

Fine-grained acceleration control for autonomous intersection management using deep reinforcement learning

Hamid Mirzaei
Dept. of Computer Science
University of California, Irvine
mirzaei@uci.edu

Tony Givargis
Dept. of Computer Science
University of California, Irvine
givargis@uci.edu

Abstract—Recent advances in combining deep learning and Reinforcement Learning have shown a promising path for designing new control agents that can learn optimal policies for challenging control tasks. These new methods address the main limitations of conventional Reinforcement Learning methods such as customized feature engineering and small action/state space dimension requirements. In this paper, we leverage one of the state-of-the-art Reinforcement Learning methods, known as Trust Region Policy Optimization, to tackle intersection management for autonomous vehicles. We show that using this method, we can perform fine-grained acceleration control of autonomous vehicles in a grid street plan to achieve a global design objective.

I. INTRODUCTION

Previous works on autonomous intersection management (AIM) in urban areas have mostly focused on intersection arbitration as a shared resource among a large number of autonomous vehicles. In these works [1][2], high-level control of the vehicles is implemented such that the vehicles are self-contained agents that only communicate with the intersection management agent to reserve space-time slots in the intersection. This means that low-level vehicle navigation which involves acceleration and speed control is performed by each individual vehicle independent of other vehicles and intersection agents. This approach is appropriate for minor arterial roads where a large number of vehicles utilize the main roads at similar speeds while the adjacent intersections are far away.

In scenarios involving local roads, where the majority of the intersections are managed by stop signs, the flow of traffic is more efficiently managed using a fine-grained vehicle control methodology. For example, when two vehicles are crossing the intersection of two different roads at the same time, one vehicle can decelerate slightly to avoid collision with the other one or it can take another path to avoid confronting the other vehicle completely. Therefore, the nature of the AIM problem is a combination of route planning and real-time acceleration control of the vehicles. In this paper, we propose a novel AIM formulation which is the combination of route planning and fine-grained acceleration control. The main objective of the control task is to minimize travel time of the vehicles while avoiding collisions between them and other obstacles. In this

context, since the movement of a vehicle is dependent on the other vehicles in the same vicinity, the motion data of all vehicles is needed in order to solve the AIM problem.

To explain the proposed AIM scheme, let us define a “zone” as a rectangular area consisting of a number of intersections and segments of local roads. An agent for each zone collects the motion data and generates the acceleration commands for all autonomous vehicles within the zone’s boundary. All the data collection and control command generation should be done in real-time. This centralized approach cannot be scaled to a whole city, regardless of the algorithm used, due to the large number of vehicles moving in a city which requires enormous computational load and leads to other infeasible requirements such as low-latency communication infrastructure. Fortunately, the spatial independence (i.e., the fact that navigation of the vehicles in one zone is independent of the vehicles in another zone that is far enough away) makes AIM an inherently local problem. Therefore, we can assign an agent for each local zone in a cellular scheme.

The cellular solution nevertheless leads to other difficulties that should be considered for a successful design of the AIM system. One issue is the dynamic nature of the transportation problem. Vehicles can enter or leave a zone controlled by an agent or they might change their planned destinations from time to time. To cope with these issues, the receding horizon control method can be employed where the agent repeatedly recalculates the acceleration command over a moving time horizon to take into account the mentioned changes. Additionally, two vehicles that are moving toward the same point on the boundary of two adjacent zones simultaneously might collide because the presence of each vehicle is not considered by the agent of the adjacent zone. This problem can be solved by adequate overlap between adjacent zones. Furthermore, any planned trip for a vehicle typically crosses multiple zones. Hence, a higher level planning problem should be solved first that determines the entry and exit locations of a vehicle in a zone.

In this paper we focus on the subproblem of acceleration control of the vehicles moving in a zone to minimize the total travel time. We use a deep reinforcement learning (RL) approach to tackle the fine-grained acceleration control

problem since conventional control methods are not applicable because of the non-convex collision avoidance constraints [3]. Furthermore, if we want to incorporate more elements into the problem, such as obstacles or reward/penalty terms for gas usage, passenger comfort, etc., the explicit modeling becomes intractable and an optimal control law derivation will be computationally unattainable.

RL methods can address the above mentioned limitations caused by the explicit modeling requirement and conventional control method limitations. The main advantage of RL is that most of the RL algorithms are “model-free” or at most need a simulation model of the physical system which is easier to develop than an explicit model. Moreover, the agent can learn optimal policies just by interacting with the environment or executing the simulation model. However, conventional RL techniques are only applicable in small-scale problem settings and require careful design of approximation functions. Emerging Deep RL methods [4] that leverage the deep neural networks to automatically extract features seem like promising solutions to shortcomings of the classical RL methods.

The main contributions of this paper are: (1) Definition and formulation of the AIM problem for local road settings where vehicles are coordinated by fine-grained acceleration commands. (2) Employing TRPO proposed in [5] to solve the formulated AIM problem. (3) Incorporating collision avoidance constraint in the definition of RL environment as a safety mechanism.

II. RELATED WORK

Advances in autonomous vehicles in recent years have revealed a portrait of a near future in which all vehicles will be driven by artificially intelligent agents. This emerging technology calls for an intelligent transportation system by redesigning the current transportation system which is intended to be used by human drivers. One of the interesting topics that arises in intelligent transportation systems is AIM. Dresner et al. have proposed a multi-agent AIM system in which vehicles communicate with intersection management agents to reserve a dedicated spatio-temporal trajectory at the intersection [2].

