

# Reinforcement Learned Distributed Multi-Robot Navigation with Reciprocal Velocity Obstacle Shaped Rewards

Ruihua Han<sup>1,2</sup>, Shengduo Chen<sup>2</sup>, Shuaijun Wang<sup>2</sup>, Zeqing Zhang<sup>1</sup>, Rui Gao<sup>2</sup>, Qi Hao<sup>2,†</sup>, Jia Pan<sup>1,†</sup>

**Abstract**—The challenges to solving the collision avoidance problem lie in adaptively choosing optimal robot velocities in complex scenarios full of interactive obstacles. In this paper, we propose a distributed approach for multi-robot navigation which combines the concept of reciprocal velocity obstacle (RVO) and the scheme of deep reinforcement learning (DRL) to solve the reciprocal collision avoidance problem under limited information. The novelty of this work is threefold: (1) using a set of sequential VO and RVO vectors to represent the interactive environmental states of static and dynamic obstacles, respectively; (2) developing a bidirectional recurrent module based neural network, which maps the states of a varying number of surrounding obstacles to the actions directly; (3) developing a RVO area and expected collision time based reward function to encourage reciprocal collision avoidance behaviors and trade off between collision risk and travel time. The proposed policy is trained through simulated scenarios and updated by the actor-critic based DRL algorithm. We validate the policy in complex environments with various numbers of differential drive robots and obstacles. The experiment results demonstrate that our approach outperforms the state-of-art methods and other learning based approaches in terms of the success rate, travel time, and average speed. Source code of this approach is available at [https://github.com/hanruihua/rl\\_rvo\\_nav](https://github.com/hanruihua/rl_rvo_nav).

**Index Terms**—Collision Avoidance, Reinforcement Learning, Multi-Robot Systems.

## I. INTRODUCTION

MULTI-ROBOT navigation systems have been widely used in many applications to improve productivity and reduce labor costs. In general, there are two types of multi-robot navigation systems: *centralized* and *distributed* methods. In a centralized system, the controller can flexibly coordinate multiple robots in the same workspace to avoid

collisions given complete information about the whole swarm. The commercial centralized planners normally plan paths for robots without considering collisions at first and then employ scheduling schemes to avoid collisions at potential conflict points, such as the crossroad of the planned paths [1]. However, it is known that the centralized system requires more computation budget when the number of robots increases, and it may also suffer from the signal delay or instability between each individual robot and the central controller.

The distributed multi-robot system allows each robot to make decisions independently based on onboard sensors. As such, it is suitable for deploying a large number of robots with a relative low computational budget. One of its major challenges is how to achieve reliable collision avoidance with limited sensing information and determine optimal velocities independently with sufficient safety and high efficiency.

Some approaches design the dynamically-feasible trajectories for each robot within a time interval in real time to achieve collision-free navigation, followed by replanning the new trajectories in the next time horizon [2]. However, the finer discretization of time requires more computation resources to get more accurate results. Other approaches calculate the optimal velocities directly with low costs, such as the potential field based approaches, where the concept of potential field, including artificial attraction and repulsion, is utilized to find the collision-free and time-efficient velocities [3]. Nevertheless, the problem of handling the local minima is challenging.

Velocity obstacle (VO) based approaches as well as their extensions [4]–[7] are widely used for dynamic collision avoidance, which predicts collision regions for robots and determines the velocity of each robot to avoid reaching the union of these regions in real time. In contrast, deep reinforcement learning (DRL) based approaches represent environment state with suitable data structure and train the neural networks with manually designed rewards [8]–[10]. DRL approaches are advantageous in converting the rich training experiences into capabilities of taking multiple steps ahead into account and achieving more aggressive movement decisions.

However, there are three challenges when developing a fully functional DRL based multi-robot navigation system:

- 1) **Interactive environment state representation.** How to develop a proper form of environment state representation which can explicitly describe the collision avoidance interactions between those robots is still an open problem.
- 2) **Efficient mapping from sequential environment states to continuous control actions.** How to achieve the optimal

Manuscript received: September, 25, 2021; Revised January, 25, 2022; Accepted March, 4, 2022.

This paper was recommended for publication by Editor Hanna Kurniawati upon evaluation of the Associate Editor and Reviewers' comments.

<sup>1,2</sup> R. Han is with the Department of Computer Science, The University of Hong Kong, Hong Kong, and also with the Department of Computer Science and Engineering, Southern University of Science and Technology, Shenzhen, Guangdong, China [hanrh@connect.hku.hk](mailto:hanrh@connect.hku.hk)

<sup>2</sup> S. Chen, S. Wang, R. Gao, and Q. Hao are with the Department of Computer Science and Engineering, Southern University of Science and Technology, and are partially supported by the Shenzhen Fundamental Research Program (No: JCYJ20200109141622964) [hao.q@sustech.edu.cn](mailto:hao.q@sustech.edu.cn)

<sup>1</sup> Z. Zhang and J. Pan are with the Department of Computer Science, University of Hong Kong, and are partially supported by HKSAR RGC GRF 11202119, 11207818, and the Innovation and HKSAR Technology Commission under the InnoHK initiative.

<sup>†</sup> denotes the corresponding authors

Digital Object Identifier (DOI): see top of this page.

and precise control actions for collision avoidance given the current sequential environment state using neural networks at a low computational cost is still challenging.

- 3) **Reward design for collision avoidance behavior regulation.** There is still no systematic method to design rewards according to the observations for representing the collision risk precisely and guiding robots to achieve reciprocal collision avoidance (RCA) behaviors.

Some methods use a set of vectors containing positions and velocities from multiple robots as the policy input [11], [12], which does not directly describe the collision avoidance interaction constraints among robots and hence demands extra capabilities of DRL networks to derive those constraints. To map the state of a varying number of robots to actions, prior works utilize the raw sensor data to train the policy [9]. Other approaches use recurrent neural networks (RNNs) to extract invariant features from the sequential input and output the probability mass function to select the optimal discrete actions [10], [13]. However, the high computational cost of the control action limits the applications of these approaches in the real world, which has been studied to reduce by proposing a novel DRL framework [14]. For the reward design, a dynamic window approach (DWA) is used to encourage those local robots' behaviors towards dynamically feasible areas [15]. However, such a method does not encourage reciprocal collision avoidance behavior among multiple robots.

