

Counterfactual Fairness Filter for Fair-Delay Multi-Robot Navigation

Hikaru Asano*
The University of Tokyo
Tokyo, Japan
asano-hikaru19@g.ecc.u-tokyo.ac.jp

Ryo Yonetani
OMRON SINIC X
Tokyo, Japan
ryo.yonetani@sinicx.com

Mai Nishimura
OMRON SINIC X
Tokyo, Japan
mai.nishimura@sinicx.com

Tadashi Kozuno
OMRON SINIC X
Tokyo, Japan
tadashi.kozuno@sinicx.com

ABSTRACT

Multi-robot navigation is the task of finding trajectories for a team of robotic agents to reach their destinations as quickly as possible without collisions. In this work, we introduce a new problem: fair-delay multi-robot navigation, which aims not only to enable such efficient, safe travels but also to equalize the travel delays among agents in terms of actual trajectories as compared to the best possible trajectories. The learning of a navigation policy to achieve this objective requires resolving a nontrivial credit assignment problem with robotic agents having continuous action spaces. Hence, we developed a new algorithm called Navigation with Counterfactual Fairness Filter (NCF2). With NCF2, each agent performs counterfactual inference on whether it can advance toward its goal or should stay still to let other agents go. Doing so allows us to effectively address the aforementioned credit assignment problem and improve fairness regarding travel delays while maintaining high efficiency and safety. Our extensive experimental results in several challenging multi-robot navigation environments demonstrate the greater effectiveness of NCF2 as compared to state-of-the-art fairness-aware multi-agent reinforcement learning methods.

Project webpage: <https://omron-sinicx.github.io/ncf2/>

KEYWORDS

Multi-agent navigation; Multi-agent reinforcement learning; Counterfactual inference

ACM Reference Format:

Hikaru Asano*, Ryo Yonetani, Mai Nishimura, and Tadashi Kozuno. 2023. Counterfactual Fairness Filter for Fair-Delay Multi-Robot Navigation. In *Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, United Kingdom, May 29 – June 2, 2023, IFAAMAS, 9 pages.

1 INTRODUCTION

Time is equally precious to everyone. Having only certain people experience delays in services where efficiency is critical would lead to unfairness. Familiar examples include food delivery and cab dispatch services: imagine a situation in which your afternoon

meeting is about to start, but only the food that you ordered has not been delivered. Fairness regarding time is also be important in life-threatening situations, such as disaster relief and evacuation.

We are interested in achieving such temporal fairness among multiple autonomous mobile robotic agents. In particular, this work proposes *fair-delay multi-robot navigation*, in which each robot must travel to its destination as fairly as possible in terms of temporal delays, *e.g.*, by avoiding a situation where one agent is an hour late in arriving while the other agents are only 10 minutes late. In the fields of AI and robotics, multi-robot navigation has long been studied as a practical application of multi-agent reinforcement learning (MARL) (*e.g.*, [29, 42]) or multi-agent pathfinding (MAPF) [36]. Typically, each robot must navigate as quickly as possible to its destination while avoiding collisions with obstacles (*e.g.*, walls, shelves) and other robots in motion. Furthermore, the layout of obstacles may have changed since an environmental map was first constructed. Even in such unknown environments, the robots still have to cooperate with each other to reach their destinations efficiently and safely.

By considering fairness for travel delays on top of such a navigation objective, fair-delay multi-robot navigation can be viewed as a novel extension of fairness-aware MARL problems for robot navigation tasks. Specifically, we aim to *find navigation policies that can equalize the delays among agents between their actual trajectories and the best possible trajectories they could have taken by ignoring the presence of other agents, while also pursuing high efficiency and safety.*

Achievement of the above goal requires solving a challenging credit assignment problem of evaluating the impact of an agent’s actions on the success of other agents with continuous action spaces [44]. In multi-robot navigation, situations that require fairness consideration emerge naturally but unpredictably. An agent giving way in one place may cause another agent’s travel to become blocked at another time in another place. For prior fairness-aware MARL methods that simply learn fairness through a reward [22, 48], it would be hard to reason about which actions contribute to fairness. Such reasoning could be possible via *counterfactual inference*, which compares actual events to events that did not occur but could have. Nevertheless, existing counterfactual methods for multi-agent tasks often assume only discrete action spaces [26, 46], which makes it nontrivial how to define *default actions* (*i.e.*, replacement actions for when agents do not take planned actions) for robotic agents that are typically commanded with continuous values. Furthermore,

*Work done during an internship at OMRON SINIC X.

Proc. of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023), A. Ricci, W. Yeoh, N. Agmon, B. An (eds.), May 29 – June 2, 2023, London, United Kingdom. © 2023 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

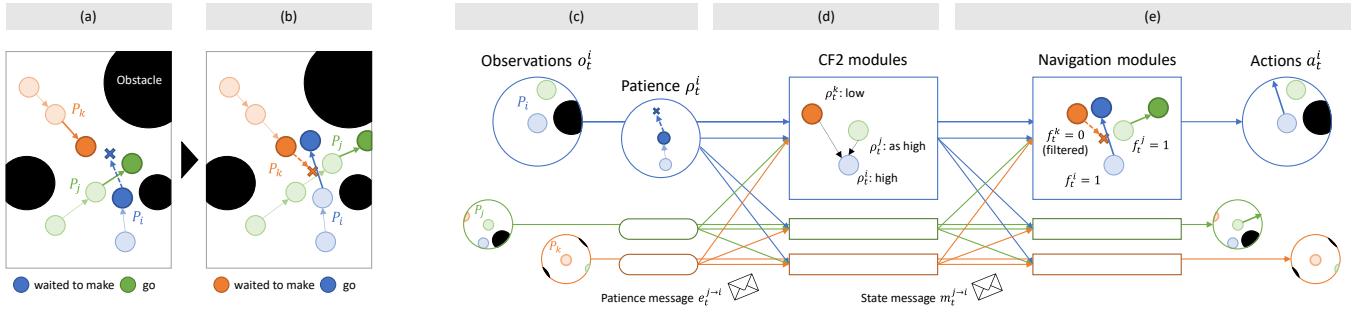


Figure 1: Example of fair-delay multi-robot navigation by NCF2. (a, b) Three agents perform a navigation task while giving way to equalize their travel delays. (c) To do so, the agents first observe the situations around their current positions and share their “patience,” which represents how much delay they have experienced so far. (d) The CF2 module processes the patience information to decide and share whether each agent should move. (e) The agents take their next actions according to the navigation module.

most counterfactual methods require centralized training with a global critic evaluating the actions of all agents [11, 26, 37], which is not feasible as the number of agents increases.

