

Comparing Exploration Approaches in Deep Reinforcement Learning for Traffic Light Control

Yaniv Oren¹, Rolf A. N. Starre¹, and Frans A. Oliehoek¹

Technical University Delft, Mekelweg 5 Delft, the Netherlands {y.oren@student.,
r.a.n.starre@, f.a.oliehoek@}tudelft.nl

Abstract. Identifying the most efficient exploration approach for deep reinforcement learning in traffic light control is not a trivial task, and can be a critical step in the development of reinforcement learning solutions that can effectively reduce traffic congestion. It is common to use baseline dithering methods such as ϵ -greedy. However, the value of more evolved exploration approaches in this setting has not yet been determined. This paper addresses this concern by comparing the performance of the popular deep Q-learning algorithm using one baseline and two state of the art exploration approaches, and their combination. Specifically, ϵ -greedy is used as a baseline, and compared to the exploration approaches Bootstrapped DQN, randomized prior functions, and their combination. This is done in three different traffic scenarios, capturing different traffic profiles. The results obtained suggest that the higher the complexity of the traffic scenario, and the larger the size of the observation space of the agent, the larger the gain from efficient exploration. This is illustrated by the improved performance observed in the agents using efficient exploration and enjoying a larger observation space in the complex traffic scenarios.

Keywords: Reinforcement learning · Traffic optimization · Exploration.

1 Introduction

Traffic congestion is a global predicament. For instance, in the EU alone its cost is estimated to be 1% of the EU's GDP [15]. One approach for reducing this cost is optimization of traffic flows by improving traffic light control policies. To find such policies, reinforcement learning (RL) has a strong appeal as a paradigm that is able to find high performance solutions to sophisticated problems. Research has been done into the application of RL to the problem of traffic light control optimization in the past [1], [4], [13], [14], [18], often specifically concerning application of deep RL algorithms [4], [13], [18].

A fundamental principle of RL is exploration, and the balance between exploration and exploitation. Namely, how much does the agent explore its environment, versus how much it opts for actions that it expects to return the most cumulative reward. Many different exploration approaches exist for reinforcement learning [3], [6], [10]–[12]. These approaches often perform differently in different settings, in addition to having different computational costs [11], [12]. It has been shown that for some RL settings, simple exploration approaches such as ϵ -greedy are insufficient for RL to be able to perform well, or at all [11], [12], which can be caused by the reward function used and the complexity of the specific problem tackled. This illustrates the importance of identifying effective exploration techniques for specific RL settings. To the best of our knowledge, there has not been an attempt to investigate the importance of efficient exploration in deep RL in the setting of traffic light control.

This paper investigates a comparison between different exploration approaches in deep RL for traffic light control. This, to facilitate better deep RL by identifying the value of evolved exploration approaches in this setting, such as higher sample efficiency, or higher final policy score. For that purpose, this paper compares the performance of the popular deep Q-learning algorithm (DQN) [10] using one baseline and two state of the art exploration approaches, and their combination. Specifically, ϵ -greedy is used as a baseline, and compared to the exploration approaches Bootstrapped DQN, randomized prior functions, and their combination. This is done in three different traffic scenarios, ranging from simplified to simulating real traffic, in order to investigate the effect of exploration in different traffic profiles.

This paper first introduces a theoretical background for deep RL and exploration. Second, a description of the exploration techniques compared, along with the modeling of traffic light control as an RL problem, are given. This is then followed by an explanation of the research methodology and the experimental setup, leading to a presentation of the results obtained and their analysis. Last, ethical and epistemic concerns are considered, implications of the work are discussed and conclusions are drawn.

Altogether, the results obtained suggest a link between the complexity of the traffic scenario, the amount of information accessible to the agent, and the gain from efficient exploration. This is illustrated by the improved performance observed in the agents using efficient exploration and enjoying a large observation space, in the complex traffic scenarios.

2 Background

This section introduces background information relevant to the work presented in this paper. First, an overview of reinforcement learning (RL), is given, laying the basis for an introduction to deep RL and a description of the deep Q-learning (DQN) algorithm [10] that follows. Last, the principle of exploration is explained, followed by an overview of dithering [12], deep and directed exploration.

2.1 Reinforcement learning

In RL, an agent operates in an environment. The environment provides information regarding the state the agent is in and what actions it can execute. The environment is usually described in the form of a Markov decision process (MDP), a stochastic control process often used to model decision making in partially stochastic domains. A Markov decision process is represented as a four-tuple, $M = (S, A, P, R)$, where S represents the state space, A the action space, P the transition function and R the reward function.

The agent interacts with the environment by observing a state $s_t \in S$, executing an action $a_t \in A$, and receiving a reward $r_t \in R$ for the action executed. The agent is attempting to learn a policy π , such that the expected reward over time is maximised.