In [6], authors have proposed a self-organizing control framework in which a cooperative multi-agent control scheme is employed in addition to each vehicle’s autonomy. The authors have proposed a priority-level system to determine the right-of-way through intersections based on vehicles’ characteristics or intersection constraints.

Zohdy et al. presented an approach in which the Cooperative Adaptive Cruise Control (CACC) systems are leveraged to minimize delays and prevent clashes [7]. In this approach, the intersection controller communicates with the vehicles to recommend the optimal speed profile based on the vehicle’s characteristics, motion data, weather conditions and intersection properties. Additionally, an optimization problem is solved to minimize the total difference of actual arrival times at the intersection and the optimum times subject to conflict-free temporal constraints.

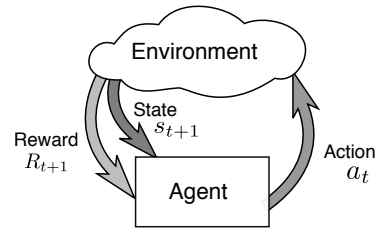


Fig. 1: Agent-Environment interaction model in RL

A decentralized optimal control formulation is proposed in [8] in which the acceleration/deceleration of the vehicles are minimized subject to collision avoidance constraints.

Makarem et al. introduced the notion of fluent coordination where smoother trajectories of the vehicles are achieved through a navigation function to coordinate the autonomous vehicles along predefined paths with expected arrival time at intersections to avoid collisions.

In all the aforementioned works, the AIM problem is formulated for only one intersection and no global minimum travel time objective is considered directly. Hausknecht et al. extended the approach proposed in [2] to multi-intersection settings via dynamic traffic assignment and dynamic lane reversal [1]. Their problem formulation is based on intersection arbitration which is well suited to main roads with a heavy load of traffic.

In this paper, for the first time, we introduce fine-grained acceleration control for AIM. In contrast to previous works, Our proposed AIM scheme is applicable to local road intersections. We also propose an RL-based solution using Trust Region Policy Optimization to tackle the defined AIM problem.

III. REINFORCEMENT LEARNING

In this section, we briefly review RL and introduce the notations used in the rest of the paper. In Fig. 1, the agent-environment model of RL is shown. The “agent” interacts with the “environment” by applying “actions” that influence the environment state at the future time steps and observes the state and “reward” in the next time step resulting from the action taken. The “return” is defined as the sum of all the rewards from the current step to the end of current “episode”:

$$G_t = \sum_{i=t}^T r_i \quad (1)$$

where r_i are future rewards and T is the total number of steps in the episode. An “episode” is defined as a sequence of agent-environment interactions. In the last step of an episode the control task is “finished.” Episode termination is defined specifically for the control task of the application.

For example, in the cart-pole balancing task [9], the agent is the controller, the environment is the cart-pole physical system, the action is the force command applied to the cart, and the reward can be defined as $r = 1$ as long as the pole is nearly

in an upright position and a large negative number when the pole falls. The system states are cart position, cart speed, pole angle and pole angular speed. The agent task is to maximize the return G_t , which is equivalent to prevent pole from falling for the longest possible time duration.

In RL, a control policy is defined as a mapping of the system state space to the actions:

$$a = \pi(s) \quad (2)$$

where a is the action, s is the state and π is the policy. An optimal policy is one that maximizes the return for all the states, i.e.:

$$v_{\pi^*}(s) \geq v_{\pi}(s), \quad \text{for all } s, \pi \quad (3)$$

where v is the return function defined as the return achievable from state s by following policy π . Equation (3) means that the expected return under optimal policy π^* is equal to or greater than any other policy for all the system states.

The concepts mentioned above to introduce RL are all applicable to deterministic cases, but generally we should be able to deal with inherent system uncertainty, measurement noise, or both. Therefore, we model the system as a Markov Decision Process (MDP) assuming that the environment has the *Markov property* [9]. However, contrary to most of the control design methods, many RL algorithms do not require the models to be known beforehand. The elimination of the requirement to model the system under control is a major strength of RL.

A system has the Markov property if at a certain time instant, t , the system history can be captured in a set of *state variables*. Therefore, the next state of the system has a distribution which is only conditioned on the current state and the taken action at the current time, i.e.:

$$s_{t+1} \sim P(s_{t+1}|s_t, a_t) \quad (4)$$

The Markov property holds for many cyber-physical system application domains and therefore MDP and RL can be applied as the control algorithm. We can also define the stochastic policy which is a generalized version of (2) as a probability distribution of actions conditioned on the current state, i.e.:

$$a_t \sim \pi(a_t|s_t) \quad (5)$$

The expected return function which is the expected value of ‘return’ defined in (1) can be written as:

$$v_{\pi}(s_t) = \mathbb{E}_{a_{\tau} \sim \pi, \tau \geq t} \left[\sum_{i=0}^{\infty} \gamma^i r_{t+i} \right] \quad (6)$$

This equation is defined for infinite episodes and the constant $0 < \gamma < 1$ is introduced to ensure that the defined expected return is always a finite value, assuming the returns are bounded.