To address this challenge, we propose a decentralized, counterfactual, multi-robot navigation algorithm called **Navigation with Counterfactual Fairness Filter (NCF2)**. The key idea is to equip each agent with a policy comprising a *navigation module* and a *counterfactual fairness filter (CF2) module*. While the navigation module enables agents to take continuous actions to reach their destinations efficiently and safely, the CF2 module enables them to perform counterfactual inference on action choices to improve fairness. Fig. 1 illustrates a typical case targeted by NCF2. In Fig. 1b, agent P_i wants to advance according to its navigation module, while agent P_k is trying to cut in along the way. As shown in Fig. 1a, P_i 's travel already involved a delay to make room for another agent P_j . Here, each agent's CF2 module observes the local situations around the agent and asks P_k to suspend its next action and let P_i go ahead. By repeatedly resolving such small problems in every local situation, we can expect to eventually achieve fairness across all agents.

In NCF2, agents learn a shared policy in a decentralized fashion by communicating with each other as follows. At each timestep, each agent partially observes the environment around its current position and creates messages about its ‘patience’. This notion of patience indicates how much delay an agent has experienced so far in comparison to the best possible actions it could have taken (Fig. 1c). An agent's CF2 module receives the patience information and decides whether agents should move according to the difference in the patience scores between each agent and its surroundings (Fig. 1d). The CF2 module outputs are again shared across agents, and the navigation module determines the actual next action while considering which agents are allowed to move (Fig. 1e). As a result of the action, each agent receives a reward for efficiency, safety, and the counterfactual measurement of fairness. The experience from this procedure is accumulated over all agents and steps, and it is used to learn the policy to maximize the cumulative rewards via a reinforcement learning (RL) algorithm.

To evaluate the effectiveness of our approach, we constructed a set of challenging, fair-delay, multi-agent navigation environments

involving multiple wheeled robotic agents with a simple but practical kinematics model and a continuous action space. Our results show that the policy learned with NCF2 outperforms state-of-the-art fairness-aware MARL methods [22, 48] in terms of the success rate, travel time, and degree of fairness.

2 PROBLEM STATEMENT

This section introduces the proposed fair-delay multi-robot navigation problem step by step. We address a multi-agent setting with a swarm of N robotic agents, P_1, \dots, P_N , to simulate autonomous mobile robots in a 2D environment with some static obstacles. We begin by explaining a simpler case with only a single agent P_i and then extend it to the fair-delay multi-robot navigation problem.

Single-Robot Navigation: This problem is formulated as a tuple $\Pi_i = (\mathcal{S}_i, \mathcal{A}_i, \mathcal{O}_i, I_i, T_i, R_i, \Omega_i, t_{\max})$, where \mathcal{S}_i , \mathcal{A}_i , and \mathcal{O}_i are the (individual) state, action, and observation spaces, respectively.¹ Specifically, $\mathcal{S}_i = \text{SE}(2) = \mathbb{R}^2 \times \mathbb{S}^1$, where $\text{SE}(2)$ is the special Euclidean group whose elements represent the robot's position and orientation. It comprises a set of invalid states, I_i , and a set of valid states, \mathcal{V}_i . Invalid states are those where the agent is too close to an obstacle and is considered to be crashing into it, whereas all other states are valid.² At the beginning of an “episode,” a start state s_1^i and a goal region center $c_i \in \mathbb{R}^2$ are sampled from a probability distribution I_i over $\mathcal{V}_i \times \mathbb{R}^2$. A ball (whose radius is environment-dependent) around c_i is defined as the goal region, and valid states whose position belongs to the goal region are called goal states. The set of all goal states is denoted by $g_i \subset \mathcal{V}_i$. The agent's aim is to reach a goal state without entering I_i . At each timestep t , the agent receives partial information $o_t^i = \Omega_i(s_t^i) \in \mathcal{O}_i$ about the environment around its current location, where s_t^i is the current state. Then, it chooses and executes an action $a_t^i \in \mathcal{A}_i \subset \mathbb{R}^2$, which entails the robot's velocity and angular velocity. As a consequence, the state s_t^i transitions to $s_{t+1}^i = T_i(s_t^i, a_t^i, c_i)$, and the agent receives a reward $r_t^i = R_i(s_t^i, a_t^i, s_{t+1}^i, c_i)$. When the agent reaches an invalid or goal state, its next state is the same regardless of any action, i.e.,

¹We omit “individual” when it is clear from context.

²The definition of a “crash” depends on the environment. All environment-specific details, such as the definitions of a crash and the goal region, are given in Sec. 4.1.

$T_i(s_t^i, a_t^i, c_i) = s_{t+1}^i$. Each episode is reset at $t+1 = t_{\max}$. If a trajectory, $\tau_i = (s_t^i)_{t=1}^{t_{\max}}$, ends at a goal state, the trajectory is considered valid. The length of a valid trajectory is defined as $l_i = \min\{t | s_t^i \in g_i\}$, which indicates the travel time. For simplicity, we assume a *homogeneous* setup where $\Pi_1 = \dots \Pi_N = \Pi'$. Because this formulation can be understood as a goal-conditional version [2] of a partially observable Markov decision process (POMDP) [24], we refer to it as an individual POMDP (for P_i).

Multi-Robot Navigation: This problem is formulated as a tuple $\Pi = (N, \mathcal{S}, \mathcal{A}, O, I, T, R, \Omega, t_{\max})$, where $\mathcal{S} = \times_{i=1}^N \mathcal{S}_i$, $\mathcal{A} = \times_{i=1}^N \mathcal{A}_i$, and $O = \times_{i=1}^N \mathcal{O}_i$ are the (joint) state, action, and observation spaces, respectively.^{3 4} While the state space can be partitioned as before into a set of invalid states, \mathcal{I} , and a set of valid states, \mathcal{V} , \mathcal{I} now also includes states where an agent is crashing into another agent. An initial state $s_1 = (s_1^i)_{i=1}^N \in \mathcal{V}$, a goal center $c = (c_i)_{i=1}^N \in (\mathbb{R}^2)^N$, and a goal state $g = (g_i)_{i=1}^N \in \mathcal{V}$ are determined as before.

⁵ The probability distribution I of (s_1, c) is simply chosen to be consistent with this construction. The aim of the agents in the swarm is to reach g without entering \mathcal{I} . At each timestep t , the agents receive an observation $o_t = (o_t^i)_{i=1}^N = \Omega(s_t) \in \mathcal{O}$, where $s_t = (s_t^i)_{i=1}^N \in \mathcal{S}$ is the current state. Then, the agents exchange messages, possibly multiple times, to cooperate with each other. Each message from P_j to P_i at timestep t is given by a D -dimensional real-valued vector $m_t^{j \rightarrow i} \in \mathbb{R}^D$. After message exchanges, the agents choose and execute an action $a_t = (a_t^i)_{i=1}^N \in \mathcal{A}$. As a consequence, the state s_t transitions to $s_{t+1} = (s_{t+1}^i)_{i=1}^N = T(s_t, a_t, c)$, and the agents receive a global reward $r_t = R(s_t, a_t, s_{t+1}, c)$. We assume that $s_{t+1}^i = T_i(s_t^i, a_t^i)$, unless an agent (say P_i) is crashing into another agent, in which case $s_{t+1}^i = s_t^i$. Each episode is reset at $t+1 = t_{\max}$. A trajectory $\tau = (s_t)_{t=1}^{t_{\max}}$ is valid if it ends at a goal state, *i.e.*, if all agents arrived at their own goal. The length of a valid trajectory is given by $l = \min\{t | s_t \in g\} = \max\{\min\{t_i | s_t^i \in g_i\}\} = \max\{l_i\}$, where l_i here is the travel time of P_i . Note that this trajectory length is referred to as a *makespan* in the context of MAPF [36]. This formulation can be viewed as a goal-conditional version of a decentralized POMDP [14], hereafter simply called a Dec-POMDP.