In the Q-learning algorithm [17], the value of state-action pairs is estimated by the agent, using iterative Bellman updates: $Q_{t+1}(s_t, a_t) = Q_t(s_t, a_t) + \alpha[y_t - Q_t(s_t, a_t)]$, where α is the learning rate, and the target $y_t = r_t + \gamma \max_a Q_t(s_{t+1}, a)$. s_{t+1} denotes the new state arrived at after choosing action a_t in state s_t , a any action available at state s_{t+1} , and $0 \leq \gamma \leq 1$, is a discount factor. In many scenarios however, the state-action pair space is too large for the computation or memorisation of each value $Q(s, a)$ to be tractable. To avoid this problem, function estimators such as neural networks can be used to estimate the Q value. This gives rise to the use of neural networks in reinforcement learning, and specifically the deep Q-Learning (or deep Q-networks) algorithm, commonly referred to as DQN [7], [10].

2.2 Deep reinforcement learning with DQN

The DQN algorithm uses deep neural networks to estimate a mapping from states to Q-values [10]. Instead of saving or computing each $Q(s, a)$ value separately, the algorithm learns a parameterized value function $Q(s, a; \theta_t)$. As a result, rather than learning the Q-values directly, the algorithm learns the parameter set θ of the Q-function. The previous Q-learning update then becomes:

$$\theta_{t+1} = \theta_t + \alpha(y_t - Q(s_t, a_t; \theta_t)) \nabla_{\theta_t} Q(s_t, a_t; \theta_t)$$

Here, $y_t = r_{t+1} + \gamma \max_a Q_t(s_{t+1}, a; \theta_t)$. To prevent instabilities, the DQN algorithm uses an additional network, termed the target network, θ^- [10]. The target network is the same as the regular, or online, network. However, it only updates every certain τ time-steps, by copying the parameters θ of the online network. This target network is used by the DQN algorithm in the target term, which becomes instead: $y_t = r_{t+1} + \gamma \max_a Q_t(s_{t+1}, a; \theta_t^-)$.

Double DQN [16] is a commonly used modification suggested to the original DQN algorithm, which aims to reduce overoptimism caused by estimation errors, which DQN is prone to. This is done by decoupling the selection of an action from its evaluation [16]. In vanilla DQN, in the term $y_t = r_{t+1} + \gamma \max_a Q_t(s_{t+1}, a; \theta_t^-)$, the agent uses the same network θ^- for both selecting and evaluating an action. The double-DQN algorithm proposes using the online network θ to choose the action, and the target network θ^- to evaluate the choice. The target term y_t used in the double DQN update then becomes:

$$y_t = r_{t+1} + \gamma Q_t(s_{t+1}, \arg \max_a Q_t(s_{t+1}, a; \theta_t); \theta_t^-)$$

2.3 Exploration in reinforcement and deep reinforcement learning

In order to find an optimal policy through experience alone, which is the general premise of RL, the agent must encounter the rewards that are part of an optimal policy at least once. This leads directly to a necessity to explore the environment - if the agent does not explore, how will it encounter valuable rewards that do not lie over its existing policy's path? However, the agent is also expected to efficiently converge into an optimal policy, and not only explore its environment. This leads to one of the fundamental principles of RL - exploration vs exploitation. Different approaches have been developed - ranging from dithering, random action choosing exploration [10], to more evolved notions such as deep and directed exploration [11], [12], and others. These approaches each attempt to achieve efficient exploration through different means - from simplicity of computation to effective analysis of the agent's knowledge and uncertainty.

Dithering exploration The common baseline exploration strategy used in DQN is a dithering exploration method, or ϵ -greedy. In ϵ -greedy, the agent takes a random action (i.e., explores) with probability ϵ , and with probability $1 - \epsilon$ the agent takes the best action according to its current Q value estimation. ϵ -greedy achieves state of the art performance against many popular benchmarks [12], [16]. However, as discussed in [12], in environments where rewards are scarce and distanced in the state-action space, and their values have a large spread, the dithering exploration of ϵ -greedy can take an exponentially long time to arrive at high-valued rewards. This raises the necessity for a more advanced type of exploration, that can be directed over over multiple time steps. These concepts have been coined 'directed exploration' and 'deep exploration' [12].

Deep & directed exploration Directed exploration attempts to improve the efficiency of the agent’s exploration by directing it. For example, to previously unexplored or under-explored areas of the state-action space. To achieve a measure of directed exploration, an uncertainty measure can be used - the more uncertain the agent is about the value of some state or action, the more it can prioritize exploration. For directed exploration to be effective, it often doesn’t suffice for it to be directed over one, or small number of time steps [12], but must be directed over multiple time steps. The term ‘deep exploration’ is used to describe such an exploration approach, that is directed over multiple time steps [12].