To this end, we propose a distributed multi-robot navigation approach combining the benefits of the VO concept with the DRL framework to achieve the reciprocal collision avoidance task under limited surrounding information in the shared workspace. Specifically, we utilize VO vectors to model both dynamic agents and static obstacles and employ bidirectional gated recurrent units (GRUs) based neural network for feature extraction and action calculation under the continuous space. The reward function is also designed based on the RVO area. The main contributions of this work include:

- 1) We represent the information of surrounding robots and obstacles in terms of RVO and VO vectors respectively to achieve a unified environment state representation, which enables each robot to explicitly describe its (reciprocal) collision avoidance interactions with both dynamic agents and static obstacles.
- 2) We develop a bidirectional gated recurrent units (BiGRUs) based neural network to extract environment features from the sequential input and map the features to the continuous control action directly with low cost.
- 3) We develop a novel reward function based on the RVO or VO areas and expected collision time, which encourages robots to learn reciprocal local collision avoidance behaviors under diverse situations.

Section II reviews the related work, followed by the introduction of the system framework and the problem statement in Section III. Then, Section IV describes the deep reinforcement learning network structure, and Section V provides the training algorithm in detail. The proposed method is validated by experiments in Section VI. Finally, the conclusion and future work are presented in Section VII.

## II. RELATED WORK

### A. Reciprocal Velocity Obstacle

The central concept of VO is to formulate a potential collision area for a moving obstacle using the relative velocity and position. A velocity outside this area is then chosen for the robot to complete the collision avoidance task. Based on this concept, RVO has been developed to achieve reciprocal collision avoidance (RCA) [16], where each robot has a similar navigation policy and uses an equal effort to avoid each other. RVO has been further developed into optimal reciprocal collision avoidance (ORCA) [4], where several half-plane constraints in the velocity space are used to help the robot find the optimal velocity through linear programming. Non-Holonomic ORCA (NH-ORCA) combines the kinematic model with ORCA to guarantee the ORCA performance under non-holonomic constraints [6]. However, the perfect sensing assumption limits the performance of these approaches in real-world applications. Therefore, some approaches combine VO based methods with robot kinematic constraints to account for the uncertainties in localization and control by increasing the collision radius of the robot [5], [17]. Other approaches propose a probabilistic variant of RVO (PRVO) to overcome Gaussian uncertainties [18]. Despite all, most of these VO based approaches are overly conservative, which inevitably reduces navigation efficiency.

### B. Deep Reinforcement Learning

On the other hand, DRL based collision avoidance approaches can take a large number of training experiences into account and are advantageous in tackling complex scenarios with high efficiency and robustness. Specifically, collision avoidance with DRL (CADRL) employs a value network to generate a collision-free path toward the goal [19]. Some approaches focus on the sensor-level collision avoidance policy that maps raw sensor data to the robot control vector with end-to-end training [9], [20], [21].

Compared with sensor-level methods, agent-level DRL approaches use environment models other than raw sensor data to achieve high computational efficiency as well as flexibility for sensor modalities and kinematic/dynamic details [11]. However, the dimension of the input data for a neural network is required to be fixed. Thus, for the environment model with the time-varying number of surrounding robots, some approaches assume that the number of obstacles is a constant and has an upper limit [12]. RNNs are able to tackle a variable number of moving obstacles, such as in GA3C-CADRL [10], where the exteroceptive measurements at each step are rearranged as the sequential input data through the long short-term memory (LSTM) module to produce a fixed-size feature vector of the environment. However, such a fixed-size vector likely gives more weights to the final section of the input sequences, which limits the performance in the obstacle-dense environment. Thus, the socially aware RL (SARL) has been proposed [22], which utilizes the self-attention mechanism to infer the relative importance of the surrounding dynamic obstacles but without emphasizing the reciprocal collision avoidance interactions and the ability of tackling noncircular

obstacles. In this paper, we use RVO vectors as input instead of robot positions and velocities, which can better describe the reciprocal collision interactions among robots and model the obstacles with various shapes. A bidirectional GRUs based neural network is designed to map the state of a varying number of surrounding robots to the continuous action directly with low-cost computation.

Moreover, a proper reward function is the central part of neural policy training. For most DRL based navigation approaches, the reward functions are designed based on the distance among robots, goals, and obstacles, which represent the collision possibility [9], [23]. Both behaviors of moving to the goal and keeping away from obstacles can be assigned positive and negative rewards easily. In this paper, the reward function is designed based on the RVO area and using the expected collision time to represent the collision risk, which encourages the RCA regular behaviors.

### III. SYSTEM SETUP AND PROBLEM STATEMENT

#### A. System Framework

In our system, a group of differential drive robots navigates in a shared workspace with a series of positions and velocities along  $x$  and  $y$  directions. Each robot has no communication with the others but can sense the surrounding robots within a certain range. The information for each robot includes the robot radius, current velocity, current orientation, and desired velocity. The desired velocity is the maximum velocity for the robot moving from the current position to its goal directly without considering obstacles. Similarly, the information of the surrounding robots consists of the robot radiuses, relative positions, and velocities, which can be represented in terms of RVO vectors [16]. Moreover, the static line obstacles can also be represented by VO vectors [24]. These vectors are fed into a neural network to train the collision avoidance policy. The reward function, which helps evaluate the states and actions, is designed based on the VO area and expected collision time under the current velocities. The navigation policy is trained in simulated scenarios and optimized by using the PPO algorithm. Whereas the policy output of most RL based approaches is the control vector for the robot directly, that of our approach is an increment of the current velocity at each time step, that is, the change rate of the current velocity, which can achieve smoother and efficient robot control. Finally, the output velocity is decomposed into the linear and angular components to control the differential drive robot.