Cooperative Policy: Given an observation $o_t^i \in \Omega_i$ and a collection of messages from neighboring agents, $\{m_t^{j \rightarrow i} | j \in \mathcal{N}_t^i\}$, where the $\mathcal{N}_t^i \subset \{1, \dots, N\}$ are the indices of agents regarded as P_i 's neighbors at time t , each agent executes an action a_t^i sampled from a shared *cooperative policy* $\pi(\cdot | o_t^i, c_i, \{m_t^{j \rightarrow i} | j \in \mathcal{N}_t^i\})$ to move one step toward the goal g_i . If π is trained properly and is running for an episode with a start state $s_1 = (s_1^i)_{i=1}^N$ and a goal state $g = (g_i)_{i=1}^N$, then a valid trajectory τ^π with length $l^\pi = \max\{l_i^\pi\}$ is obtained.

Delay Measurement with Solitary Policy: For a valid trajectory τ^π , each agent's goal time l_i^π can be delayed in comparison to the case when agents do not have to avoid collisions with each other. To quantify this delay, we assume that another deterministic policy, the *solitary policy* $\mu : \mathcal{O}_i \times c_i \rightarrow \mathcal{A}_i$, is pre-trained for an individual

POMDP Π_i to ensure safe, efficient actions in the absence of other agents. By considering the interaction of a single agent P_i with the multi-robot environment Π and the solitary policy μ while freezing the other agents, we obtain the *the best possible trajectory* τ_i^μ , with length l_i^μ , that P_i could have taken. (The start and goal states here are s_1^i and g_i , respectively.) Then, the delay for agent P_i with the cooperative policy π can be calculated as $\delta_i = l_i^\pi - l_i^\mu$.

Fair-Delay Multi-Robot Navigation Problem: For a problem instance modeled by a Dec-POMDP Π , the navigation task is regarded as *successful* if we can find a valid cooperative trajectory τ^π . By contrast, navigation fails if any agent fails to reach its goal within t_{\max} or collides with a static obstacle, or if any two agents collide with each other. For a successful task, we define the degree of (un)fairness in terms of the *the variance of delays*, *i.e.*, $\frac{1}{N} \sum_i (\delta_i - \frac{1}{N} \sum_j \delta_j)^2$: the smaller this variance is, the fairer the solution is. We also measure the travel (in)efficiency in terms of the valid trajectory length l^π (*i.e.*, the makespan), where a shorter length is better.

The goal of the proposed *fair-delay multi-robot navigation* is to find a cooperative policy that is successful with the highest probability possible (*i.e.*, the policy achieves a high safety level) while also making the unfairness and inefficiency scores as low as possible for the successful instances.

3 NCF2: NAVIGATION WITH COUNTERFACTUAL FAIRNESS FILTER

We construct the proposed cooperative policy, hereafter referred to as the *NCF2 policy*, via two modules. The *navigation module* produces agent actions in a continuous space for travel efficiency and safety, while the *counterfactual fairness filter (CF2) module* judges whether agents should move by considering fairness and efficiency. The CF2 module's action is restricted to be binary, which makes it easy to incorporate counterfactual reasoning while allowing agents to move with the continuous action space. Below, we first describe how to generate fairness-aware actions with the NCF2 policy in Sec. 3.1; we then introduce reward designs to learn the policy in a decentralized fashion via an RL algorithm in Sec. 3.2.

3.1 Generation of Fairness-Aware Actions

The NCF2 policy generates actions in the following three steps, with multiple message exchanges between agents (see also Fig. 1).

Step 1: Calculate Patience as Proxy for Delays (Fig. 1c): The first step is to measure the delay $\delta_i = l_i^\pi - l_i^\mu$ as defined in Sec. 2. However, the exact value is unavailable during task execution because the actual trajectory length l_i^π is unknown until each agent arrives at the goal. Instead, we introduce a proxy measurement, referred to as *patience*, which indicates how patient an agent has been in being cooperative (*i.e.*, moving by communicating with others) rather than solitary (*i.e.*, moving while ignoring others). We calculate the patience as follows by using the action-value function of the solitary policy μ on an individual POMDP Π_i , *i.e.*, Q_μ :

$$\rho_t^i = \sum_{t'=1}^{t-1} \left(Q_\mu(o_{t'}^i, \hat{a}_{t'}^i, c_i) - Q_\mu(o_{t'}^i, a_{t'}^i, c_i) \right), \quad (1)$$

where $\hat{a}_{t'}^i = \mu(o_{t'}^i, c_i)$ and $a_{t'}^i \sim \pi(\cdot | o_{t'}^i, c_i, \{m_t^{j \rightarrow i} | j \in \mathcal{N}_{t'}^i\})$ are the actions produced by the solitary and cooperative policies,

³We omit "joint" when it is clear from context.

⁴Without loss of generality, we assume that each \mathcal{O}_i contains not only observations in the single-agent case but also the multi-agent case considered here. Indeed, an agent's observation is the same in both cases when there are no neighboring agents.

⁵To avoid pathological cases, we assume that the goal centers are sufficiently separated from each other.

respectively. Because μ is trained to enable efficient navigation in a single-agent environment, as assumed in the previous section, its action-value function evaluates how promising input actions are from this perspective. As a small $Q_\mu(o_t^i, a_t^i, c_i)$ indicates that action a_t^i will not move the agent efficiently toward the goal, a large ρ_t^i implies that the delay will be as large.

