

Feedback-efficient Active Preference Learning for Socially Aware Robot Navigation

Ruiqi Wang¹, Weizheng Wang^{1,2}, and Byung-Cheol Min¹

Abstract—Socially aware robot navigation, where a robot is required to optimize its trajectory to maintain comfortable and compliant spatial interactions with humans in addition to reaching its goal without collisions, is a fundamental yet challenging task in the context of human-robot interaction. While existing learning-based methods have achieved better performance than the preceding model-based ones, they still have drawbacks: reinforcement learning depends on the handcrafted reward that is unlikely to effectively quantify broad social compliance, and can lead to reward exploitation problems; meanwhile, inverse reinforcement learning suffers from the need for expensive human demonstrations. In this paper, we propose a feedback-efficient active preference learning approach, FAPL, that distills human comfort and expectation into a reward model to guide the robot agent to explore latent aspects of social compliance. We further introduce hybrid experience learning to improve the efficiency of human feedback and samples, and evaluate benefits of robot behaviors learned from FAPL through extensive simulation experiments and a user study (N=10) employing a physical robot to navigate with human subjects in real-world scenarios. Source code and experiment videos for this work are available at: <https://sites.google.com/view/san-fapl>.

I. INTRODUCTION

Advances in artificial intelligence-embedded robotics are increasingly enabling robots to work in environments that necessitate human-robot interaction (HRI). Delivery robots around university campuses, guide robots in shopping malls, elder care robots at nursing homes, and other such applications all require robots to perform socially aware navigation in human-rich environments, wherein the robots must not only consider how to complete navigation tasks successfully but also recognize and follow social etiquette to sustain comfortable spatial interaction with humans [1], [2]. For example, when navigating in a human-filled environment as depicted in Fig. 1, beyond simply reaching the final goal without collisions, the robot must maintain an acceptable distance from other pedestrians and adjust its movements to generate a comfortable interaction experience for humans.

Existing research in the field of socially aware navigation can be divided into two main categories: model-based and learning-based approaches. Model-based methods [3]–[5] aim to model the social conventions and dynamics of crowds to serve as additional parameters for traditional multi-agent collision-free robot navigation algorithms. However, it is quite difficult to emulate with one single model the precise



Fig. 1. An example of real-world socially aware navigation with a mobile robot in a human-filled environment.

rules followed by all pedestrians [1], not to mention that oscillatory trajectories can be produced [6].

As an alternative, learning-based approaches have achieved better performance [1]. Within this category, there are two main types: reinforcement learning (RL) [6]–[8] and inverse reinforcement learning (IRL) [9]–[11]. A primary shortcoming of RL methods is that engineering a meticulous handcrafted reward which emulates non-quantifiable social compliance is a non-trivial endeavor. Another problem stemming from handcrafted rewards is reward exploitation, that is, robots learn to achieve high rewards via some undesired and unnatural action that impairs human comfort. On the other hand, IRL methods, where a policy or reward is learned from human demonstrations, can avoid reward engineering and exploitation and allow experts to introduce human insights and comfort into robot policy. However, obtaining sufficient and accurate demonstrations is expensive, and careful feature engineering is required to achieve sound performance [12].

In this paper, we present a Feedback-efficient Active Preference Learning approach (FAPL) for socially aware robot navigation. The framework and procedure are illustrated in Fig. 2. Our main idea is to efficiently distill a reward model, one having human intelligence and comfort embedded, via active preference learning. Then, the learned reward model imitates a human teacher in the subsequent RL process to guide the robot agent to explore the latent space of social compliance. To reduce the heavy human workload typically required in active preference learning, we introduce hybrid experience learning, which consists of curious exploration based on [13]–[15] and expert demonstration. Diverse exploration experiences provide a good state coverage while expert experiences guarantee positive benchmarks, enabling human teachers to give instructional feedback efficiently and accelerating the learning process.

To evaluate the performance of FAPL, we conducted

¹SMART Laboratory, Department of Computer and Information Technology, Purdue University, West Lafayette, IN, USA. [wang5357,minb]@purdue.edu.

²College of Mechanical and Electrical Engineering, Beijing University of Chemical Technology, Beijing (BUCT), China. wz.w.robot@gmail.com.