#### B. Reciprocal Velocity Obstacle

This section summarizes the geometrical definition of the VO and RVO. For disc-shaped robot  $A$  and robot  $B$  with radiuses  $R_a$  and  $R_b$ , the positions and velocities can be denoted as  $\mathbf{p}_A$ ,  $\mathbf{p}_B$ ,  $\mathbf{v}_A$ ,  $\mathbf{v}_B$ , respectively. The VO area of the robot  $A$  generated by the robot  $B$  can be given by the following formula:

$$VO_B^A(\mathbf{v}_B) = \{\mathbf{v}_A | \lambda(\mathbf{p}_A, \mathbf{v}_A - \mathbf{v}_B) \cap B \oplus -A \neq \emptyset\}, \quad (1)$$

where  $\lambda(\mathbf{p}, \mathbf{v})$  denotes the ray with a starting point of  $\mathbf{p}$  and in the direction of  $\mathbf{v}$ , and  $\oplus$  is the Minkowski sum. As shown

in Fig. 1 (a), a VO area can be constructed by an apex, where  $\mathbf{v}_p = \mathbf{v}_B$ , and two direction vectors,  $\mathbf{v}_l$ ,  $\mathbf{v}_r$ . It represents the velocity set that might incur collisions for a robot within a period. To achieve collision avoidance from obstacles, robot  $A$  should select a velocity outside the VO area, that is  $\mathbf{v}'_A \notin VO_B^A(\mathbf{v}_B)$ . Similarly, the VO area of the static line obstacle is constructed using the same definition, which is shown in Fig. 1 (b). In addition, for the situation that the static obstacle blocks the path from the current position of a robot to its goal, a global navigation strategy is required to generate a rational desired velocity.

As VO based methods have been widely used for avoiding dynamic obstacles, RVO is more suitable for a group of robots to avoid each other actively. As illustrated in Fig. 1 (c), RVO is extended from the VO concept geometrically and can be described using the following formula:

$$RVO_B^A(\mathbf{v}_B, \mathbf{v}_A) = \{\mathbf{v}'_A | 2\mathbf{v}'_A - \mathbf{v}_A \in VO_B^A(\mathbf{v}_B)\}. \quad (2)$$

Actually, a RVO area can be obtained by translating from a VO area whose apex is moved from  $\mathbf{v}_B$  to  $\frac{\mathbf{v}_A + \mathbf{v}_B}{2}$ . Thus, both VO and RVO can be represented by a six-dimensional vector  $\mathbf{c} = [\mathbf{v}_p, \mathbf{v}_l, \mathbf{v}_r] \in \mathbb{R}^6$ , where  $\mathbf{v}_p = [v_x, v_y]$  denotes the coordinates of the apex,  $\mathbf{v}_l = [vl_x, vl_y]$  and  $\mathbf{v}_r = [vr_x, vr_y]$  describe the direction of left and right rays respectively. Within a robot group, each robot selects a velocity outside such a joint VO and RVO areas to achieve the collision avoidance task cooperatively.

#### C. Problem Statement

The multi-robot collision avoidance problem can be defined as an optimization problem that finds a series of optimal velocities to minimize the travel time under the collision avoidance constraints. For a group of  $n$  differential drive robots navigating in the shared workspace, each robot  $i$  with radius  $R_i$ , state  $\mathbf{p}_t^i = [p_{xt}^i, p_{yt}^i]$ , and velocity  $\mathbf{v}_t^i = [v_{xt}^i, v_{yt}^i]$  can sense  $m$  surrounding robots  $j$  with a limited range  $d_l$  at time  $t$ . The robot's observation has two parts, the proprioceptive measurement  $\mathbf{o}_{\text{self}}$  of the state of the ego-robot and the exteroceptive measurement  $\mathbf{o}_{\text{sur}}$  of the surrounding environment. The proprioceptive observation includes the robot's current velocity, orientation, its desired velocity, and its virtual radius for collision avoidance, *i.e.*,  $\mathbf{o}_{\text{self}} = [v_t, \text{ori}, \mathbf{v}_t^{\text{des}}, R_c]$ . The exteroceptive observation contains the RVO vectors  $\mathbf{c}$ , as well as the relative distance  $d$  and expected collision time  $t_e$  between the ego-robot and its  $j$ -th neighboring robot or obstacle under the current velocity. While the expected collision time will be an infinite value when there is no collision possible under the current situation, which will influence the subsequent calculations. Thus, we use the reciprocal value  $r_e$  to replace that, *i.e.*,  $r_e = 1/(t_e + 0.2)$ , where 0.2 is a constant value. Hence, the exteroceptive observation should be:  $\mathbf{o}_{\text{sur}}^j = [c_j, d_j, r_{e_j}]$ , where  $j = 0, 1, \dots, m$ . Given the proprioceptive/exteroceptive observations as the input, the policy neural network  $\pi_\theta$  parameterized by  $\theta$  will output the actions  $\mathbf{a}_t$ , which are the robot's velocity increment in each time step, *i.e.*,  $\Delta \mathbf{v}_t = [\Delta v_{xt}, \Delta v_{yt}]$ . The computed velocity increments are optimal if the resulting velocity in the next




Fig. 1: An illustration of velocity obstacle and reciprocal velocity obstacle for the dynamic agent and static line obstacle.

time step is not only collision-free but also has a minimal difference to the robot's desired velocity. In other words, we need to solve the following constrained optimization problem at each time step  $\Delta t$  for each robot:

$$\begin{aligned} \arg \min_{\pi_\theta} \quad & \|\mathbf{v}_t - \mathbf{v}_t^{\text{des}}\|, \\ \text{s.t.} \quad & \Delta \mathbf{v}_t \sim \pi_\theta(\mathbf{a}_t \mid \mathbf{o}_{\text{self}}, \mathbf{o}_{\text{sur}}), \\ & \mathbf{v}_t = \mathbf{v}_{t-1} + \mu \cdot \Delta \mathbf{v}_t, \\ & \mathbf{p}_t = \mathbf{p}_{t-1} + \Delta t \cdot \mathbf{v}_t, \\ & d_j = \|\mathbf{p}_t - \mathbf{p}_{jt}\|, \\ & \forall j \in [1, m], d_j > R_c + R_{jc}, \end{aligned} \quad (3)$$

where  $\|\cdot\|$  is the Euclidean norm operation for a vector,  $\mu$  is the hyperparameter adjusting the range of the velocity increment, and  $d_j$  is the distance between a robot and its  $j$ -th neighbor with position  $\mathbf{p}_{jt}$  and collision radius  $R_{jc}$ . All the robots share the same navigation policy  $\pi_\theta$  and find the optimal velocity independently.