3 Exploration In Deep Reinforcement Learning for Traffic Light Control

This section will first motivate the choice to evaluate exploration techniques that focus on deep and directed exploration. Following that, the agent used in the experiments is outlined, and the exploration approaches it implements are described. Last, the modelling of traffic light control as an RL problem used in this paper is presented.

3.1 Motivation

This paper opts to specifically identify the value of exploration approaches that focus on achieving deep exploration, in the setting of traffic light control. These approaches have been chosen not only for being state of the art in this field, but also for their potential in this setting. In heavy traffic scenarios, suboptimal actions may carry a long term effect. They may immediately cause congestion, and once there, it may be difficult to return to less congested states [13]. Deep exploration may be able to improve the agent’s ability to escape such scenarios, by directing its exploration along a specific path. While this path will not necessarily pay in the short run, it may allow the agent to recover from congestion in the longer run. Additionally, the ability of the agent to more efficiently explore areas of the state action space that lie beyond areas plagued by negative rewards, as a result of employing deep exploration, may allow the agent to learn optimal policies that will otherwise be unlikely for a dithering agent to ever achieve.

3.2 The agent

The agent used in the experiments presented in this paper is a DQN agent, using the double-DQN [16] modification. The agent implements the following exploration mechanisms: ϵ -greedy, Bootstrapped DQN (BDQN) [12] and randomized prior functions [11]. The implementation used in this paper, based on [8] and modified for the setting of traffic light control, allows the agent to use any combination of the three different mechanisms listed above. ϵ -greedy has been described in section 2.3, and has been chosen as it is the common baseline exploration used in DQN [10],[16]. BDQN and randomized prior functions have been chosen as state of the art exploration approaches that aim to achieve efficient deep exploration, and in addition, for their ability to elegantly combine for even better exploration. BDQN and randomized prior functions are described below.

Bootstrapped DQN BDQN has been developed in an attempt to achieve a measure of deep exploration, discussed in section 2.3. In order to achieve deep exploration, BDQN approximates a distribution over Q-values, using a bootstrap. Bootstrapping is a technique used to approximate a population distribution from a sample distribution, using random sampling with replacement [5].

The bootstrap can be implemented efficiently using a shared neural network with several heads. The shared network’s role is to learn a feature representation, while each head is providing an independent Q-values estimation. A visualization can be found in figure 1. In learning, the algorithm randomly samples an estimator (a ‘head’) out of the bootstrap, and follows the policy which is optimal for that estimator for some number of steps greedily or ϵ -greedily. In the experiments done in this paper ϵ -greedy is used. The resulting experiences are gathered in a buffer, and are available for all estimators to learn from, under some probability that decides which experiences will be available to which estimator. Each estimator is trained against its own target network / target network head.

In evaluation, an ensemble voting policy is used to evaluate which action has been chosen by most heads. If there is no majority vote, an arbitrary choice is made between the actions chosen by the most heads. The action is then chosen and executed.




Fig. 1: The BDQN architecture proposed in [12].

BDQN attempts to achieve a measure of deep exploration by following the policy of one of the estimators for some number of steps. For this to be effective, the agent must guarantee that in areas of uncertainty (under-explored areas of the state action space), the different estimators will have different estimations. However, in BDQN this uncertainty, or variety in the Q-value estimations, is only based on the observed data [11]. This can be problematic, because in environments where rewards are very scarce, the agent may learn to believe that there is no reward, and lose all uncertainty, rather than direct its exploration to remote, unexplored areas of the state action space in the hope that they may contain rewards. Such ‘prior’ drive for exploration, that is independent from the data, is proposed in [11] in the form of randomized prior functions.

Randomized prior functions While usable with a regular DQN agent, the randomized prior functions algorithm is designed to be combined with the BDQN model. To achieve independent uncertainty, the randomized prior functions model consists of one additional neural network for each Q-value estimator, or one shared neural network with one head for each estimator. This additional network or head p is combined with the original estimator f to form the final output Q , through a scaling factor β : $Q = f + \beta p$ [11]. Q is then used in the learning process to minimize the training loss. This results in uncertainty that

is independent from the data: the Q value estimation always includes a neural network initialized with random parameters. No matter the uniformity of experiences the agent encounters (for example, only similar, negative rewards), it will still consider some additional prior 'assumption', in the form of the prior function, in regard to previously un-encountered states. As a result, each estimator will always approximate the Q value of as yet un-encountered states differently. This allows the agent to better direct its exploration, by guaranteeing diversity between the estimations of the different bootstrap heads for previously un-encountered states.