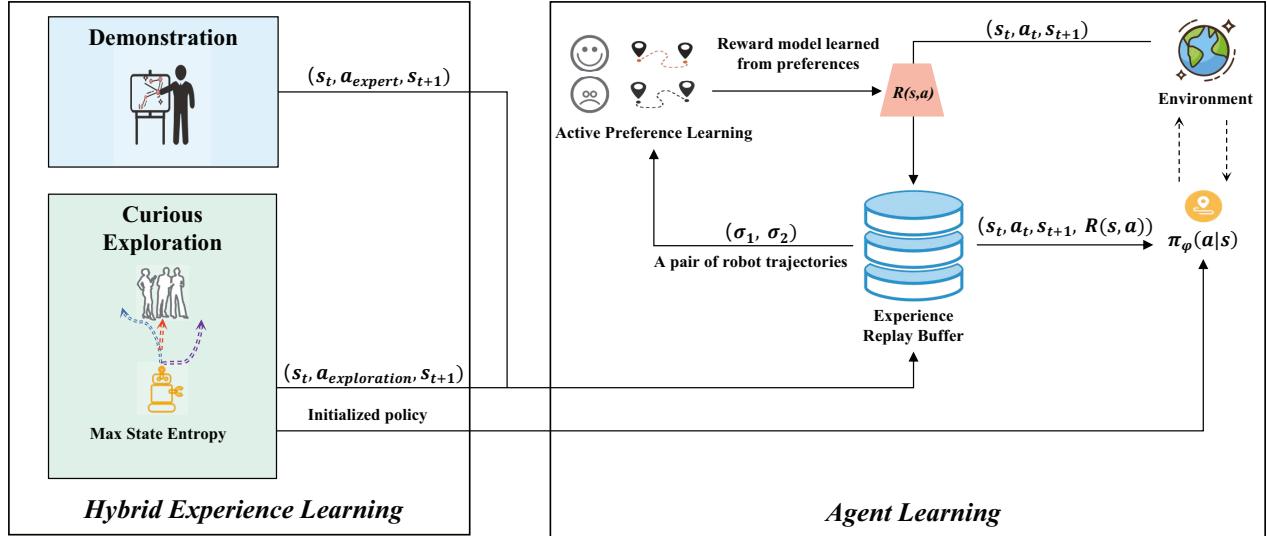


Fig. 2. Framework of the proposed FAPL method. The FAPL is composed of two parts: 1) Hybrid Experience Learning (Section III-B) and 2) Agent Learning (Section III-C, III-D). The robot agent begins with curious exploration, where it is encouraged to take different actions and reach diverse states through a state-entropy-based primitive reward. The collected exploration experiences are stored in a replay buffer along with the expert experiences from human demonstrations. Next, human teachers express preferences for a pair of robot navigation trajectories in the buffer, based on which the reward model is learned. Then, all samples are updated with a new reward value from each new generation of the reward model. At the end, the robot agent utilizes the updated samples to optimize a policy that gains maximum return from the model distilled from human preferences through off-policy RL.

extensive experiments to compare it with other four state-of-the-art baseline methods, [3], [6]–[8], and with one ablation model in both simulation and real-world scenarios for quantitative measurement and qualitative analysis.

Our main contributions are: 1) To the best of our knowledge, this paper is the first to introduce active preference learning for socially aware navigation, such a purely data-driven method can tailor robot behaviors to human expectations and comfort, without the massive reward or feature engineering required in previous research; 2) The hybrid experience learning we introduce improves the efficiency of human feedback and samples; and 3) Our experiments show that FAPL can lead to more desirable and natural robot navigation behaviors.

II. BACKGROUND

Related Work. In optimizing robot trajectories for human comfort, existing RL-based methods [6]–[8] rely on a hand-crafted reward function as shown in Eq. 1, which penalizes collisions and uncomfortable distances.

$$R_t = \begin{cases} -0.25 & \text{if collision} \\ -0.1 + \frac{d_t}{2} & \text{else if distance from human } d_t < 0.2 \\ 1 & \text{else if reaching the goal} \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

However, distance is not the only consideration in robot movement that influences human comfort [1], [2]. As such, a handcrafted reward function is unlikely to lead to robot behaviors that satisfy the broad and non-quantifiable human expectations of robots in socially compliant navigation. Handcrafted reward functions are also conducive to reward exploitation, where robots learn undesired and unnatural but highly-rewarded behaviors that impair human comfort, e.g., frequently turning 180 degrees to avoid collisions or invasive proximity.

Another mainstream approach concerns IRL-based methods [9]–[11], which aim to learn a policy or reward function through learning from human demonstration. Such human-in-loop learning approaches introduce expert intelligence to guide the robot agent, avoiding the reward engineering and exploitation characteristic of RL, and introducing naturalness to robot trajectories. Nevertheless, as the expert demonstrations constitute the only sample resources, it is very expensive to obtain sufficient and accurate samples so as to cover all latent aspects of social compliance. For instance, learning from demonstrations that lack negative data, e.g., collisions, may result in a policy that only values comfort without safety [12]. Moreover, extensive feature engineering is required to achieve reasonable performance [9].

Active preference or feedback learning [16], [17], which relies on human feedback rather than demonstrations to provide corrective and adaptable instructions to guide the learning process, can serve as a good alternative for addressing the above-described challenges. Such interactive reinforcement learning can tailor robot behaviors to human preference naturally, without reward engineering and exploitation. In addition, it does not require and therefore is not limited by expensive human demonstrations. However, achieving good performance with active preference learning does require sufficiently frequent human feedback. Thus, it is essential to improve feedback efficiency so as to reduce the requisite human workload.