#### IV. REINFORCEMENT LEARNING FRAMEWORK

An actor-critic reinforcement learning framework consists of observation space, action space, reward function, policy actor, and policy critic. The policy actor in the form of a deep neural network maps the observed information into an action to control the robot. The policy critic in the form of a deep neural network utilizes the reward function to evaluate each action made by the actor.

##### A. Observation Space and Action Space

The observation space contains the proprioceptive and exteroceptive measurements, as mentioned in Section III,  $\mathbf{o} = [\mathbf{o}_{\text{self}}, \mathbf{o}_{\text{sur}}^j], j = 0, \dots, m$ . The action space is the velocity increment within the  $x-y$  plane,  $\mathbf{a} = [\Delta v_x, \Delta v_y]$ . Thus, the control vector at the next time step should be  $\mathbf{v}_{t+1} = \mathbf{v}_t + \mu \cdot \mathbf{a}$ . Specifically, the velocity is clipped within a range between a maximum and a minimum value,  $\mathbf{v}_t \in [\mathbf{v}_{\text{min}}, \mathbf{v}_{\text{max}}]$ . However, all the robots are non-holonomic and controlled by the transitional and rotational velocity,  $\mathbf{v}_t^c = [v_t, \omega_t]$ . Thus, we convert the orthogonal velocity  $\mathbf{v}_t = [v_{xt}, v_{yt}]$  to  $\mathbf{v}_t^c$  as follows:

$$\begin{cases} v_t = \|\mathbf{v}_t\| \cdot \cos(\varsigma) \\ \omega_t = -\varsigma/\tau \end{cases}, \quad (4)$$

where  $\varsigma$  is the radian difference between the orientations of the robot and velocity  $\mathbf{v}_t$ , and  $\tau$  is the guaranteed time to adjust the

rotation rate. Thus, the transitional velocity is determined by the speed component along with the robot orientation. And the rotational velocity is to make the robot orientation consistent with the direction of orthogonal velocity.

##### B. Neural Network Architecture

The architecture of the deep neural network is illustrated in Fig. 2. For the variable-length input sequences, existing work uses the LSTM to produce a fixed-length representation, which processes inputs in a strict order [10]. By contrast, we use the bidirectional GRU module, which consists of two GRUs to process the inputs with both forward and backward directions, respectively. Such a scheme helps the network more accurately to find the underlying relationship among input vectors better depending on the limited information, resulting in a better presentation of the input sequences. At each step, the exteroceptive measurements  $\mathbf{o}_{\text{sur}}^j$  in terms of a series of vectors are fed into the BiGRU in the first ascending order of  $r_e$  and the second descending order of the distance  $d$ . There are two final hidden states,  $\mathbf{h}_{\text{for}}$  and  $\mathbf{h}_{\text{back}}$ , from the forward and backward GRU, respectively. They are added as one 256-D fixed-size vector  $\mathbf{h}_m$ . This output is concatenated with the proprioceptive measurement  $\mathbf{o}_{\text{self}}$  as the integrated fixed-length observation  $\mathbf{o}$ . To achieve a normalized representation to speed up the training process, where all parts have a common scale, we use the Layer Normalization method [25] to process fixed-length observation. The output is used as the input data of two neural networks. One is for the policy actor  $\pi_\theta$ , which consists of two hidden fully connected (FC) layers with 256 rectifier linear units. The output of the final layer is a two-dimensional vector, which is the mean of velocity  $\mathbf{v}_{\text{mean}}$ , activated by a hyperbolic tangent activation function (Tanh). The action is sampled from a Gaussian distribution represented by its mean,  $\mathbf{v}_{\text{mean}}$ , and log standard deviation  $\mathbf{v}_{\text{log std}}$ , where  $\mathbf{v}_{\text{log std}}$  is updated independently. Another neural network is for the policy critic  $V_\psi(\mathbf{o})$ , which is also composed of two hidden fully-connected layers with 256 rectifier linear units. The output of its final layer activated by an identity function is a scalar value used to evaluate each state-action pair.

##### C. Reward Function

In the reinforcement learning framework, the design of the reward function is vital to policy performance. The common

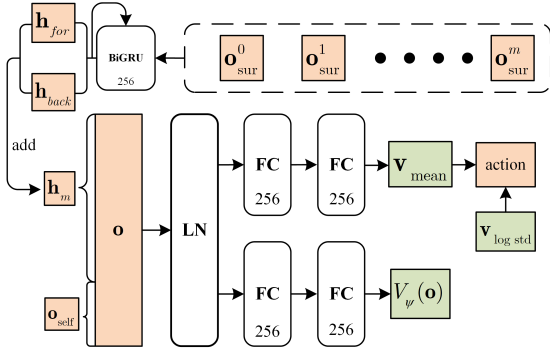


Fig. 2: An illustration of the proposed BiGRU based neural network architecture of the navigation policy actor and critic.