Another important concept in RL is the action-value function, $Q_{\pi}(s, a)$ defined as the expected return (value) if

action a_t is taken at time t under policy π :

$$Q_{\pi}(s_t, a_t) = \mathbb{E}_{a_{\tau} \sim \pi, \tau > t} \left[\sum_{i=0}^{\infty} \gamma^i r_{t+i} \right] \quad (7)$$

There are two main categories of methods to find the optimal policy. In the first category, $Q_{\pi}(s, a)$ is parameterized as $Q_{\pi}^{\theta}(s, a)$ and the optimal action-value parameter vector θ is estimated in an iterative process. The optimal policy can be defined implicitly from $Q_{\pi}(s, a)$. For example, the greedy policy is the one that maximizes $Q_{\pi}(s, a)$ in each step:

$$a_t = \arg \max_a \{Q_{\pi}(s_t, a)\} \quad (8)$$

In the second category, which is called policy optimization and has been successfully applied to large-scale and continuous control systems [4], the policy is parameterized directly as $\pi^{\theta}(a_t|s_t)$ and the parameter vector of the optimal policy θ is estimated. The Trust Region Policy Method (TRPO) [5] is an example of the second category of methods that guarantees monotonic policy improvement and is designed to be scalable to large-scale settings. In each iteration of TRPO, a number of MDP trajectories are simulated (or actually experienced by the agent) and θ is updated to improve the policy. A high level description of TRPO is shown in Algorithm 1.

IV. PROBLEM STATEMENT

There is a set of vehicles in a grid street plan area consisting of a certain number of intersections. For simplicity, we assume that all the initial vehicle positions and the desired destinations are located at the intersections. There is a control agent for the entire area. The agent’s task is to calculate the acceleration command for the vehicles in real-time (see Fig. 2). We assume that there are no still or moving obstacles other than vehicles’ or street boundaries.

The input to the agent is the real-time state of the vehicles which consists of their positions and speeds. We are assuming that vehicles are point masses and their angular dynamics are ignored. However, to take the collision avoidance in the problem formulation, we define a safe radius for each vehicle and no objects (vehicles or street boundaries) should be closer than the safe radius to the vehicle.

Algorithm 1: High-Level description of Trust Region Optimization

Data: \mathcal{S} ▷ Actual system or Simulation model
 π^{θ} ▷ Parameterized Policy
Result: θ^* ▷ Optimal parameters

```

1 repeat
2   Use  $\mathcal{S}$  to generate trajectories of the system using
   current  $\pi^{\theta}$ ;
3   Perform one iteration of policy optimization using
   Monte Carlo method to get  $\theta_{new}$ ;
4    $\theta \leftarrow \theta_{new}$ 
5 until no more improvements;
6 return  $\theta$ 

```

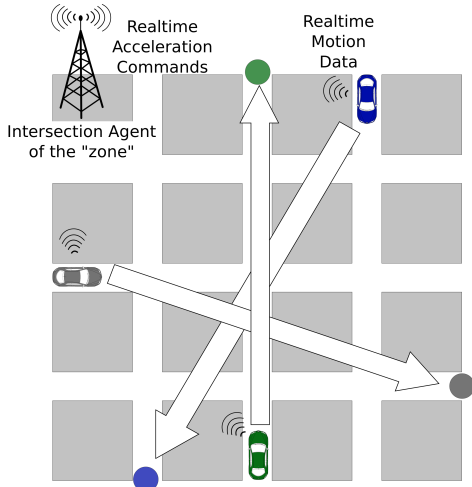


Fig. 2: Intersection Management Problem. The goal of the problem is to navigate the vehicles from the sources to destinations in minimum time with no collisions.

The objective is to drive all the vehicles to their respective destinations in a way that the total travel time is minimized. Furthermore, no collision should occur between any two vehicles or a vehicle and the street boundaries.

To minimize the total travel time, a positive reward is assigned to the terminal state in which all the vehicles approximately reach the destinations within some tolerance. A discount factor γ strictly less than one is used. Therefore, the agent should try to reach the terminal state as fast as possible to maximize the discounted return. However, by using only this reward, too many random walk trajectories are needed to discover the terminal state. Therefore, a negative reward is defined for each state, proportional to the total distance of the vehicles to their destinations as a hint of how far the terminal state is. This negative reward is not in contradiction with the main goal which is to minimize total travel time.

To avoid collisions, two different approaches can be considered: we can add large negative rewards for the collision states or we can incorporate a collision avoidance mechanism into the environment model. Our experiments show that the first approach makes the agent too conservative about moving the vehicles to minimize the probability of collisions. This might lead to extremely slow learning which makes it computationally infeasible. Furthermore, collisions are inevitable even with large negative rewards which limits the effectiveness of learned policies in practice.

For the above mentioned reasons, the second approach is employed, i.e. the safety mechanism that is used in practice is included in the environment definition. The safety mechanism is activated whenever two vehicles are too close to each other or a vehicle is too close to the street boundary. In these cases, the vehicle built-in collision avoidance system will control the vehicle's acceleration and the acceleration commands from the RL agent are ignored as long as the distance is near the allowed safe radius of the vehicle. In the agent learning process these cases are simulated in a way that the vehicles

come to a full stop when they are closer than the safe radius to another vehicle or boundary. By applying this heuristic in the simulation model, the agent should avoid any "near collision" situations explained above because the deceleration and acceleration cycles take a lot of time and will decrease the expected return.