Step 2: Process Patience with CF2 Module (Fig. 1d): The patience is included in a message that is shared among the agents and used in the CF2 module to judge whether each agent should really move. We define this message, hereafter referred to as a *patience message*, as follows:

$$e_t^{j \rightarrow i} = \left(\Delta(s_t^j, s_t^i); \frac{\rho_t^j - \rho_t^i}{\sum_{j \in \mathcal{N}_t^i} \rho_t^j}; \Delta(\hat{s}_{t+1}^j, s_t^i); \frac{\rho_t^j - \rho_t^i}{\sum_{j \in \mathcal{N}_t^i} \rho_t^j} \right). \quad (2)$$

Here, $\Delta(a, b)$; $a, b \in \mathcal{S}_i$ is the relative position of state a with respect to b , \hat{s}_{t+1}^i is the next state sampled from an individual (*i.e.*, single-agent) POMDP Π^i with action \hat{a}_t^i , and $(u; v)$ is a vector obtained by vertically concatenating column vectors u and v . Intuitively, a patience message expresses where each agent is and where it would go if it were allowed to ignore other agents, given the current relative patience. We model the CF2 module as a sub-policy $\psi(\cdot \mid o_t^i, c_i, \{e_t^{j \rightarrow i} \mid j \in \mathcal{N}_t^i\})$, which provides a binary decision $f_t^i \in \{0, 1\}$ based on the current observation and patience messages. Once learned properly, it gives $f_t^i = 1$ to allow agents that have been sufficiently patient to move, and it acts like a message filter in the next step.

Step 3: Decide Next Actions (Fig. 1e): Given the output f_t^i from the CF2 module, we define a *state message* as follows:

$$m_t^{j \rightarrow i} = \left(f_t^j \Delta(s_t^j, s_t^i); f_t^j \Delta(\hat{s}_{t+1}^j, s_t^i) \right), \quad (3)$$

which represents the relative state of agent P_j with respect to P_i if P_j is allowed to move (*i.e.*, $f_t^j = 1$) or contains all-zeros otherwise. This message is again exchanged among the agents, and used to generate each agent's the next action. The navigation module ϕ is modeled as another sub-policy to sample a continuous action a_t^i from an observation o_t^i and state messages $\{m_t^{j \rightarrow i} \mid j \in \mathcal{N}_t^i\}$. The sampling is filtered by f_t^i to stop agents that are not allowed to move:

$$a_t^i = f_t^i \cdot v, \quad v \sim \phi \left(\cdot \mid o_t^i, c_i, \{m_t^{j \rightarrow i} \mid j \in \mathcal{N}_t^i\} \right). \quad (4)$$

Encoding Messages: In the above action generation procedure, collections of messages $\{e_t^{j \rightarrow i} \mid j \in \mathcal{N}_t^i\}$ and $\{m_t^{j \rightarrow i} \mid j \in \mathcal{N}_t^i\}$ are each encoded into a single fixed-length vector representation via the scaled dot-product attention mechanism [41] (see Appendix B.1 for more details). This is an effective approach for handling sets with variable numbers of messages in multi-agent modeling [8].

3.2 Decentralized Learning of NCF2 Policy

The NCF2 policy is learned in a decentralized fashion to ensure scalability for the number of agents. Specifically, to improve the fairness regarding delays while maintaining navigation efficiency and safety, we design a local reward r_t^i that is given to each agent, instead of explicitly designing a global reward function R for the Dec-POMDP Π .

Crucially, the navigation and CF2 modules each act as sub-policies that pursue different objectives, and thus, they should be evaluated differently. To this end, we use the hybrid reward architecture [40] and define the local reward r_t^i as the sum of two rewards: an *efficiency-safety reward* and a *fairness-efficiency reward*. As explained below, they are given *separately* to the two modules.

Efficiency-Safety Reward: The navigation module is trained with the following reward:

$$\tilde{r}_t^i = \begin{cases} r_{\text{goal}} - r_{\text{time}} & \text{upon reaching the goal} \\ -r_{\text{crash}} - r_{\text{time}} & \text{if crashing into an obstacle or another agent} \\ 0 & \text{after reaching the goal or crashing} \\ -r_{\text{time}} & \text{otherwise.} \end{cases} \quad (5)$$

Here, $r_{\text{goal}} > 0$ and $r_{\text{crash}} > 0$ are respectively the goal reward and crash penalty given sparsely, and $r_{\text{time}} > 0$ is a time penalty to encourage agents to reach the goal in the shortest possible time.

Fairness-Efficiency Reward: On the other hand, the CF2 module is given a counterfactual reward to improve fairness while maintaining efficiency. We consider the most cooperative action that could happen if all agents were allowed to move by their CF2 modules as a *default action*, *i.e.*, a replacement for the planned action:

$$\bar{a}_t^i \sim \phi \left(\cdot \mid o_t^i, c_i, \left\{ \left(\Delta(s_t^j, s_t^i), \Delta(\hat{s}_{t+1}^j, s_t^i) \right) \mid j \in \mathcal{N}_t^i \right\} \right). \quad (6)$$

This enables us to define the *improvement* for an actual action a_t^i by applying the action-value function of the solitary policy Q_μ to an individual POMDP Π_i as follows:

$$\xi_t^i = Q_\mu(o_t^i, a_t^i, c_i) - Q_\mu(o_t^i, \bar{a}_t^i, c_i). \quad (7)$$

Similarly to how the patience ρ_t^i is characterized by Q_μ in Eq. (1), this improvement ξ_t^i indicates how much better the actual action a_t^i has become compared to the default action, by preventing some agents from moving via $m_t^{j \rightarrow i}$ in Eq. (4). It is then aggregated across neighboring agents, while being weighted by the relative patience, $\rho_t^j - \rho_t^i$, to form a sparse *fairness-efficiency* reward that is given when a CF2 module takes $f_t^i = 0$:

$$\tilde{r}_t^i = (1 - f_t^i) \left(\alpha \frac{\sum_{j \in \mathcal{N}_t^i} (\rho_t^j - \rho_t^i) \cdot \xi_t^j}{\sum_{j \in \mathcal{N}_t^i} \rho_t^j} - \beta \frac{\rho_t^i}{\sum_{j \in \mathcal{N}_t^i} \rho_t^j} \right). \quad (8)$$

Intuitively, this reward becomes higher for agents that decide to stop when their surrounding agents have been more patient and also can improve their action values by moving. The second term is a penalty to prevent excessive encouragement of agents to stop and sacrifice efficiency. Here, α, β are user-specified constants to control the balance between fairness and efficiency. Further analysis of this reward design is presented in Appendix A.

4 EXPERIMENTS

4.1 Environments

Our work targets multi-robot navigation tasks in which robotic agents with practical kinematics move in a 2D environment with some static obstacles. This kind of task requires a new multi-agent environment that is completely different from previous ones that only support discrete movements on a grid map, as what was used

in [22]. To this end, we extend a multi-agent path planning environment [31, 45] with crowd-aware single-robot navigation [5] to construct environments with the following features.

Agent Design: Each agent simulates a wheeled robot equipped with a simulated lidar sensor. The sensor perceives the distances to surrounding obstacles and other agents at equal-angle intervals with a maximum scan range of $0.1 \times \text{map-size}$ to obtain 64-dimensional scan data. The robot is modeled as a circle of fixed size ($0.02 \times \text{map-size}$), and we determine that a collision occurs if an obstacle or another agent enters the circle. The agent can exchange messages with other agents within a range of $0.15 \times \text{map-size}$. Finally, an agent’s motion follows a non-holonomic kinematics model that is characterized by a maximum velocity ($0.05 \times \text{map-size}/\text{timestep}$) and a maximum angular velocity ($0.25\pi/\text{timestep}$). The agent accepts continuous-valued velocity and angular velocity commands to move forward.