distance based reward function, such as the penalties for the increasing distance to obstacles or the awards for the decreasing distance to the goal [9], [10], is improper for our RVO observations, which do not include the information of positions and velocities directly. Thus, a novel reward function based on RVO areas is developed, which helps each robot to learn a reciprocal collision avoidance behavior. In addition, based on our previous experiences, the expected collision time performs better than the distance-to-obstacle in terms of representing the collision risk, as the expected collision time takes both the distance-to-obstacle and relative velocity between the robot and the obstacle into account. Based on this concept, the RVO reward function  $r_{rvo}^t$  at time  $t$  is designed as follows, which represents the quality of the selected velocity  $\mathbf{v}_t$  judged by the joint RVO area:

$$r_{rvo}^t = \begin{cases} a - b * diff_v & \text{if } \mathbf{v}_t \notin RVO \text{ or } \xi > 5 \\ c - d * (\xi + f)^{-1} & \text{if } \mathbf{v}_t \in RVO \text{ and } \xi > 0.1 \\ -e * (\xi + f)^{-1} & \text{if } \xi \leq 0.1 \end{cases}, \quad (5)$$

where  $diff_v$  is the distance between the selected velocity and desired velocity,  $diff_v = \|\mathbf{v}_t - \mathbf{v}_t^{des}\|$ , and  $\xi$  is the expected minimum time that the robot has a collision with an obstacle under the current velocity. Whether the current velocity  $\mathbf{v}_t$  is within the RVO area is judged by a six-dimensional vector  $\mathbf{c} = [\mathbf{v}_p, \mathbf{v}_l, \mathbf{v}_r]$  as follows:

$$\begin{cases} \mathbf{v}_t \in RVO := (\mathbf{v}_t - \mathbf{v}_p) \times \mathbf{v}_l \geq 0 \wedge (\mathbf{v}_t - \mathbf{v}_p) \times \mathbf{v}_r \leq 0 \\ \mathbf{v}_t \notin RVO := (\mathbf{v}_t - \mathbf{v}_p) \times \mathbf{v}_l < 0 \vee (\mathbf{v}_t - \mathbf{v}_p) \times \mathbf{v}_r > 0 \end{cases}, \quad (6)$$

where  $\times$  is cross product.

The characters  $a, b, c, d, e, f$  are all constant values and tunable to adjust the policy performance.  $a$  and  $c$  are the basic awards to encourage the behavior when the current velocity is close to the desired velocity, or the current collision risk is low, which tend to be positive values.  $b$  is the coefficient to adjust the weight of  $diff_v$ , which represents the motion of moving with the desired velocity. While  $d, e$  are the coefficients to adjust the weight of  $(\xi + f)^{-1}$ , which denotes the motion of keeping longer expected collision time. The range of  $(\xi + f)^{-1}$  is  $[0, f^{-1}]$ . It is the substitute of  $\xi$ , which is inappropriate for policy training directly because of the range  $[0, +\infty]$ . To summarize, for the crowded scenarios, bigger

---

### Algorithm 1: Policy Training Algorithm

---

```

1 Initialize the policy actor  $\pi_\theta$  and critic  $V_\psi$ ;
2 for epoch  $\leftarrow 1, 2, \dots$  do
3   for robot  $i \leftarrow 1$  to  $n$  do // collect data
4     Run policy  $\pi_{\theta_{old}}$  for  $T$  steps;
5     Collect data  $\{\mathbf{o}_t^i, \mathbf{a}_t^i, r_t^i\}$  for  $T$  steps;
6     Compute GAE:  $\hat{A}_1^i, \dots, \hat{A}_T^i$ ;
7     Add data into buffer  $i$ ;
8   for buffer  $j \leftarrow 1$  to  $n$  do // update policy
9     for  $k \leftarrow 1$  to  $K_{iter}$  do
10      Optimize  $L^{CLIP}(\theta)$  w.r.t.  $\theta$  with  $l_a$ ;
11      if KL-divergence  $> \overline{KL}$  then break;
12       $\theta_{old} \leftarrow \theta$ ;
13     for  $h \leftarrow 1$  to  $H_{iter}$  do
14      Optimize  $L^V(\psi)$  w.r.t.  $\psi$  with  $l_v$ ;
15       $\psi_{old} \leftarrow \psi$ ;

```

---

values of  $d$  and  $e$  are recommended to guarantee the collision avoidance ability. The values for our policy training in this paper are  $a = 0.3, b = 1.0, c = 0.3, d = 1.2, e = 3.6, f = 0.2$ . The main idea is that when the velocity is in the joint RVO area, the robot has to pay attention to the potential collision judged by the expected collision time. The constant values of the reward function are all tunable parameters, which can be adjusted to improve the training performance.

## V. POLICY TRAINING

PPO is a popular optimization method for DRL algorithms, which is advantageous for training the policy associated with continuous action spaces. In this work, we utilize the PPO algorithm to train and update the multi-robot collision avoidance policy successively. The detail of the training algorithm is summarized in Algorithm 1. First, this algorithm begins with the initialization of the neural network parameters, including the policy actor  $\pi_\theta$  and critic  $V_\psi$ . Then, during the training loop, each robot runs in the environment with the navigation policy  $\pi_\theta$  for  $T$  timesteps. The data of observation, reward, and action of each robot at each time step  $t$  is collected. Generalized advantage estimator (GAE) is used to estimate the advantage of the action of robot  $i$  using the reward function  $r_t^i$  and value function  $V_\psi$ . The related observation, action, and reward data are stored in the buffer for the following policy update. The buffered data are used to construct the clip surrogate objective  $L^{CLIP}(\theta)$  and loss of value function  $L^V(\psi)$  [26], which are optimized with the Adam optimizer and learning rate  $l_a, l_v$  respectively. The learning rates determine the change rate of the policy model within each update step. The Kullback-Leiber (KL)-divergence measures the difference between two probability distributions of the previous and current policies. The update loop will stop when the KL-divergence value is over a threshold (denoted by  $\overline{KL}$ ), which means that the model changes too fast, and hence such an update process is unsuccessful.

## VI. EXPERIMENTS AND RESULTS

### A. Simulation Setup

In this work, the navigation policy RL-RVO is trained through the simulated scenario and implemented with Pytorch (Python 3.8). The simulated scenarios are developed via the OpenAI Gym interface [27], which is a popular toolkit for developing reinforcement learning algorithms. We use several disc-shaped robots to train the collision avoidance policy through the circle scenario only. In the circle scenario, all robots with random orientations are uniformly arranged along with a circle shape. And the goal position is on the opposite side, which leads to a rich interaction for the robots, as illustrated in Fig. 3 (a). Specifically, the sensing range  $d_l$  and the maximum input number of neighbors  $\kappa$  of each robot are limited to be  $4m$  and 5, respectively. All robots share the same policy and collect the observations for the parameter update by the PPO algorithm after each epoch. The environment will be reset when there are collisions or over the maximum episode length. To speed up the training convergence, we use two stages to perform the training process. First, the policy is trained in a  $9m \times 9m$  circle scenario with a few robots (2 or 4) for  $e_1$  epochs to achieve basic functionality. Then, the policy continues to be trained with 10 robots in the same circle scenario for about  $e_2$  epochs. The training process will be early stopped depending on the success rate and step cost. The values of the hyperparameters in the training process are chosen, as listed in Table I. Additionally, the training process is performed via a computer with CPU i7-9700 and GPU Nvidia GTX 1080 for about 10 hours.