Based on the problem statement explained above, we can describe the RL formulation in the rest of the subsection. The state is defined as the following vector:

$$\mathbf{s}_t = \left(x_t^1, y_t^1, v_{x_t}^1, v_{y_t}^1, \dots, x_t^n, y_t^n, v_{x_t}^n, v_{y_t}^n \right)^\top \quad (9)$$

where (x_t^i, y_t^i) and $(v_{x_t}^i, v_{y_t}^i)$ are the position and speed of vehicle i at time t . The action vector is defined as:

$$\mathbf{a}_t = \left(a_{x_t}^1, a_{y_t}^1, \dots, a_{x_t}^n, a_{y_t}^n \right)^\top \quad (10)$$

where $(a_{x_t}^i, a_{y_t}^i)$ is the acceleration command of vehicle i at time t . The reward function is defined as:

$$r(s) = \begin{cases} 1 & \text{if } \|(x^i - d_x^i, y^i - d_y^i)^\top\| < \eta \quad (1 \leq i \leq n) \\ -\alpha \sum_{i=1}^n \|(x^i - d_x^i, y^i - d_y^i)^\top\| & \text{otherwise} \end{cases} \quad (11)$$

where (d_x^i, d_y^i) is the destination coordinates of vehicle i , η is the distance tolerance and α is a positive constant.

Assuming no collision occurs, the state transition equations for the environment are defined as follows:

$$\begin{aligned} x_{t+1}^i &= \text{sat}_{\underline{x}, \bar{x}}(x_t^i + h v_{x_t}^i) \\ y_{t+1}^i &= \text{sat}_{\underline{y}, \bar{y}}(y_t^i + h v_{y_t}^i) \\ v_{x_{t+1}}^i &= \text{sat}_{\underline{v}_m, \bar{v}_m}(v_{x_t}^i + h a_{x_t}^i) \\ v_{y_{t+1}}^i &= \text{sat}_{\underline{v}_m, \bar{v}_m}(v_{y_t}^i + h a_{y_t}^i) \end{aligned} \quad (12)$$

where h is the sampling time, $(\underline{x}, \bar{x}, \underline{y}, \bar{y})$ defines area limits, v_m is the maximum speed and $\text{sat}_{\underline{w}, \bar{w}}(\cdot)$ is the saturation function defined as:

$$\text{sat}_{\underline{w}, \bar{w}}(x) = \begin{cases} \underline{w} & x \leq \underline{w} \\ \bar{w} & x \geq \bar{w} \\ x & \text{otherwise.} \end{cases} \quad (13)$$

To model the collisions, we should check certain conditions and set the speed to zero. A more detailed description of collision modeling is presented in Algorithm 2.

A. Solving the AIM problem using TRPO

The simulation model can be implemented based on the RL formulation described in Section IV. To use TRPO, we need a parameterized stochastic policy, $\pi^\theta(a_t|s_t)$, in addition to the simulation model. The policy should specify the probability distribution for each element of the action vector defined in (10) as a function of the current state s_t .

We have used the sequential deep neural network (DNN) policy representation as described in [5]. The input layer receives the state containing the position and speed of the vehicles (defined in (9)). There are a number of hidden layers, each followed by tanh activation functions [10]. Finally, the

output layer generates the mean of a gaussian distribution for each element of the action vector.

To execute the optimal policy learned by TRPO in each sampling time, the agent calculates the forward-pass of DNN using the current state. Next, assuming that all the action elements have the same variance, the agent samples from the action gaussian distributions and applies the sampled actions to the environment as the vehicle acceleration commands.

V. EVALUATION

A. Baseline Method

To the best of our knowledge there is no other solution proposed for the fine-grained acceleration AIM problem introduced in this paper. Therefore, we use conventional optimization methods to study how close the proposed solution is to the optimal solution. Furthermore, we will see that the conventional optimization is able to solve the AIM problem only for very small-sized problems. This confirms that the proposed RL-based solution is a promising alternative to the conventional methods.

Theoretically, the best solution to the problem defined in Section IV can be obtained if we reformulate it as a conventional optimization problem. The following equations and inequalities describe the AIM optimization problem:

$$\mathbf{a}_t^* = \arg \max_{\mathbf{a}_t} \sum_{t=0}^{T-1} \sum_{i=1}^n \|(x_t^i - d_x^i, y_t^i - d_y^i)^\top\| \quad (14)$$

$$\text{s. t. } \underline{x} \leq x_t^i \leq \bar{x} \quad (1 \leq i \leq n) \quad (15)$$

$$\underline{y} \leq y_t^i \leq \bar{y} \quad (1 \leq i \leq n) \quad (16)$$

$$\underline{v}_m \leq v_{x_t}^i \leq \overline{v}_m \quad (1 \leq i \leq n) \quad (17)$$

$$\underline{v}_m \leq v_{y_t}^i \leq \overline{v}_m \quad (1 \leq i \leq n) \quad (18)$$