3.3 The model

The modeling of the traffic light control problem as an RL problem used in this paper is follows work done in [13]. The problem is modelled as an MDP $M = (S, A, P, R)$, where S is the state space, A is the action space, P the transition function and R the reward function. The open source traffic simulator SUMO [9] is used to generate the environment.

State space, action space & transition function The state provided to the agent is represented as a set of stacked frames of size $x \in \mathbb{Z}^+$. Each frame is a matrix containing current locations of vehicles in the agent's observation space, and the current traffic light configuration. The observation space of the agent is a square centered at the intersection controlled. Each location of a vehicle is marked with a 1, and empty locations with a 0. The traffic lights configuration is presented in the matrix as numbers between 0 and 1 chosen arbitrarily. Stacking several frames allows the agent to extrapolate vehicle speed from the state representation. $x = 4$ was used in the experiments, to achieve a balance between too much information (complicating the learning process), and too little information (hindering the capacity of the agent to learn effective policies).

The actions available to the agent at each time step are one of two traffic light configurations, representing which lanes receive a green light. The transition function is defined by SUMO.

Reward function The reward function used is a modified version of the reward function developed in [13]. At each time-step t , the agent receives a reward r_t , computed by iterating over all vehicles currently in the agent's observation space, and summing different penalties:

$$r_t = -1.5c - 0.2 \sum_{i=1}^N e_i - 0.3 \sum_{i=1}^N d_i - 0.3 \sum_{i=1}^N w_i$$

This, where i represents the vehicle index. c is a penalty for switching the light configuration, to prevent flickering. e_i is a penalty for sharp decelerations, to penalize emergency stops. d_i is a penalty for the 'delay' of a vehicle, defined as $1 - \frac{\text{vehicle speed}}{\text{allowed speed}}$. Finally, w_i is a waiting penalty, defined as 0.5 for the first step of a car standing still, and 1 for any consecutive step.

The modification included removal of a term that punishes teleportation of vehicles (used in SUMO to mark traffic collisions) due to implementation challenges. Additionally, the coefficient of the term c was increased from 0.1 to 1.5, after observation that otherwise the penalty for light switching is barely noticeable with almost any number of cars.

4 Experimental Setup

This section first describes and motivates the methodology used to evaluate the exploration approaches in the setting of traffic light control. This is followed by a description of the three traffic scenarios used to evaluate the exploration approaches.

4.1 Comparison of different exploration methods

To evaluate the impact of exploration in the setting of traffic light control, this paper compares the performance of agents using ϵ -greedy, BDQN [12], randomized prior functions [11] and a combination of the above for exploration, in three different scenarios. The performance is evaluated and averaged over multiple repetitions.

The compared agents As all three exploration approaches investigated are designed to be combined, the performance of the following six agents is compared: a regular DQN agent and regular DQN agent with a randomized prior function, both employing ϵ -greedy; Two BDQN agents with increasing bootstrap size: 4 & 10 bootstrap samples (or neural network 'heads'); Last, two similar BDQN agents, combining randomized prior functions in their bootstrap mechanism. The number of bootstrap samples has been chosen based on a relation between computational complexity (the larger the bootstrap, the larger the complexity), and gain from the bootstrap (the larger the bootstrap, the better the average performance). As shown empirically in [12], the relative gain from sizes larger than 10 becomes insignificant very quickly. The experimentation with different combinations of those techniques allows to evaluate, in essence, even more exploration techniques, and is the reason it is done in this paper.

The parameters of the agents investigated are not tuned for the specific setting of traffic light control. For the purpose of full reproducibility, a full list of the experiment parameters and agents' hyper-parameters used is available with the code base used in the experiments.

The evaluation The evaluation is done as follows: an experiment is done, for each agent in each scenario and each intersection considered. Each experiment consists of N learning episodes. Every *evaluation_frequency* learning episodes, an evaluation phase is ran. In order to reduce sensitivity to stochastic noise from the random nature of the traffic used in the experiments, in the evaluation phase the agent's policy is evaluated over *number_evaluations* evaluation episodes, with $\epsilon = 0$. The average episodic reward is then used for evaluation. The exact parameters *evaluation_frequency*, *number_evaluations* used in each experiment are detailed in section 5. To further reduce the impact of stochasticity, the entire experiment is repeated X times for each agent and the results averaged. Finally, the performances of the different agents are plotted against each other, in the form of their averaged evaluations' rewards. The results are presented in section 5.

The different parameters mentioned above were chosen in the following way: N was chosen from experimentation, as the range within which the agents' learning starts to plateau, in order to present the differences between the evaluations of the agents' in the clearest way. The *evaluation_frequency* used for each experiment set is chosen to achieve balance between the number of episodes each experiment is ran for, and the number of total evaluation episodes in the experiment. The *number_evaluations* parameter has been chosen as a balance between the total number of evaluation phases and the total number of episodes

in each experiment. The larger the number of learning episodes, and lower the evaluation frequency, the larger the *number_evaluations* parameter. To balance computational costs and time constraints with reliability of results, each experiment was chosen to be averaged over $X = 25$ repetitions.