Fig. 3: The simulated scenarios for multi-robot navigation: (a) circle scenario, (b) random scenario, (c) corridor scenario.

TABLE I: The hyperparameters of the training process

Para.	Val.	Para.	Val.	Para.	Val.	Para.	Val.
$d_l$	4	$l_a$	$4e-6$	$K_{iter}$	50	$\kappa$	5
$\tau$	0.2	$l_v$	$5e-5$	$\overline{KL}$	0.01	$\mu$	1.0
$R_c$	0.3	$e_1$	200	$H_{iter}$	50	$T$	450
$R_{jc}$	0.3	$e_2$	1000	$\mathbf{v}_{min}$	-1.5	$\mathbf{v}_{max}$	1.5

### B. Results and Discussions

1) *Metrics*: Three metrics are utilized to evaluate the policy performance, including success rate, travel time, and average speed. The success rate is a ratio of the successful cases without any collision or being stuck somewhere during the navigation, which describes the policy's ability of collision avoidance. The travel time refers to the amount of time, represented by the iteration step in the simulation when all robots arrive at the goal positions, which reflects the policy's

efficiency. The average speed of the navigation process of the whole robot team measures the policy's performance on effective velocity selection. In addition, the average computational cost of a control action for one robot is also compared.

2) *Simulated experiments*: To validate the collision avoidance performances of our policy, we compare it with SARL [22], GA3C-CADRL [10], and NH-ORCA [6] in both circle and random scenarios. The random scenarios are used to test the generalization ability of a navigation policy trained from a specific scenario, where both initial and target positions are generated randomly with a minimum interval ( $1m$ ) and change randomly in each test episode, which are different from the training scenarios, as shown in Fig. 3 (a) and (b). Thus, the performance tests in the random scenario are fair for all policies. To test the ability of collision avoidance under limited information, the sensing range is set to be  $d_l=4m$ . All policies are performed for 100 episodes with various numbers of robots (from 6 to 20). The average results with standard deviations (std) of 100 cases in  $9m \times 9m$  circle scenarios and  $10m \times 10m$  random scenarios are listed in Table II. The trajectories of four approaches are compared visually in Fig. 4. Finally, the average computational cost to achieve single control action for each robot is shown as a bar figure in Fig. 5.

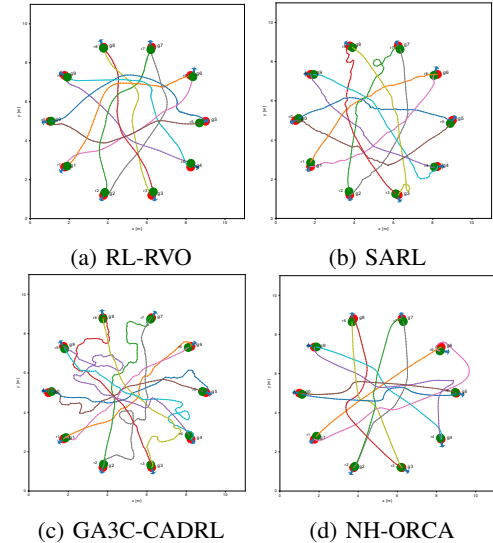


Fig. 4: An illustration of the trajectories generated by four approaches respectively in the circle scenario.




Fig. 5: A comparison of the average computational cost to achieve single control action for each robot.

Generally, as the number of robots increases, the success rate and average speed decrease, whereas the travel time increases. In the circle scenario, compared with the other

TABLE II: A performance comparison of four approaches in the circle and random scenarios with different numbers of robots.

Robot Number	Success Rate				Travel Time (iteration step) / std				Average Speed (m/s) / std			
	RL-RVO	SARL	GA3C-CADRL	NH-ORCA	RL-RVO	SARL	GA3C-CADRL	NH-ORCA	RL-RVO	SARL	GA3C-CADRL	NH-ORCA
<i>sensing range: 4m, circle scenario</i>												
6	<b>1</b>	1	0.98	1	<b>78.29/4.48</b>	78.79/11.32	101.0/35.08	79.26/6.63	<b>1.09/0.06</b>	1.08/0.12	1.05/0.24	1.15/0.09
10	<b>0.99</b>	0.89	0.82	1	<b>90.23/4.16</b>	92.31/18.32	101.96/7.59	91.32/20.3	<b>0.98/0.04</b>	0.92/0.18	0.97/0.15	0.94/0.12
14	<b>0.97</b>	0.82	0.75	0.97	<b>103.13/5.04</b>	115.04/43.61	125.7/27.98	116.25/16.59	<b>0.88/0.04</b>	0.91/0.24	0.93/0.17	0.77/0.1
16	<b>0.93</b>	0.76	0.61	0.89	<b>111.75/7.88</b>	128.85/38.52	141.4/33.76	130.43/25.04	<b>0.83/0.05</b>	0.90/0.31	0.86/0.1	0.72/0.12
20	<b>0.90</b>	0.71	0.34	0.80	<b>128.62/9.52</b>	138.73/33.05	152/30.23	143.51/27.89	<b>0.76/0.07</b>	0.80/0.19	0.75/0.34	0.70/0.13
<i>sensing range: 4m, random scenario</i>												
6	<b>1</b>	0.93	0.81	0.94	<b>75.24/17.59</b>	88.22/22.83	90.49/52.73	77.41/15.49	<b>0.75/0.13</b>	0.58/0.27	0.75/0.24	0.82/0.12
10	<b>0.98</b>	0.89	0.72	0.92	<b>85.88/13.49</b>	89.64/19.78	110.79/45.53	92.53/23.38	<b>0.70/0.11</b>	0.49/0.28	0.62/0.15	0.68/0.19
14	<b>0.97</b>	0.71	0.61	0.84	<b>95.88/17.12</b>	99.38/29.29	123.67/48.93	97.38/32.55	<b>0.63/0.1</b>	0.56/0.28	0.52/0.17	0.62/0.17
16	<b>0.96</b>	0.69	0.42	0.70	<b>106.91/25.1</b>	116.68/33.85	147.2/41.49	109.41/35.18	<b>0.60/0.12</b>	0.49/0.27	0.43/0.1	0.57/0.21
20	<b>0.92</b>	0.60	0.36	0.65	<b>115.25/21.22</b>	122.11/26.72	151.33/43.21	122.33/39.96	<b>0.56/0.13</b>	0.51/0.36	0.42/0.34	0.52/0.21