Algorithm 2: State Transition Function

- | | |
|--------------------------|-------------------------|
| Data: s_t | ▷ State at time t |
| a_t | ▷ Action at time t |
| Result: s_{t+1} | ▷ State at time $t + 1$ |
- 1 $a_{x_t}^i \leftarrow \text{sat}_{\underline{a}_m, \overline{a}_m}(a_{x_t}^i)$;
 - 2 $a_{y_t}^i \leftarrow \text{sat}_{\underline{a}_m, \overline{a}_m}(a_{y_t}^i)$;
 - 3 $s_{t+1} \leftarrow$ updated state using (12);
 - 4 $v_{c1} \leftarrow$ find all the vehicles colliding with street boundaries;
 - 5 speed elements of v_{c1} in $s_{t+1} \leftarrow 0$;
 - 6 location elements of v_{c1} in $s_{t+1} \leftarrow$ closest point on the street boundary with the margin of ϵ ;
 - 7 $v_{c2} \leftarrow$ find all the vehicles colliding with some other vehicle;
 - 8 speed elements of v_{c2} in $s_{t+1} \leftarrow 0$;
 - 9 location elements of v_{c2} in $s_{t+1} \leftarrow$ pushed back location with the distance of $2 \times$ safe radius to the collided vehicle;
 - 10 **return** s_{t+1} ;
-

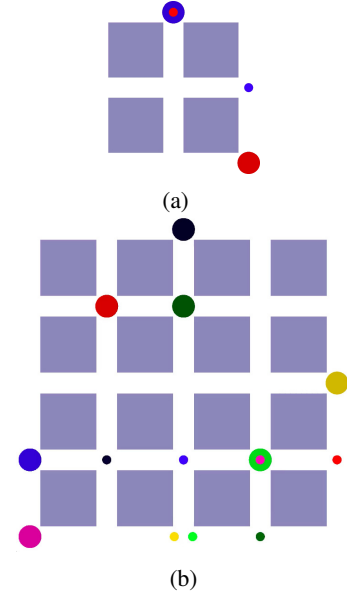


Fig. 3: Initial setup of each episode. Small circles are the sources and big circles are the destinations. (a) small example (b) large example

$$\underline{a}_m \leq a_{x_t}^i \leq \overline{a}_m \quad (1 \leq i \leq n) \quad (19)$$

$$\underline{a}_m \leq a_{y_t}^i \leq \overline{a}_m \quad (1 \leq i \leq n) \quad (20)$$

$$-\lfloor N/2 \rfloor \leq r_t^i \leq \lfloor N/2 \rfloor \quad (1 \leq i \leq n, r_t^i \in \mathbb{Z}) \quad (21)$$

$$-\lfloor M/2 \rfloor \leq c_t^i \leq \lfloor M/2 \rfloor \quad (1 \leq i \leq n, c_t^i \in \mathbb{Z}) \quad (22)$$

$$x_0^i = s_x^i, y_0^i = s_y^i \quad (1 \leq i \leq n) \quad (23)$$

$$x_{T-1}^i = d_x^i, y_{T-1}^i = d_y^i \quad (1 \leq i \leq n) \quad (24)$$

$$v_{x_0}^i = 0, v_{y_0}^i = 0 \quad (1 \leq i \leq n) \quad (25)$$

$$x_{t+1}^i = x_t^i + v_{x_t}^i \cdot h \quad (1 \leq i \leq n) \quad (26)$$

$$y_{t+1}^i = y_t^i + v_{y_t}^i \cdot h \quad (1 \leq i \leq n) \quad (27)$$

$$v_{x_{t+1}}^i = v_{x_t}^i + a_{x_t}^i \cdot h \quad (1 \leq i \leq n) \quad (28)$$

$$v_{y_{t+1}}^i = v_{y_t}^i + a_{y_t}^i \cdot h \quad (1 \leq i \leq n) \quad (29)$$

$$(x_t^i - x_t^j)^2 + (y_t^i - y_t^j)^2 \geq (2R)^2 \quad (1 \leq i < j \leq n) \quad (30)$$

$$|x_t^i - c_t^i \cdot b_w| \leq \left(\frac{l}{2} - R\right) \text{ or}$$

$$|y_t^i - r_t^i \cdot b_h| \leq \left(\frac{l}{2} - R\right) \quad (1 \leq i \leq n) \quad (31)$$

where r_t^i and c_t^i are the row number and column number of vehicles at time t , respectively, assuming the zone is a perfect rectangular grid; N and M are the number of rows and columns, respectively; b_w and b_h are block width and block height; l is the street width; R is the vehicle clearance radius; T is number of sampling times; and (s_x^i, s_y^i) is the source coordinates of vehicle i .

In the above mentioned problem setting, (15) to (20) are the physical limit constraints. (23) to (25) describe the initial and final conditions. (26) to (29) are dynamic constraints. (30) is the vehicle-to-vehicle collision avoidance constraint and finally (31) is the vehicle-to-boundaries collision avoidance constraint.

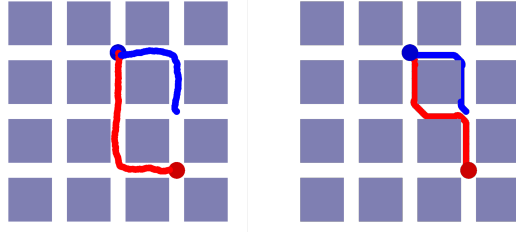


Fig. 4: Learnt policy by (left)RL agent and (right)the baseline method for the small example