For the purpose of full reproducibility, all random generators used are fully seeded, and the seeds logged. Each experiment is initiated with a different random seed, to guarantee random initialization of the agents' neural networks' weights. The environment's traffic generator however is seeded with the same set of seeds for all experiments. This is done to guarantee that while all agents experimented with are different, they are tested against the same traffic simulations.

4.2 The traffic scenarios

The impact of the exploration approaches investigated in this paper is evaluated using the traffic simulator SUMO [9], in three different traffic scenarios. In the first scenario, one set of experiments is done. In the second and third scenarios, two separate experiment sets are done, evaluating the agent against traffic of slower and faster average speed.

Scenario 1: The grid The first scenario is a basic grid-like road network, with one intersection in the center, and four roads going one in each direction from the intersection: north, east, south and west. A visualization of the grid scenario is presented in figure 2 a. The first scenario is meant to capture a simple, independent intersection profile, that does not consider or experience the behavior of other neighboring intersections. The traffic in this scenario is generated randomly, based on a set number of vehicles over a set spawning time.

Scenario 2: Simulation of real traffic in Manhattan, New York The second scenario, visualized in figure 2 b, is based on a section of the road map of Manhattan, New York, and is a more complex network containing several interconnected intersections. The specific road map used in our experiments is a 700 m^2 section centered around the corner of Waring and Woodhull Avenues. The map has been imported using SUMO's web-wizard. This scenario means to evaluate the effect of efficient exploration in a more complex traffic profile, where the agent may observe and consider the behavior of neighboring intersections. Manhattan enjoys a grid road-map design, that for the purpose of this work serves as both realistic and practical to use.

The red squares marked in figure 2 point to the two intersections given to the agent to control, as two separate experiments sets in this scenario. These intersections were chosen due to their encapsulation of different traffic profiles. The top intersection, Waring-Woodhull, presents a gentler form of traffic with significantly lower vehicle speed average. The bottom intersection, Eastchester-Waring, experiences much higher average vehicle speed, and more strongly resembles a central road. The throughput of both intersections is rather similar, with a slightly heavier load going through Eastchester-Waring. In every experiment, all intersections in the network except the one controlled by the agent are controlled by SUMO. To generate traffic for the Manhattan scenario, a random routes generator is used, based on real population distributions in the area, for the time of day 08:30 AM to 09:00 AM. The traffic data is imported using SUMO's web-wizard as well.

Scenario 3: Simulation of custom traffic in Manhattan, New York A third scenario is introduced in order to evaluate exploration under heavier traffic settings. This scenario uses the same map, and experiments with the same intersections illustrated in figure 2 b. However, in this scenario a random traffic generator is used for traffic generation, in order to introduce much heavier traffic loads than the ones generated to simulate real traffic. Again, two different sets of experiments are done in this scenario, one on each of the two marked intersections in figure 2 b. The difference of the traffic profiles between the two intersections is similar to the one in the second scenario: the top intersection enjoys slower average speed, and the bottom faster.

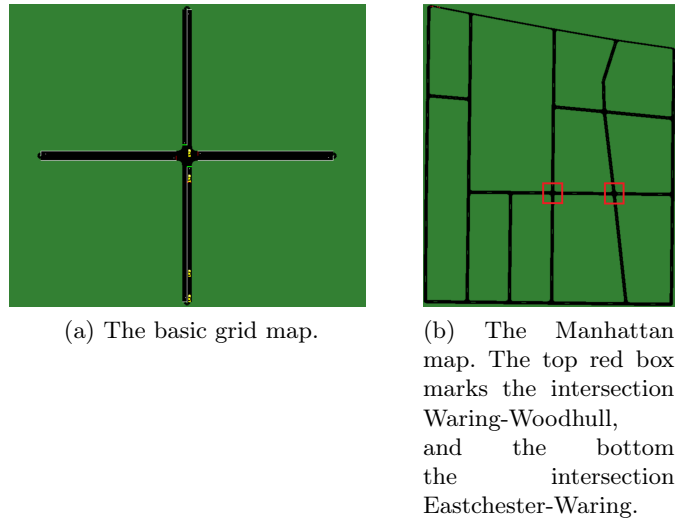


Fig. 2: The road maps used in the traffic scenarios.