Environment Design: The environment is a map of size 128×128 (i.e., $\text{map-size} = 128$) and includes static circular obstacles. In our experiments, the goal radius was set to $0.02 \times \text{map-size}$, and the time limit was set to $t_{\max} = 100$. As shown in Fig. 2, we constructed the following four environments:

- **UNIFORM- N -25** comprised $N \in \{8, 12, 16\}$ agents with start and goal positions sampled uniformly in the environment. It included 25 circular obstacles of random size ($[0.05, 0.08] \times \text{map-size}$) and location.
- **CORNER- N -25** also comprised N agents and 25 random circular obstacles, but the agent’s start positions were sampled randomly from the four corner regions, and the goal positions were diagonally opposite the starts. This environment requires agents to give way to each other in the center, thus giving a more challenging case than UNIFORM- N -25.
- **UNIFORM- N -50** and **CORNER- N -50** increased the number of random obstacles from 25 to 50, making the navigation task more difficult.

4.2 Implementation Overview

Because of space limitations, we give only a brief overview of the implementation here. Further details are described in Appendix B.

Navigation Module: We constructed the navigation module by extending a state-of-the-art RL-based navigation algorithm called residual reactive navigation [32], which is an application of residual RL [23, 35] to navigation tasks. Specifically, we used the classical dynamic window approach (DWA) [12] as a base controller, and we then applied learning of a residual policy given by a multi-layer perceptron (MLP), which added residual actions to DWA’s outputs. An observation of this module, o_t^i , comprised the current state, a lidar scan as described in Sec. 4.1, the relative displacement of the goal position with respect to the current position, and the DWA output.

CF2 Module: The CF2 module was also modeled by a standard MLP but had a binary output given by the softmax activation. An observation was the same as for the navigation module.

Solitary Policy: The solitary policy had the same form as that of the navigation module, but without taking messages from other agents as input.

Reward: The reward and penalty values for the safety-efficiency reward in Eq. (5) were $r_{\text{goal}} = 3.0$, $r_{\text{crash}} = 10.0$, and $r_{\text{time}} = 0.1$, such that crashes were penalized more than timeouts. The constants α , β in the fairness-efficiency reward in Eq. (8) were set empirically to $\alpha = 0.5$, $\beta = 0.1$, where α was selected from $\alpha \in [0.1, 5]$ in our pilot experiment.

Learning Algorithm: We used soft actor-critic (SAC) [16, 17] as the RL algorithm to achieve high sample efficiency and learning stability. Crucially, our implementation adopted the distributed experience-generating strategy of Ape-X [19] to achieve high data throughput, which we found effective for alleviating the non-stationarity of gathered experiences in off-policy MARL. To learn the NCF2 policy, we first warmed up the navigation module by itself for 1M iterations and then conducted learning by both modules for 1M iterations. This allowed the CF2 module to make a valid decision to pursue fairness by leveraging the safety and efficiency provided by the trained navigation module. The solitary policy was trained for 1M iterations by limiting the number of agents to one in the learning environment. All experiments were conducted with a single GPU (NVIDIA V100) and a single CPU (Intel Xeon Gold 6252).

4.3 Baselines

We used the following state-of-the-art fairness-aware MARL methods as baselines.

- **Fair-Efficient Network (FEN)** [22] is a hierarchical MARL algorithm that simultaneously learns fairness and efficiency. FEN comprises a controller and multiple sub-policies, where the controller first selects a sub-policy and then the selected sub-policy outputs an action. While one of the sub-policies is learned with an efficiency-safety reward \hat{r}_t^i as defined in Eq. (5), the remaining sub-policies receive the log-probability of being selected by the controller as a reward. Furthermore, the average reward over the elapsed timesteps is calculated as a utility, and to pursue fairness, the controller is learned to minimize the difference between its own utility and the average utilities of all agents.
- **Self-Oriented Team-Oriented Network (SOTO)** [48] is another hierarchical algorithm, which seeks to maximize the generalized Gini social welfare function (GGF) to simultaneously quantify efficiency and fairness. SOTO learns two policies: a self-oriented policy and a team-oriented policy. The self-oriented policy first samples an action, which is then used as an input to the team-oriented policy to sample a cooperative action. While the self-oriented policy is learned from the efficiency-safety reward, the team-oriented policy receives as a reward comprising the weighted sum of the other agent’s cumulative rewards to maximize the GGF. Here, the weights are determined such that agents with larger cumulative rewards are given smaller values.

Note that in our initial study, we realized that direct adaptations of the above methods did not solve our navigation tasks at all. To mitigate this problem, we extended these methods by (1) replacing the base RL algorithm from proximal policy optimization (PPO) [33] with our SAC implementation to achieve high data throughput; (2) limiting each agent’s communication to a few agents in its vicinity; and (3) using the warmed-up navigation module as one of these




Figure 2: Qualitative comparisons of the methods and environments (best viewed in color). Start positions and goal centers are represented by filled and white circles, respectively, with different colors for different agents. The histories of the agent positions are indicated by small circles connected by line segments. For the case of NCF2, moments when the CF2 module decided to stop an agent are annotated by gray circles. Collisions and timeouts are shown by gray crosses and triangles respectively. The black circles are obstacles.

methods’ sub-policies. For another baseline, we used a degraded version of our approach, referred to as **Nav-Only**, which used only the trained navigation module to control the agents. To make the experiment fair, the hyperparameters of the RL algorithm were identical to those of the proposed method. The number of training iterations was doubled for Nav-Only to compensate for the warm-ups of the other methods.

4.4 Evaluation Metrics

For each method, a cooperative policy was learned five times with different random seeds, and we evaluated each learned policy on a set of 100 new episodes. The following metrics were calculated and averaged over all episodes to measure the fairness, efficiency, and safety of multi-robot navigation by each method.

- **Success Rate (SR)** measures the proportion of successful evaluation episodes, which indicates a method’s safety.
- **Makespan (MS)** is the length of a valid trajectory l^π for successful episodes, which indicates a method’s efficiency.
- **Variance of Delays (VD)** is our fairness indicator that was introduced in Sec. 2. It is calculated for each successful episode as the variance of the differences between the lengths of a cooperative trajectory l_i^π and the best possible trajectory τ_i^H obtained by the solitary policy.
- **Maximum Delay (MAXD) and Mean Delay (MEAND)** are the maximum and mean delay for each successful episode, thus giving another set of efficiency metrics.