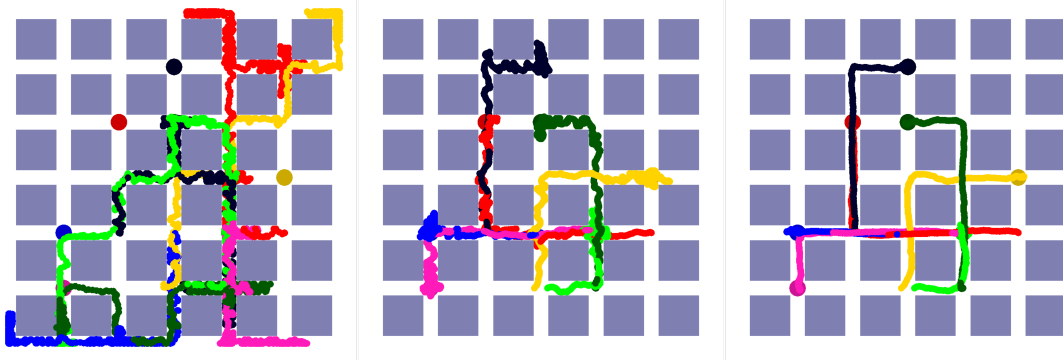


Fig. 5: Learnt policy by the AIM agent at different iterations of the training. left: begging of the training, middle: just after the fast learning phase in Fig. 7. right: end of the training.

The basic problem with the above formulation is that constraint (30) leads to a non-convex function and convex optimization algorithms cannot solve this problem. Therefore, a Mixed-Integer Nonlinear Programming (MINLP) algorithm should be used to solve this problem. Our experiments show that even a small-sized problem with two vehicles and 2×2 grid cannot be solved with an MINLP algorithm, i.e. AOA[11], in a reasonable time. To overcome this issue, we should reformulate the optimization problem using 1-norm and introduce new integer variables for the distance between vehicles using the ideas proposed in [12].

To achieve the best convergence and execution time by using a Mixed-integer Quadratic Programming (MIQP), the cost function and all constraints should be linear or quadratic. Furthermore, the ‘‘or’’ logic in (31) should be implemented using integer variables. The full MIQP problem can be written as the following equations and inequalities:

$$\mathbf{a}_t^* = \arg \max_{\mathbf{a}_t} \sum_{t=0}^{T-1} \sum_{i=1}^n (x_t^i - d_x^i)^2 + (y_t^i - d_y^i)^2 \quad (32)$$

s. t. (15) to (29)

$$b_{x_t}^i, b_{y_t}^i \in \{0, 1\} \quad (1 \leq i \leq n) \quad (33)$$

$$b_{x_t}^i + b_{y_t}^i \geq 1 \quad (1 \leq i \leq n) \quad (34)$$

$$c_{x_t}^{i,j}, c_{y_t}^{i,j}, d_{x_t}^{i,j}, d_{y_t}^{i,j} \in \{0, 1\} \quad (1 \leq i < j \leq n) \quad (35)$$

$$c_{x_t}^{i,j} + c_{y_t}^{i,j} + d_{x_t}^{i,j} + d_{y_t}^{i,j} \geq 1 \quad (1 \leq i < j \leq n) \quad (36)$$

$$x_t^i - x_t^j \geq 2Rc_{x_t}^{i,j} - \mathcal{M}(1 - c_{x_t}^{i,j}) \quad (1 \leq i < j \leq n) \quad (37)$$

$$x_t^i - x_t^j \leq -2Rd_{x_t}^{i,j} + \mathcal{M}(1 - d_{x_t}^{i,j}) \quad (1 \leq i < j \leq n) \quad (38)$$

$$y_t^i - y_t^j \geq 2Rc_{y_t}^{i,j} - \mathcal{M}(1 - c_{y_t}^{i,j}) \quad (1 \leq i < j \leq n) \quad (39)$$

$$y_t^i - y_t^j \leq -2Rd_{y_t}^{i,j} + \mathcal{M}(1 - d_{y_t}^{i,j}) \quad (1 \leq i < j \leq n) \quad (40)$$

$$x_t^i - c_t^i b_w \leq \left(\frac{l}{2} - R\right) b_{x_t}^i + \mathcal{M}(1 - b_{x_t}^i) \quad (1 \leq i \leq n) \quad (41)$$

$$x_t^i - c_t^i b_w \geq -\left(\frac{l}{2} - R\right) b_{x_t}^i - \mathcal{M}(1 - b_{x_t}^i) \quad (1 \leq i \leq n) \quad (42)$$

$$y_t^i - c_t^i b_w \leq \left(\frac{l}{2} - R\right) b_{y_t}^i + \mathcal{M}(1 - b_{y_t}^i) \quad (1 \leq i \leq n) \quad (43)$$

$$y_t^i - c_t^i b_w \geq -\left(\frac{l}{2} - R\right) b_{y_t}^i - \mathcal{M}(1 - b_{y_t}^i) \quad (1 \leq i \leq n) \quad (44)$$

where \mathcal{M} is a large positive number.

(37) to (40) represent the vehicle-to-vehicle collision avoidance constraint using 1-norm:

$$\|(x_t^i, y_t^i)^\top - (x_t^j, y_t^j)^\top\|_1 \geq 2R \quad (45)$$

for any two distinct vehicles i and j . This constraint is equivalent to the following:

$$|x_t^i - x_t^j| \geq 2R \text{ or } |y_t^i - y_t^j| \geq 2R \quad \forall t, (1 \leq i < j \leq n) \quad (46)$$

The absolute value function displayed in (46) should be replaced by logical ‘‘or’’ of two linear conditions to avoid

nonlinearity. Therefore, we obtain the following constraint which is represented by (37) to (40):

$$\begin{aligned} x_t^i - x_t^j &\geq 2R \text{ or } x_t^i - x_t^j \leq -2R \text{ or} \\ y_t^i - y_t^j &\geq 2R \text{ or } y_t^i - y_t^j \leq -2R \quad \forall t, (1 \leq i < j \leq n) \end{aligned} \quad (47)$$

(35) implements the “or” logic required in (47).