Observation space specification An important difference between the scenarios is the observation space provided to the agent in each one. The observation space provided in the experiments done with the grid scenario and the Eastchester-Waring intersection in either scenario, was 50m^2 . This is done because in all three a larger observation space would be outside the bounds of the environment. Due to the structure of the scenario, it was possible to provide the agent with a larger observation space in the experiments done with the Waring-Woodhull intersection in either scenario. An observation space of 84m^2 was chosen, to balance complexity (the larger the observation, the more complex the learning) with observation distance. This is done to allow the agent access to more information, which (1) contains the adjacent intersections, enabling the agent to take their behavior into account, and (2) providing the agent with the ability to react to incoming traffic earlier, as a result of the larger observation space.

5 Results

This section presents and analyzes the results obtained in the experiments described in section 4, divided between the different scenarios investigated and intersections controlled.

For each set of experiments, the results presented are the episodic rewards attained in the evaluations, as described in section 4.1. This is presented alongside a plot of the 95% confidence interval of the mean, computed using the standard error of the mean (SEM) [2] of the different experiments for two sample agents. These agents are chosen independently for each experiment: the one that performed, on average, the best, and the one that performed the worst. This is done to illustrate how significant are the differences observed between the evaluations of the different agents in each experiment. Only two agents are presented in order to reduce clutter in the plot. A simple moving average (SMA) of window size 5 is applied to the data presented in order to smoothen the stochastic effect, to facilitate visual analysis of the results.

5.1 Scenario 1: The grid

The results of evaluating the performance of six different agents against the grid scenario can be found in figure 3. The agents’ policies are evaluated every five learning episodes, and averaged over three evaluation episodes. Additionally, the 95% confidence intervals of the means of the two sample agents are presented.

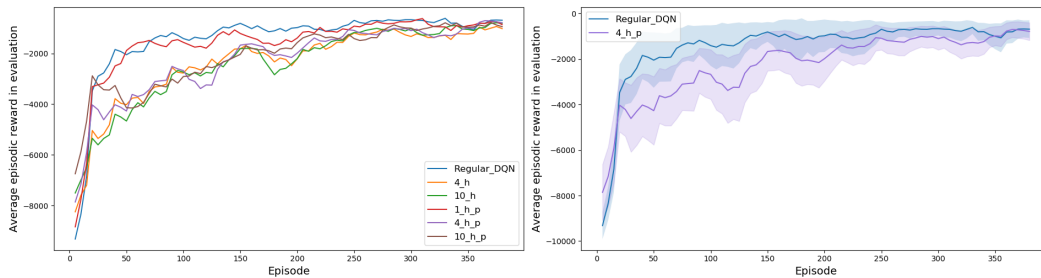


Fig. 3: Evaluating the agents against the grid scenario. The left figure presents the evaluations. The number describes the size of the bootstrap sample (number of heads), and the p whether a randomized prior has been incorporated. The right figure presents the same for two chosen agents, including the 95% confidence interval of the mean.

Figure 3 illustrates that while all agents learn, the agents that appear to have the sharpest learning rates are the regular DQN with and without prior function applied. However, as can be seen in the right plot in figure 3, there is some overlap in the confidence intervals of their means.

5.2 Scenario 2: Simulation of real traffic in Manhattan, New York

The results of evaluating the agents against the Manhattan scenario simulating real traffic can be found below, separately for each set of experiments, controlling each of the two intersections marked in figure 2 b. The agents’ policies are evaluated every training episode, over two evaluation episodes, and averaged.

Traffic of low average speed Figure 4 presents the results of experimenting control of the intersection Waring-Woodhull, the top of the two intersections in figure 2 b, capturing a traffic profile of slower average speed.

As observable, the difference between the evaluation scores in relation to the confidence interval of the means, is mostly negligible, with the exception that the regular DQN achieves inferior scores prior to episode 20. However, these scores are still well within the confidence interval of the means of the other agent’s, and thus cannot be considered significant. This behavior is attributed to the simplicity of the traffic profile, in relation to the amount of information accessible to the agent in this scenario. As mentioned in section 4.2, the observation space of the agents in this experiment is 84 m². While the scenario can be viewed as complex (several interconnected intersections, whose behaviors directly influence each other’s traffic), which can translate to the learning process being more difficulty, the volume of the traffic, including the average speed, is rather low, and thus the policy required is not complex.




Fig. 4: Evaluating the different agents against the Waring-Woodhull intersection, simulating real traffic. The left figure presents the evaluation of the different agents. The number describes the size of the bootstrap sample (number of heads), and the p whether a randomized prior has been incorporated. The right figure presents the same results for a sample of the agents, including the 95% confidence interval of the mean.

Traffic of high average speed Figure 5 presents the results of experiments done controlling the second intersection, Eastchester-Waring. Eastchester-Waring enjoys both traffic of higher average speed, as well as higher traffic loads. No significant difference in the performance of the different agents is observed. This is further illustrated with the confidence interval of the means presented in figure 5 and their overlap.