4.5 Results

Quantitative Results: Tab. 1 summarizes the quantitative comparison results. Overall, the proposed NCF2 performed well on every metric. In contrast, for challenging environments with 50 obstacles and more than 12 agents, it was hard or almost impossible for the baseline methods to successfully complete the navigation tasks. For SOTO, this was possibly because it is difficult to resolve a credit assignment problem by learning a team-oriented sub-policy just to balance the cumulative rewards among neighboring agents. FEN could perform better than SOTO and was sometimes comparable to or better than NCF2 in terms of the MS, VD, MAXD, and MEAND scores. Nevertheless, its SR scores were quite limited as the numbers of agents and obstacles both increased; that is, fair-delay navigation was possible only in a limited number of successful episodes. For FEN, we also observed many situations in which only the efficiency-oriented sub-policy was selected by the controller, rather than other sub-policies that would have allowed agents to give way to each other, thus limiting the SR. Nav-Only, the degraded version of NCF2, demonstrated comparable success rates in environments that were not very difficult (e.g., UNIFORM-8-25, UNIFORM-12-25, UNIFORM-16-25, CORNER-8-25, CORNER-16-25). Nevertheless, its VD scores were worse than those of NCF2, especially in CORNER environments, because of the lack of a mechanism to maintain fairness regarding delays. Interestingly, in difficult environments with 50 obstacles, NCF2 could succeed in more episodes than Nav-Only, apart from the fairness. This was because the NCF2 policy gave agents a wider choice of actions, including “staying still,” and allowed the agents to consider the actions of selected agents who will move in the next step. On the other hand, the Nav-Only agents had to consider the actions of all neighboring agents. This limited the options for the Nav-Only agents’ next actions and could result in collisions or timeouts in congested environments.

Qualitative Results: Fig. 2 shows some visual examples of solutions by each method. For NCF2, agents could stop properly when another agent was moving just ahead of them, as indicated by gray circles on their trajectories. On the other hand, agents tended to get stuck (gray triangles) or even collided with other obstacles (gray crosses) with the other baselines.

Ablation Study: We additionally evaluated the following degraded versions of NCF2 in the CORNER-12-25 environment to demonstrate the importance of each component: (i) **w/o improvements**, which used $\xi_t^j = 1$ for all j and t in Eq. (8) to disable the counterfactual reasoning; (ii) **w/o limited communications**,

which used $\mathcal{N}_t^i = \{1, \dots, N\}$ for all i and t to examine how allowing message exchanges beyond an agent’s proximity affected the performance; and (iii) **fixed priority**, to examine how the performance changed if the patience was given deterministically as $\rho_t^i = i - 1$ rather than by considering an agent’s actual patience. As listed in Tab. 2, each of these versions showed decreased performance as compared to the proposed method. While the MEAND (mean of delays) did not differ among the methods, the VD (variance of delays) showed a considerable gap. This result indicates that the counterfactual reasoning, limited communication, and use of patience scores for improved fairness regarding delays were all important.

Failure Cases and Limitations: Although NCF2 worked well in many cases, our work has several limitations. We found that the NCF2 policy sometimes failed to resolve situations when agents were crowded into a small area. This was because the actions to avoid collisions are inherently limited in highly congested environments, thus making it difficult to expand other agents’ action choices simply by taking cooperative actions and improving the fairness. A possible way to mitigate such problems in congested environments would be to incorporate global planning that searches the environment for low-cost, feasible paths. This approach can address the myopia of local planning and make it possible to actively bypass congested areas. Finally, while we assumed the state transitions and agent observations to be deterministic, they are stochastic in practice because of various noise sources [38]. Extensions of our work to stochastic environments would be an interesting future direction.

5 RELATED WORK

Fairness has been actively studied in MARL. Typical tasks include job scheduling, the Matthew effect, traffic control, and so on in a cooperative setting [22, 34, 48], and fair resource allocation [7, 9, 47] and social dilemmas [20, 21, 25] in a mixed setting. More recently, the application of prediction-based fairness [4, 6] was studied in a multi-agent setting [15]. Two prior studies that provided baselines in our experiments [22, 48] especially addressed the tradeoff between fairness and efficiency. Nevertheless, these methods have limitations in resolving complex credit assignment problems to improve fairness while maintaining navigation safety and efficiency, as demonstrated in the experiments.

On the other hand, the notion of counterfactuals has recently attracted much attention as a powerful tool to address complex multi-agent credit assignment problems. Since a pioneering work [11] that used counterfactuals to compute the baseline for the actor-critic algorithm, more recent works have extended counterfactuals with a graph convolution communication [37], Shapley counterfactual credits [26], and simulation reasoning for fully-decentralized counterfactual MARL [46]. However, all these methods have been found effective only in discrete action spaces. For some methods, the specification of default actions is nontrivial in a continuous space [26, 46]. Moreover, all but the last one [46] require a centralized critic and have limited scalability with respect to the number of agents.

Finally, multi-robot navigation is a popular task in the robotics community. It has been addressed analytically [1, 3, 39], via MAPF [18, 27, 30], or via machine learning including MARL [10, 13, 28, 43], to name a few approaches. In any case, the main challenges in these studies have been successful collision avoidance for safety

Table 1: Quantitative results. SR: success rate (\uparrow); MS: makespan (\downarrow); VD: variance of delays (\downarrow); MAXD: maximum delay (\downarrow); MEAND: mean delay (\downarrow). Each SR score is the average over five trials for 100 episodes with different random seeds, while the other scores are the averages over the successful episodes among them.