policies, our policy has the highest success rate, average speed, and lowest travel time, especially in the crowded scenario, as shown in Fig. 6 (a). When there is a small number of robots, the policies perform comparably in the sparse scenario. As the number of robots increases, other policies have a heavy decrease in terms of the success rate, while our policy still works well. Similarly, in the randomly generated scenarios, our policy has better performance in various dense situations, as shown in Fig. 6 (b). SARL and GA3C-CADRL predict the future state of the dynamic obstacles for decision making. However, the selection schemes of the control action are easy to encounter a freezing point problem, which influences the success rate in the random scenario. In comparison, our policy utilizes the BiGRU module to find more underlying relationships among surrounding robots and predict the future situation of the given limited information. Besides, the VO observation and reward guide the robots to achieve RCA behavior, which also improves the ability of collision avoidance and efficiency in dense situations. The results of the standard deviations also demonstrate the stability of our approach. In addition, the computational cost of our approach is lower than the other two learning based approaches a lot, which will influence the real-time performance of the policy. To summarize, compared with other policies, our policy is more robust and efficient in crowded situations under limited information.

3) *Ablation Study*: We also compare our policy with RL-NRVO, RL-LSTM, and RL-Reward to demonstrate the functionality of components of our approach in the circle, random, and corridor scenarios, as shown in Fig. 3. All the policies are trained in the same DRL framework but with different components. RL-NRVO is the policy using the original information of robots as input instead of RVO vectors. This information includes relative positions/velocities and robot radiuses, as introduced in our previous work [12], *i.e.*,  $\mathbf{o}_{\text{sur}} = [p_x, p_y, v_x, v_y, R]$ . Similarly, RL-LSTM is the policy which replaces the BiGRUs with LSTM to tackle the input information of robots with varying numbers. And RL-Reward utilizes the distance-based reward function to train the policy under the same observations, as described in [9]. The results are listed in Table III. It can be seen that other policies can tackle the circle scenario well but have a dramatic performance decrease in the random and corridor scenarios. The reason is that RL-NRVO does not learn an appropriate and regular behavior from the non-RVO input. Using LSTM to represent the robot information with variable numbers is not enough

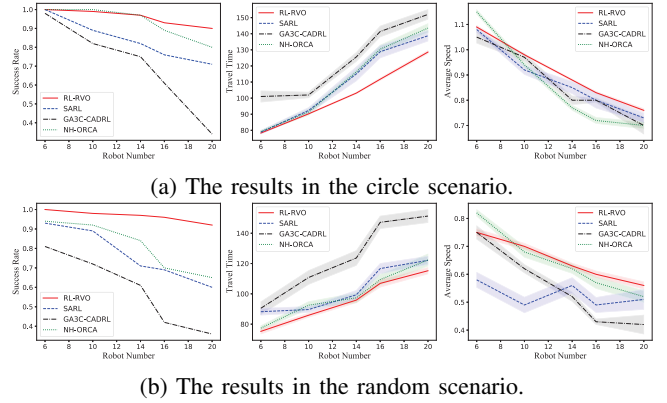


Fig. 6: The testing results of four policies in the simulated scenarios in terms of three metrics.

TABLE III: The results of ablation experiments

Setup	<i>robot number 8, sensing range: 4m, circle/random/corridor scenario</i>		
	Success Rate	Travel Time	Average Speed
<b>RL-RVO</b>	<b>1/1/0.89</b>	<b>83.75/79.71/84.23</b>	<b>1.04/0.73/0.89</b>
RL-NRVO	0.97/0.53/0.29	94.45/85.22/89.31	0.94/0.70/0.84
RL-LSTM	0.96/0.85/0.42	90.05/91.65/85.43	0.96/0.69/0.70
RL-Reward	0.93/0.71/0.58	98.54/88.87/92.61	0.91/0.68/0.67

to tackle the crowded workspace. The distance base reward function is unmatched with the RVO based observations.

4) *Real-world experiments*: The RL-RVO policy is implemented in Turtlebot robots to demonstrate the performance of our policy in the real world. Those experiments use up to 8 differential drive Turtlebots, as shown in Fig. 7. All the Turtlebots are arranged along a circle with random orientations. Each Turtlebot is equipped with a mini PC to compute velocity and a tag to localize itself by the UWB localization system. The robots receive the exteroceptive measurements through the ROS architecture. Specifically, due to the mechanical limit, the maximum velocity of each robot is set to be  $1\text{m/s}$ .

Because of the high computational cost of GA3C-CADRL and SARL, we only compare our policy with NH-ORCA using 4, 6, and 8 Turtlebots in the real world experiment with circle scenario. Statistics across 50 cases are described in Table IV. Different from simulated experiments, there are ubiquitous uncertainties in real-world experiments. NH-ORCA does not take noise and uncertainties into account, thus its performance is less robust in real-world experiments. In contrast, our policy has a better success rate, and it also takes less travel time to accomplish the same navigation tasks than NH-ORCA.