(41) to (44) describe the vehicle-to-boundaries collision avoidance constraint:

$$\begin{aligned} |x_t^i - c_t^i b_w| &\leq \left(\frac{l}{2} - R\right) b_{x_t^i} \text{ or} \\ |y_t^i - r_t^i b_w| &\leq \left(\frac{l}{2} - R\right) b_{y_t^i} \quad \forall t, (1 \leq i \leq n) \end{aligned} \quad (48)$$

which is equivalent to:

$$\begin{aligned} (x_t^i - c_t^i b_w \leq \left(\frac{l}{2} - R\right) b_{x_t^i} \text{ and } x_t^i - c_t^i b_w \geq -\left(\frac{l}{2} - R\right) b_{x_t^i}) \text{ or} \\ (y_t^i - r_t^i b_w \leq \left(\frac{l}{2} - R\right) b_{y_t^i} \text{ and } y_t^i - r_t^i b_w \geq -\left(\frac{l}{2} - R\right) b_{y_t^i}) \\ \forall t, (1 \leq i \leq n) \end{aligned} \quad (49)$$

The “or” logic in this constraint is realized in (34).

We will show in the next subsection that the explained conventional optimization formulation is not feasible except for very small-sized problems. Another limitation that makes the conventional method impractical is that this formulation works only for a perfect rectangular grid. However, the proposed RL method in this paper can be extended to arbitrary street layouts.

B. Simulation results

The implementation of the TRPO in rllab library [4] is used to simulate the RL formulation of the AIM problem described in Section IV. For this purpose, the AIM state transition and reward calculation are implemented as an OpenAI Gym [13] environment.

The neural network used to approximate the policy is an Multilayer Perceptron (MLP) which consists of three hidden layers. Each hidden layer has 100 nodes (Fig. 6). Table I lists the parameters for the simulation. To speed up simulation, normalized units are used for the physical properties of the environment instead of real-world quantities.

Fig. 3 shows the small and large grid plans used for the simulation. The small and large circles represent the source and destination locations, respectively. The vehicles are placed at the intersections randomly at the beginning of each episode.

TABLE I: Parameter value settings for the experiment

Parameter	Value
discount factor(γ)	0.999
distance reward penalty factor(α)	0.1
distance tolerance(η)	0.05
maximum speed(v_m)	0.8
maximum acceleration(a_m)	30
sampling time(h)	10 (ms)
maximum episode length	200
vehicle safe radius	0.02

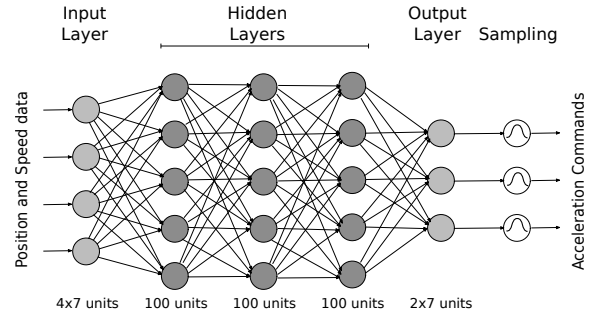


Fig. 6: Neural network used in the simulations.

TABLE II: Total travel time obtained by baseline method and proposed method

	Baseline Method	Proposed Method
Small example	1.79 (s)	2.43 (s)
Large Example	no solution	11.2 (s)

The destinations are also chosen randomly. When the simulator is reset, the same set of source and destination locations are used.

The small grid can be solved both by the baseline method and by the proposed RL method. However, the large grid can only be solved by the RL method because the MIQP algorithm could not find a feasible solution (which is not optimal necessarily) and was stopped after around 68 hours and using 21 GB of system memory. On the other hand, the RL method can solve the problem using 560 MB of system memory and 101 MB of GPU memory.

Table II and Fig. 4 show the comparison of proposed RL and baseline method results. In Table II the total travel time is provided for both methods and Fig. 4 shows the vehicles’ trajectories by running the navigation policy obtained by both solutions for the small examples.

The learning curve of the RL agent which is the expected return vs the training epoch number is shown in Fig. 7 for the large grid example. This figure shows that the learning rate is higher at the beginning which corresponds to the stage where in the agent is learning the very basics of driving and avoiding collisions, but improving the policy towards the optimal policy takes considerably more time. The increase in learning occurs after two epochs when the agent discovers the policy that successfully drives all the vehicles to the destination and the positive terminal reward is gained. Moreover, the trajectories of vehicles are depicted in Fig. 5 at three stages of the learning process, i.e. at the early stage, at epoch 2 where the learning curve slope rapidly decreases, and the end of the training.

The total number of “near collision” incidents discussed in Section IV is shown in Fig. 8. Fig. 9 shows the total travel time as a function of training iteration.