Fig. 5: Evaluating the different agents against the Eastchester-Waring intersection, simulating real traffic. The left figure presents the evaluation of the different agents. The number describes the size of the bootstrap sample (number of heads), and the p whether a randomized prior has been incorporated. The right figure presents the same results for a sample of the agents, including the 95% confidence interval of the mean.

5.3 Scenario 3: Simulation of custom traffic in Manhattan, New York

The results of evaluating the agents against the Manhattan scenario simulating random, heavy traffic can be found below. The agents' policies are evaluated every second training episodes, and averaged over two evaluation episodes.

Traffic of low average speed The results for the intersection Waring Woodhull are presented in figure 6. While the differences in the learning between most of the different agents appears negligible, all of them appear to outperform the regular DQN agent. However, as illustrated by the right plot in figure 6, the confidence intervals still have some overlap.




Fig. 6: Evaluating the different agents against the Waring-Woodhull intersection in the custom scenario. The left figure presents the evaluation of the different agents. The number describes the size of the bootstrap sample (number of heads), and the p whether a randomized prior has been incorporated. The right figure presents the same results for a sample of the agents, including the 95% confidence interval of the mean.

Traffic of high average speed The results of experimenting control over the intersection Eastchester Waring are presented in figure 7. No significant difference is observed between the evaluations of the different agents. This is illustrated strongly by the confidence intervals, and their overlap, presented in the right plot in figure 7.




Fig. 7: Evaluating the different agents against the Eastchester-Waring intersection in the custom scenario. The left figure presents the evaluation of the different agents. The number describes the size of the bootstrap sample (number of heads), and the p whether a randomized prior has been incorporated. The right figure presents the same results for a sample of the agents, including the 95% confidence interval of the mean.

5.4 Analysis

Many different factors can influence the results obtained by the different agents, ranging from the reward function, the size of the observation space and its resolution, the state abstraction, the complexity of the traffic scenario and sheer stochasticity. The type and complexity of the traffic scenario, along with the observation space, are the main parameters that will be considered in this analysis.

Here, the complexity of the scenario, while not defined exactly, considers the volume and variety in the traffic and the number of surrounding interconnected intersections. In particular, the "grid" scenario can be viewed as having lower complexity, as it has a medium traffic volume, of low speeds and no neighbouring intersections. In the experiments done in the scenario simulating real traffic, both intersections could be viewed as having higher complexity: "Waring Woodhull", which is surrounded by a number of neighbouring intersections, with lower traffic volume and a slower speed; and Eastchester-Waring, which is surrounded by a similar number of neighbouring intersections, with somewhat higher traffic volume with significantly higher average speed. Last, the scenario simulating custom traffic could be viewed as having higher complexity as well, as it has similar parameters to the simulation of real traffic, with the exception of significantly increased traffic loads. In addition, it is worth mentioning that in the experiments done with the Waring Woodhull intersection the observation space was larger than in any of the other experiments, as specified in section 4.2.

It appears that when the complexity of the traffic is low along with a smaller observation space (the grid scenario), the simpler agents converge much faster, and more stably, to a highly evaluated policy. However, when the agent is given sufficient information of sufficient complexity (the Waring Woodhull intersection in both scenarios), deep exploration approaches are able to outperform the ϵ -greedy approach, by achieving faster learning. This may imply that advanced exploration approaches may play a significantly more critical role in much more sophisticated scenarios, such as an intersection balancing many lanes, experiencing many different traffic profiles, and provided a large observation space. The differences observed were generally small and within a 95% confidence interval of the means however, and as such, the strength of this implication is limited.

We note, though, that for the Eastchester-Waring intersection no improvements due to more sophisticated exploration were observed, even though the traffic patterns are more complex here than in the simple grid. The reason for this could be that as a result of the faster traffic and the smaller observation space, along with sufficient complexity in the road map and variety in the loads and directions of the traffic, the different costs and gains from the different exploration approaches balanced out.

6 Discussion

The results obtained suggest a link between the complexity of the scenario, the information accessible to the agent and the gain from deep exploration, as discussed in section 5.4. However, the advanced exploration approaches investigated are expected to have significant gain especially when utilized in environments with scarcity of significant positive rewards, and abundance of smaller negative rewards [11],[12]. As a consequence, the reward function used can almost directly dictate the gain from different exploration strategies. The reward function used was not designed for any specific exploration approach. It is therefore plausible that under a tailored reward function the gain from deep exploration would have been much

higher. This can be especially relevant in the case of reward functions that are known for promoting good policies, while hindering learning.

7 Conclusions and Future Work

This paper investigated the value of deep exploration in the setting of traffic light control, by comparing agents using different exploration approaches in three different traffic scenarios of rising complexity. Specifically, the state of the art approaches Bootstrapped DQN [12] and randomized prior functions [11] were compared to a baseline ϵ -greedy approach. This, to facilitate better deep RL in traffic light control, by identifying the value of evolved exploration approaches in this setting, such as higher sample efficiency or higher final policy score.