TABLE IV: Results of Real-World Experiment

Test Setup Robot Number	realistics scenario					
	Success Rate		Travel Time (s) / std		Average Speed (m/s) / std	
	RL-RVO	NH-ORCA	RL-RVO	NH-ORCA	RL-RVO	NH-ORCA
4	1	1	<b>13.96/1.09</b>	14.98/1.16	<b>0.81/0.11</b>	0.79/0.21
6	<b>0.91</b>	0.85	<b>16.01/0.92</b>	18.45/1.11	<b>0.74/0.12</b>	0.72/0.17
8	<b>0.89</b>	0.68	<b>22.18/1.25</b>	28.94/1.57	<b>0.69/0.15</b>	0.65/0.23




Fig. 7: An illustration of real-world experiments.

## VII. CONCLUSION

In this paper, we have presented a DRL based multi-robot navigation approach, which utilizes RVO techniques to tackle collision avoidance problems. The proposed RVO based dynamic and static environment state representation can better describe the reciprocal interactions. The developed BiGRUs based neural network can extract environment features despite the varying number of moving robots and map the features to the control actions directly with low cost. The RVO area and expected collision time based reward function can achieve reciprocal collision behaviors and a trade-off between the collision risk and travel time. Both simulated and real-world experiments have been performed respectively to evaluate the policy's navigation performance. Four policies, including RL-RVO, SARL, GA3C-CADRL, and NH-ORCA have been compared in terms of success rate, travel time, average speed, and computational cost in both circle and random scenarios. The experiment results have shown that the proposed approach has a superior collision avoidance capability and time efficiency over other methods in the crowded environment. The ablation study demonstrates the functionality of the individual components in our approach. Furthermore, it has been shown that our policy has been generalized well to tackle situations with more numbers of robots than those used for training. Our future work includes extending such an approach to addressing navigation problems with more challenging uncertainties in the real world.

## REFERENCES

- [1] Z. Zhang, R. Han, and J. Pan, "An efficient centralized planner for multiple automated guided vehicles at the crossroad of polynomial curves," *IEEE Robotics and Automation Letters*, vol. 7, no. 1, pp. 398–405, 2021.
- [2] S. Kousik, S. Vaskov, F. Bu, M. Johnson-Roberson, and R. Vasudevan, "Bridging the gap between safety and real-time performance in receding-horizon trajectory design for mobile robots," *The International Journal of Robotics Research*, vol. 39, no. 12, pp. 1419–1469, 2020.
- [3] M. Guerra, D. Efimov, G. Zheng, and W. Perruquetti, "Avoiding local minima in the potential field method using input-to-state stability," *Control Engineering Practice*, vol. 55, pp. 174–184, 2016.
- [4] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*. Springer, 2011, pp. 3–19.
- [5] D. Hennes, D. Claes, W. Meeussen, and K. Tuyls, "Multi-robot collision avoidance with localization uncertainty," in *AAMAS*, 2012, pp. 147–154.
- [6] J. Alonso-Mora, A. Breitenmoser, M. Ruffli, P. Beardsley, and R. Siegwart, "Optimal reciprocal collision avoidance for multiple non-holonomic robots," in *Distributed autonomous robotic systems*. Springer, 2013, pp. 203–216.

- [7] R. Han, S. Chen, and Q. Hao, "A distributed range-only collision avoidance approach for low-cost large-scale multi-robot systems," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 8020–8026.
- [8] R. Lowe, Y. I. Wu, A. Tamar, J. Harb, O. Pieter Abbeel, and I. Mordatch, "Multi-agent actor-critic for mixed cooperative-competitive environments," *Advances in neural information processing systems*, vol. 30, 2017.
- [9] P. Long, T. Fan, X. Liao, W. Liu, H. Zhang, and J. Pan, "Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 6252–6259.
- [10] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 3052–3059.
- [11] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1343–1350.
- [12] R. Han, S. Chen, and Q. Hao, "Cooperative multi-robot navigation in dynamic environment with deep reinforcement learning," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 448–454.
- [13] M. Everett, Y. F. Chen, and J. P. How, "Collision avoidance in pedestrian-rich environments with deep reinforcement learning," *IEEE Access*, vol. 9, pp. 10 357–10 377, 2021.
- [14] L. Liu, D. Dugas, G. Cesari, R. Siegwart, and R. Dubé, "Robot navigation in crowded environments using deep reinforcement learning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 5671–5677.
- [15] A. J. S. Utsav Patel, Nithish K Sanjeev Kumar and D. Manocha, "Dwa-rl: Dynamically feasible deep reinforcement learning policy for robot navigation among mobile obstacles," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021.
- [16] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *2008 IEEE International Conference on Robotics and Automation*. IEEE, 2008, pp. 1928–1935.
- [17] D. Claes and K. Tuyls, "Multi robot collision avoidance in a shared workspace," *Autonomous Robots*, vol. 42, no. 8, pp. 1749–1770, 2018.
- [18] B. Gopalakrishnan, A. K. Singh, M. Kaushik, K. M. Krishna, and D. Manocha, "Prvo: Probabilistic reciprocal velocity obstacle for multi robot navigation under uncertainty," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017, pp. 1089–1096.
- [19] Y. F. Chen, M. Liu, M. Everett, and J. P. How, "Decentralized non-communicating multiagent collision avoidance with deep reinforcement learning," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 285–292.
- [20] P. Wenzel, T. Schön, L. Leal-Taixé, and D. Cremers, "Vision-based mobile robotics obstacle avoidance with deep reinforcement learning," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 14 360–14 366.
- [21] T. Fan, P. Long, W. Liu, and J. Pan, "Distributed multi-robot collision avoidance via deep reinforcement learning for navigation in complex scenarios," *The International Journal of Robotics Research*, vol. 39, no. 7, pp. 856–892, 2020.
- [22] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6015–6022.
- [23] C. Wang, X. Zhang, L. ze Cong, J. Li, and J. Zhang, "Research on intelligent collision avoidance decision-making of unmanned ship in unknown environments," *Evolving Systems*, pp. 1–10, 2019.
- [24] J. Snape, J. Van Den Berg, S. J. Guy, and D. Manocha, "The hybrid reciprocal velocity obstacle," *IEEE Transactions on Robotics*, vol. 27, no. 4, pp. 696–706, 2011.
- [25] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [26] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [27] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.