li, Jun Shen, Linyuan Lü, Bo Du, and Ke Zhang. 2023. Transformerlight: A novel sequence modeling based traffic signaling mechanism via gated transformer. In *Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining*. 2639–2647.
- [37] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. 2020. Adversarial Counterfactual Learning and Evaluation for Recommender System. *Advances in Neural Information Processing Systems* 33.
- [38] Liang Zhang, Qiang Wu, Jun Shen, Linyuan Lü, Bo Du, and Jianqing Wu. 2022. Expression might be enough: representing pressure and demand for reinforcement learning based traffic signal control. In *Proceedings of the 39th International Conference on Machine Learning*, Vol. 162. 26645–26654.