VI. CONCLUSION

In this paper, we have shown that Deep RL can be a promising solution for the problem of intelligent intersection management in local road settings where the number of vehicles is limited and fine-grained acceleration control and

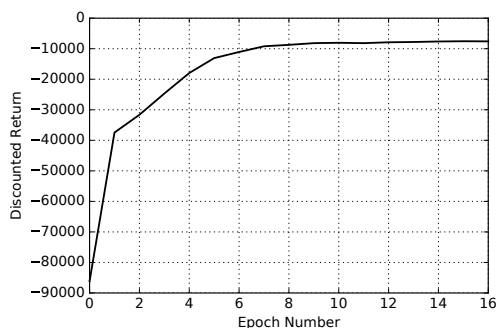


Fig. 7: Learning curve of the AIM agent for large grid example. The discounted return is always a negative value because the return is the accumulated negative distance rewards and there is only one positive reward in the terminal state in which all the vehicles are at the destinations.

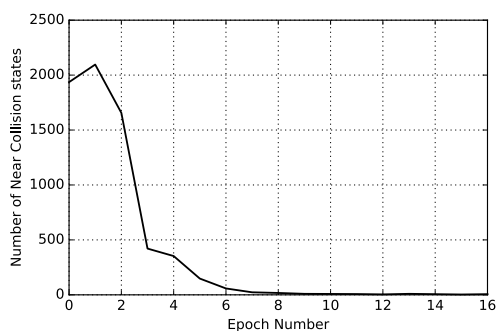


Fig. 8: Number of near collision incidents vs training iteration number.

motion planning can lead to a more efficient navigation of the autonomous vehicles. We proposed an RL environment definition in which collisions are avoided using a safety mechanism. Using this method instead of large penalties for collision in the reward function, the agent can learn the optimal policy faster and the learned policy can be used in practice where the safety mechanism is actually implemented. The experiments show that the conventional optimization methods are not able to solve the problem with the sizes that are solvable by the proposed method.

Similar to the learning process of human beings, the main benefit of the RL approach is that an explicit mathematical modeling of the system is not required and, more importantly, the challenging task of control design for a complex system is eliminated. However, since the automotive systems demand a high safety requirement, training of the RL agent using a simulation model is inevitable in most cases. However, developing a simulation model for a system is considerably simpler task compared to explicit modeling especially for systems with uncertainty.

While the work at hand is a promising first step towards using RL in autonomous intersection management, a number of potential improvements can be mentioned that might be interesting to address in future work. First, the possibility of developing pre-trained neural networks similar to the works in

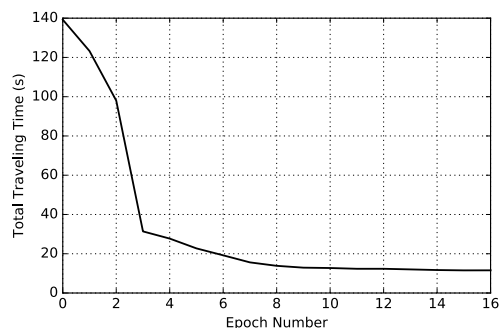


Fig. 9: Total travel time of all the vehicles vs training iteration number.

other mainstream deep learning domains that reduce learning time can be a studied in future. Furthermore, a more advanced rewards system that includes gas usage penalties is another track towards developing a practical intelligent intersection management algorithm.

VII. ACKNOWLEDGEMENT

We would like to thank Nvidia for their generous hardware donation. This work was supported in part by the National Science Foundation under NSF grant number 1563652.

REFERENCES

- [1] M. Hausknecht, T.-C. Au, and P. Stone, "Autonomous intersection management: Multi-intersection optimization," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 4581–4586.
- [2] K. Dresner and P. Stone, "A multiagent approach to autonomous intersection management," *Journal of artificial intelligence research*, vol. 31, pp. 591–656, 2008.
- [3] C. Frese and J. Beyerer, "A comparison of motion planning algorithms for cooperative collision avoidance of multiple cognitive automobiles," in *Intelligent Vehicles Symposium (IV), 2011 IEEE*. IEEE, 2011, pp. 1156–1162.
- [4] Y. Duan, X. Chen, R. Houthoofd, J. Schulman, and P. Abbeel, "Benchmarking deep reinforcement learning for continuous control," *arXiv preprint arXiv:1604.06778*, 2016.
- [5] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, "Trust region policy optimization," *CoRR*, abs/1502.05477, 2015.
- [6] M. N. Mladenović and M. M. Abbas, "Self-organizing control framework for driverless vehicles," in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. IEEE, 2013, pp. 2076–2081.
- [7] I. H. Zohdy, R. K. Kamalanathsharma, and H. Rakha, "Intersection management for autonomous vehicles using icacc," in *2012 15th International IEEE Conference on Intelligent Transportation Systems*. IEEE, 2012, pp. 1109–1114.
- [8] A. A. Malikopoulos, C. G. Cassandras, and Y. J. Zhang, "A decentralized optimal control framework for connected and automated vehicles at urban intersections," *arXiv preprint arXiv:1602.03786*, 2016.
- [9] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press Cambridge, 1998, vol. 1, no. 1.
- [10] B. Karlik and A. V. Olgac, "Performance analysis of various activation functions in generalized mlp architectures of neural networks," *International Journal of Artificial Intelligence and Expert Systems*, vol. 1, no. 4, pp. 111–122, 2011.
- [11] M. Hunting, "The aimms outer approximation algorithm for minlp," *Paragon Decision Technology, Haarlem*, 2011.
- [12] T. Schouwenaars, B. De Moor, E. Feron, and J. How, "Mixed integer programming for multi-vehicle path planning," in *Control Conference (ECC), 2001 European*. IEEE, 2001, pp. 2603–2608.
- [13] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," 2016.