The results presented in this paper suggest a link between the complexity of the traffic scenario, the size of the observation space of the agent, and the gain from efficient exploration, achieved with Bootstrapped DQN and randomized prior functions, under a specific parameters configuration. Specifically, the more complex the scenario and the larger the observation space, the more significant the gain observed from efficient exploration.

The results presented leave the following open questions, however. What is the exact relation between the gain from efficient exploration and the complexity of the scenario? How sensitive is this relation to specific parameter configurations, and specifically under parameters optimized for the specific setting? Does this conclusion apply for other exploration approaches? These questions are left for future work.

Acknowledgements

We would like to thank the students in the research-project group, Pepijn Tersmette, Cian Jansen, Emanuel Kuhn & Chris van der Werf, for their ready assistance and availability for discussion, along with their contributions to the creation of the traffic scenarios and the implementation of the reward function used. Additionally, we would like to thank the TUDelft ILDM research group for their helpful comments and support. An additional thanks goes to the TUDelft for granting us access to the INSY cluster, which was a crucial computational resource used in this research.

This project had received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 758824 —INFLUENCE).



References

1. Bakker, B., Whiteson, S., Kester, L., Groen, F.C.: Traffic light control by multiagent reinforcement learning systems. In: Interactive Collaborative Information Systems, pp. 475–510. Springer (2010)
2. Barde, M.P., Barde, P.J.: What to use to express the variability of data: Standard deviation or standard error of mean? Perspectives in clinical research **3**(3), 113 (2012)

3. Ciosek, K., Vuong, Q., Loftin, R., Hofmann, K.: Better exploration with optimistic actor critic. In: *Advances in Neural Information Processing Systems*. pp. 1785–1796 (2019)
4. Coşkun, M., Baggag, A., Chawla, S.: Deep reinforcement learning for traffic light optimization. In: *2018 IEEE International Conference on Data Mining Workshops (ICDMW)*. pp. 564–571. IEEE (2018)
5. Efron, B., Tibshirani, R.J.: *An introduction to the bootstrap*. CRC press (1994)
6. Fortunato, M., Azar, M.G., Piot, B., Menick, J., Osband, I., Graves, A., Mnih, V., Munos, R., Hassabis, D., Pietquin, O., et al.: Noisy networks for exploration. *arXiv preprint arXiv:1706.10295* (2017)
7. François-Lavet, V., Henderson, P., Islam, R., Bellemare, M.G., Pineau, J.: An introduction to deep reinforcement learning. *arXiv preprint arXiv:1811.12560* (2018)
8. Hansen, J.: *bootstrap_dqn*. GitHub repository (2019)
9. Lopez, P.A., Behrisch, M., Bieker-Walz, L., Erdmann, J., Flötteröd, Y.P., Hilbrich, R., Lücken, L., Rummel, J., Wagner, P., Wießner, E.: Microscopic traffic simulation using sumo. In: *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*. pp. 2575–2582. IEEE (2018)
10. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., et al.: Human-level control through deep reinforcement learning. *Nature* **518**(7540), 529–533 (2015)
11. Osband, I., Aslanides, J., Cassirer, A.: Randomized prior functions for deep reinforcement learning. In: *Advances in Neural Information Processing Systems*. pp. 8617–8629 (2018)
12. Osband, I., Blundell, C., Pritzel, A., Van Roy, B.: Deep exploration via bootstrapped dqn. In: *Advances in neural information processing systems*. pp. 4026–4034 (2016)
13. Van der Pol, E., Oliehoek, F.A.: Coordinated deep reinforcement learners for traffic light control. *Proceedings of Learning, Inference and Control of Multi-Agent Systems (at NIPS 2016)* (2016)
14. Rezzai, M., Dachry, W., Mouataouakkil, F., Medromi, H.: Reinforcement learning for traffic control system: Study of exploration methods using q-learning. *International Research Journal of Engineering and Technology* **4**(10), 1838–1848 (2017)
15. TO, H.R., Barker, M.M.: White paper european transport policy for 2010: time to decide (2001)
16. Van Hasselt, H., Guez, A., Silver, D.: Deep reinforcement learning with double q-learning. In: *Thirtieth AAAI conference on artificial intelligence* (2016)
17. Watkins, C.J., Dayan, P.: Q-learning. *Machine learning* **8**(3-4), 279–292 (1992)
18. Wei, H., Zheng, G., Yao, H., Li, Z.: Intellilight: A reinforcement learning approach for intelligent traffic light control. In: *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. pp. 2496–2505 (2018)