	UNIFORM-8-25					UNIFORM-12-25					UNIFORM-16-25				
	SR	MS	VD	MAXD	MEAND	SR	MS	VD	MAXD	MEAND	SR	MS	VD	MAXD	MEAND
FEN [22]	95.4	22.2	10.5	5.11	1.26	37.0	27.2	15.8	8.77	2.46	17.0	32.0	25.5	13.1	3.42
SOTO [48]	17.4	50.5	185	33	9.36	0.80	44.2	78.5	27.8	9.31	0	N/A	N/A	N/A	N/A
Nav-Only	98.6	22.3	13.3	4.85	1.26	93.0	26.1	13.1	7.52	1.74	75.8	30.9	20.9	11.5	2.86
NCF2 (Ours)	98.0	22.2	5.63	4.51	1.25	92.6	26.7	10.0	7.29	1.89	77.8	32.1	22.0	12.3	3.02
	UNIFORM-8-50					UNIFORM-12-50					UNIFORM-16-50				
	SR	MS	VD	MAXD	MEAND	SR	MS	VD	MAXD	MEAND	SR	MS	VD	MAXD	MEAND
FEN [22]	63.4	31.6	28.9	11.3	3.18	26.4	40.1	68.4	18.8	4.79	0	N/A	N/A	N/A	N/A
SOTO [48]	3.20	78.1	500	60.5	23.1	0	N/A	N/A	N/A	N/A	0	N/A	N/A	N/A	N/A
Nav-Only	89.4	29.6	21.1	8.64	2.40	65.4	38.0	54.6	16.9	4.37	30.8	50.1	101	28.1	6.39
NCF2 (Ours)	91.7	30.3	24.1	9.34	2.64	71.4	38.2	47.0	17.2	4.36	40.8	48.3	85.8	26.1	6.55
	CORNER-8-25					CORNER-12-25					CORNER-16-25				
	SR	MS	VD	MAXD	MEAND	SR	MS	VD	MAXD	MEAND	SR	MS	VD	MAXD	MEAND
FEN [22]	73.8	34.8	37.8	14.9	5.84	29.6	48.4	73.9	28.0	15.6	8.80	77.6	283	57.7	35.3
SOTO [48]	17.2	63	140	42.8	24.4	0	N/A	N/A	N/A	N/A	0	N/A	N/A	N/A	N/A
Nav-Only	91.4	42.2	85.5	23.3	14.5	77.6	58.9	133.9	39.6	25.2	52.0	64.4	170	44.3	24.5
NCF2 (Ours)	93.8	43.7	85.0	24.8	15.7	86.4	50.7	74.6	31.0	19.5	49.6	55.5	101	35.1	18.5
	CORNER-8-50					CORNER-12-50					CORNER-16-50				
	SR	MS	VD	MAXD	MEAND	SR	MS	VD	MAXD	MEAND	SR	MS	VD	MAXD	MEAND
FEN [22]	44.4	45.3	64.2	21.8	10.2	8.00	79.9	359	59.2	30.7	0	N/A	N/A	N/A	N/A
SOTO [48]	0	N/A	N/A	N/A	N/A	0	N/A	N/A	N/A	N/A	0	N/A	N/A	N/A	N/A
Nav-Only	54.4	68.0	285	46.8	29.4	20.8	78.2	323	57.1	33.2	4.00	82.6	325	61.3	33.0
NCF2 (Ours)	70.6	66.7	271	45.5	28.1	40.4	76.3	288	54.6	32.6	9.20	79.1	273	56.6	31.1

Table 2: Ablation study. (i) w/o improvements: $\xi_t^j = 1$ to disable counterfactual reasoning. (ii) w/o limited comm.: $\mathcal{N}_t^i = \{1, \dots, N\}$ to enforce agents to exchange messages with all other agents. (iii) fixed priority: $\rho_t^i = i - 1$ to avoid using patience scores.

	SR	MS	VD	MAXD	MEAND
(i) w/o improvements	84.0	58.0	99.5	36.8	20.4
(ii) w/o limited comm.	79.2	55.3	113	34.9	19.6
(iii) fixed priority	62.8	50.3	86.4	30.9	16.4
NCF2	86.4	50.7	74.6	31.0	19.5

and improvement in the trade-off between safety and efficiency, without considering how to achieve fairness as well.

6 CONCLUSION

We have introduced fair-delay multi-robot navigation, a novel extension of fairness-aware MARL, with the objective of finding a set of trajectories for a team of agents. The trajectories should be collision-free, short, and as fair as possible in terms of temporal delays as compared to the respective shortest possible trajectories

when ignoring the presence of other agents. To address this problem, we proposed a new navigation algorithm, called NCF2, that enables counterfactual inference on whether each agent should move or stay still to improve the fairness regarding delays while maintaining travel efficiency and safety. In terms of those criteria, the proposed method outperformed state-of-the-art fairness-aware MARL methods in challenging multi-robot navigation environments.

REFERENCES

- [1] Javier Alonso-Mora, Andreas Breitenmoser, Martin Rufli, Paul Beardsley, and Roland Siegwart. 2013. Optimal reciprocal collision avoidance for multiple non-holonomic robots. In *Distributed autonomous robotic systems*. Springer, 203–216.
- [2] Marcin Andrychowicz, Filip Wolski, Alex Ray, Jonas Schneider, Rachel Fong, Peter Welinder, Bob McGrew, Josh Tobin, OpenAI Pieter Abbeel, and Wojciech Zaremba. 2017. Hindsight experience replay. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*. 5055–5065.
- [3] Daman Bareiss and Jur Van den Berg. 2015. Generalized reciprocal collision avoidance. *International Journal of Robotics Research (IJRR)* 34, 12 (2015), 1501–1514.
- [4] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the Conference on Fairness, Accountability and Transparency*. 77–91.
- [5] Changan Chen, Yuejiang Liu, Sven Kreiss, and Alexandre Alahi. 2019. Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 6015–6022.
- [6] Sam Corbett-Davies and Sharad Goel. 2018. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv preprint arXiv:1808.00023*

- (2018).
- [7] Alexander D’Amour, Hansa Srinivasan, James Atwood, Pallavi Baljekar, David Sculley, and Yoni Halpern. 2020. Fairness is not static: deeper understanding of long term fairness via simulation studies. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 525–534.
 - [8] Abhishek Das, Théophile Gervet, Joshua Romoff, Dhruv Batra, Devi Parikh, Mike Rabbat, and Joelle Pineau. 2019. Tarmac: Targeted multi-agent communication. In *Proceedings of the International Conference on Machine Learning (ICML)*. 1538–1546.
 - [9] Hadi Elzayn, Shahin Jabbari, Christopher Jung, Michael Kearns, Seth Neel, Aaron Roth, and Zachary Schutzman. 2019. Fair algorithms for learning in allocation problems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*. 170–179.
 - [10] Tingxiang Fan, Pinxin Long, Wenxi Liu, and Jia Pan. 2020. Distributed multi-robot collision avoidance via deep reinforcement learning for navigation in complex scenarios. *International Journal of Robotics Research (IJRR)* 39, 7 (2020), 856–892.
 - [11] Jakob Foerster, Gregory Farquhar, Triantafyllos Afouras, Nantas Nardelli, and Shimon Whiteson. 2018. Counterfactual multi-agent policy gradients. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 2974–2982.
 - [12] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. 1997. The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine* 4, 1 (1997), 23–33.
 - [13] Julio Erasmo Godoy, Ioannis Karamouzas, Stephen J Guy, and Maria Gini. 2016. Implicit coordination in crowded multi-agent navigation. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 2487–2493.
 - [14] Claudia V Goldman and Shlomo Zilberstein. 2004. Decentralized control of cooperative systems: Categorization and complexity analysis. *Journal of Artificial Intelligence Research (JAIR)* 22 (2004), 143–174.
 - [15] Niko A Grupen, Bart Selman, and Daniel D Lee. 2022. Cooperative multi-agent fairness and equivariant policies. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 9350–9359.
 - [16] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *Proceedings of the International Conference on Machine Learning (ICML)*. 1861–1870.
 - [17] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. 2018. Soft actor-critic algorithms and applications. *arXiv preprint arXiv:1812.05905* (2018).
 - [18] Wolfgang Hönig, James A Preiss, TK Satish Kumar, Gaurav S Sukhatme, and Nora Ayanian. 2018. Trajectory planning for quadrotor swarms. *IEEE Transactions on Robotics (T-RO)* 34, 4 (2018), 856–869.
 - [19] Dan Horgan, John Quan, David Budden, Gabriel Barth-Maron, Matteo Hessel, Hado van Hasselt, and David Silver. 2018. Distributed prioritized experience replay. In *Proceedings of the International Conference on Learning and Representation (ICLR)*.
 - [20] Edward Hughes, Joel Z Leibo, Matthew Phillips, Karl Tuyls, Edgar Dueñez-Guzman, Antonio Garcia Castañeda, Iain Dunning, Tina Zhu, Kevin McKee, Raphael Koster, et al. 2018. Inequity aversion improves cooperation in intertemporal social dilemmas. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*. 3330–3340.
 - [21] Natasha Jaques, Angeliki Lazaridou, Edward Hughes, Caglar Gulcehre, Pedro Ortega, DJ Strouse, Joel Z Leibo, and Nando De Freitas. 2019. Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*. 3040–3049.
 - [22] Jiechuan Jiang and Zongqing Lu. 2019. Learning fairness in multi-agent systems. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*. 13854–13865.
 - [23] Tobias Johannink, Shikhar Bahl, Ashvin Nair, Jianlan Luo, Avinash Kumar, Matthias Loskyll, Juan Aparicio Ojea, Eugen Solowjow, and Sergey Levine. 2019. Residual reinforcement learning for robot control. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 6023–6029.
 - [24] Leslie Pack Kaelbling, Michael L. Littman, and Anthony R. Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* 101, 1 (1998), 99–134.
 - [25] Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. 2017. Multi-agent reinforcement learning in sequential social dilemmas. In *Proceedings of the International Joint Conference on Autonomous Agents & Multiagent Systems (AAMAS)*. 464–473.
 - [26] Jiahui Li, Kun Kuang, Baoxiang Wang, Furui Liu, Long Chen, Fei Wu, and Jun Xiao. 2021. Shapley counterfactual credits for multi-agent reinforcement learning. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery & Data Mining (KDD)*. 934–942.
 - [27] Jiaoyang Li, Pavel Surynek, Ariel Felner, Hang Ma, TK Satish Kumar, and Sven Koenig. 2019. Multi-agent path finding for large agents. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 7627–7634.
 - [28] Pinxin Long, Tingxiang Fan, Xinyi Liao, Wenxi Liu, Hao Zhang, and Jia Pan. 2018. Towards optimally decentralized multi-robot collision avoidance via deep reinforcement learning. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 6252–6259.
 - [29] Ryan Lowe, Yi I Wu, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. 2017. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*. 6382–6393.
 - [30] Carlos E Luis, Marijan Vukosavljev, and Angela P Schoellig. 2020. Online trajectory generation with distributed model predictive control for multi-robot motion planning. *IEEE Robotics and Automation Letters (RA-L)* 5, 2 (2020), 604–611.
 - [31] Keisuke Okumura, Ryo Yonetani, Mai Nishimura, and Asako Kanezaki. 2022. Ctrms: Learning to construct cooperative timed roadmaps for multi-agent path planning in continuous spaces. In *Proceedings of the International Joint Conference on Autonomous Agents & Multiagent Systems (AAMAS)*. 972–981.
 - [32] Krishan Rana, Ben Talbot, Vibhavari Dasagi, Michael Milford, and Niko Sünderhauf. 2020. Residual reactive navigation: Combining classical and learned navigation strategies for deployment in unknown environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 11493–11499.
 - [33] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).
 - [34] Umer Siddique, Paul Weng, and Matthieu Zimmer. 2020. Learning fair policies in multi-objective (deep) reinforcement learning with average and discounted rewards. In *Proceedings of the International Conference on Machine Learning (ICML)*. 8905–8915.
 - [35] Tom Silver, Kelsey Allen, Josh Tenenbaum, and Leslie Kaelbling. 2018. Residual policy learning. *arXiv preprint arXiv:1812.06298* (2018).
 - [36] Roni Stern, Nathan R Sturtevant, Ariel Felner, Sven Koenig, Hang Ma, Thayne T Walker, Jiaoyang Li, Dor Atzmon, Liron Cohen, TK Satish Kumar, et al. 2019. Multi-agent pathfinding: Definitions, variants, and benchmarks. In *Proceedings of the International Symposium on Combinatorial Search (SoCS)*. 151–158.
 - [37] Jianyu Su, Stephen Adams, and Peter A Beling. 2020. Counterfactual multi-agent reinforcement learning with graph convolution communication. *arXiv preprint arXiv:2004.00470* (2020).
 - [38] Sebastian Thrun. 2002. Probabilistic robotics. *Commun. ACM* 45, 3 (2002), 52–57.
 - [39] Jur Van den Berg, Ming Lin, and Dinesh Manocha. 2008. Reciprocal velocity obstacles for real-time multi-agent navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. 1928–1935.
 - [40] Harm Van Seijen, Mehdi Fatemi, Joshua Romoff, Romain Laroche, Tavian Barnes, and Jeffrey Tsang. 2017. Hybrid reward architecture for reinforcement learning. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*. 5398–5408.
 - [41] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*. 6000–6010.
 - [42] Rose E Wang, J Chase Kew, Dennis Lee, Tsang-Wei Edward Lee, Tingnan Zhang, Brian Ichter, Jie Tan, and Aleksandra Faust. 2020. Model-based reinforcement learning for decentralized multiagent rendezvous. In *Proceedings of the Conference on Robot Learning (CoRL)*, Vol. 155. 711–725.
 - [43] Pei Xu and Ioannis Karamouzas. 2021. Human-inspired multi-agent navigation using knowledge distillation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 8105–8112.
 - [44] Jiachen Yang, Alireza Nakhaei, David Isele, Kikuo Fujimura, and Hongyuan Zha. 2020. CM3: Cooperative Multi-goal Multi-stage Multi-agent Reinforcement Learning. In *Proceedings of the International Conference on Learning and Representation (ICLR)*.
 - [45] Ryo Yonetani and Keisuke Okumura. 2022. JAXMAPP: Jax-based library for multi-agent path planning in continuous spaces. <https://github.com/omron-sinix/jaxmapp>
 - [46] Yuyu Yuan, Pengqian Zhao, Ting Guo, and Hongpu Jiang. 2022. Counterfactual-based action evaluation algorithm in multi-agent reinforcement learning. *Applied Sciences* 12, 7 (2022), 3439.
 - [47] Chongjie Zhang and Julie A Shah. 2014. Fairness in multi-agent sequential decision-making. In *Proceedings of the Annual Conference on Neural Information Processing Systems (NeurIPS)*, Vol. 2. 2636–2644.
 - [48] Matthieu Zimmer, Claire Glanois, Umer Siddique, and Paul Weng. 2021. Learning fair policies in decentralized cooperative multi-agent reinforcement learning. In *Proceedings of the International Conference on Machine Learning (ICML)*. 12967–12978.