Our approach, feedback-efficient active preference learning or FAPL, can encode human comfort and expectation into a reward model and then into robot policy. To improve feedback efficiency, we introduce hybrid experience learning, which enables a human teacher to provide more critical feedback in the process of active reward learning. In contrast to existing RL methods, FAPL distills the reward model from human preferences rather than having it handcrafted, covers

social compliance more intuitively and comprehensively, and avoids reward exploitation. Also, distinct from IRL methods, FAPL introduces expert intelligence without suffering from expensive, noisy, and scarce human demonstrations, nor from extensive feature engineering.

Problem Formulation. We follow the formulation in previous work [6]–[8]. The task of navigating with humans through a spatial HRI environment is formulated as a partially observable Markov decision process (POMDP) problem in RL. Consider the movement space of all agents as a two-dimensional Euclidean space. Let h_t^n present the state of n^{th} human observable by the robot at time t . Each h_t^n consists of the position: (p_x^n, p_y^n) and angle: β_h of n^{th} human agent. Different from previous research, the human speed and radius are considered unobservable for better sim-to-real transfer, as these are expensive to be perceived precisely. Let w_t present the the robot’s state, which includes the velocity: (v_x, v_y) , maximum and minimum speed: (v_{max}, v_{min}) , position: (p_x, p_y) , radius: ρ , angle: β_r of the robot and the position of navigation goal: (g_x, g_y) . Then the joint state of the environment can be defined as $s_t^{int} = [w_t, h_t^1, h_t^2, \dots, h_t^n]$. In each time step, the robot agent starts at an initial joint state s_t^{int} , where it takes the action a_t generated by its policy $\pi(a_t | s_t)$, and then reaches the next joint state s_{t+1}^{int} via a random unknown transition function $P(\cdot | s_t, a_t)$ from the environment and gets a corresponding reward r_t . The optimal policy: $\pi^* : s_t^{int} \mapsto a_t$, is the policy which can receive maximum expected reward return $\mathbb{E}[\sum_{k=0}^{\infty} \gamma^k r_{t+k}]$, where γ is a discount factor.

Soft Actor-Critic Training Framework. In this work, Soft Actor-Critic (SAC) [18] will be utilized as the training framework. There are two alternating parts for parameter iteration in SAC: soft policy evaluation and soft policy improvement. In soft policy evaluation, a Q-function with parameters θ , is updated to minimize Bellman-residual based objective (see Eqs. (2), (3)), where \mathcal{B} presents the replay buffer, γ is a discount parameter, $\bar{\theta}$ and α are the delayed and temperature parameters respectively.

$$\mathcal{L}_Q(\theta) = \mathbb{E}_{(s_t, a_t) \sim \mathcal{B}} \left[\frac{1}{2} \left(Q_\theta(s_t, a_t) - (r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim p} [V_{\bar{\theta}}(s_{t+1})]) \right)^2 \right] \quad (2)$$

where,

$$\bar{V}(s_t) = \mathbb{E}_{a_t \sim \pi_\phi} [Q_{\bar{\theta}}(s_t, a_t) - \alpha \log \pi_\phi(a_t | s_t)]. \quad (3)$$

In soft policy improvement, the policy π_ϕ is optimized to minimize the following objective:

$$\mathcal{L}_\pi(\phi) = \mathbb{E}_{s_t \sim \mathcal{B}} [\mathbb{E}_{a_t \sim \pi_\phi} [\alpha \log (\pi_\phi(a_t | s_t)) - Q_\theta(s_t, a_t)]] \quad (4)$$

III. APPROACH

A. Overview

In this section we introduce FAPL: Feedback-efficient Active Preference Learning. A social-compliance-embedded reward function \hat{R}_μ , a Q-function Q_θ and a policy π_ϕ will be updated by following steps:

- **Expert Demonstration:** Initially, we let human experts provide demonstrations of the socially aware navigation task and store them into an initialized experience replay buffer EB as expert experiences, $(s_t, a_{expert}, s_{t+1})$. (Section III-B)
- **Curious Exploration:** We train a pre-policy π_ϕ by maximizing a state-entropy-based reward to conduct curious exploration and collect diverse samples, which are stored in the buffer EB as exploration experiences $(s_t, a_{exploration}, s_{t+1})$. Meanwhile, a temporary buffer B is set to store incoherent samples generated during the early training process as temporary experiences $(s_t, a_{temporary}, s_{t+1})$. (Section III-B)
- **Active Reward Learning:** The reward model \hat{R}_μ , encoded with human comfort and intelligence, is distilled from human feedback via active preference learning. Then all collected samples in the EB and B are updated with a new reward value by \hat{R}_μ . (Section III-C)
- **Off-policy Learning:** The Q-function Q_θ and policy π_ϕ will be optimized using all updated samples via SAC to gain maximum return from the distilled reward model \hat{R}_μ . (Section III-D)

B. Hybrid Experience Learning

At the start of traditional active preference learning, the initialized agent follows a random policy to interact with the environment and obtain samples for a human to judge. However, such a policy covers the state space badly and leads to incoherent behaviors, making it hard for the human teacher to provide meaningful feedback [14]. Thus, a lot of samples and human efforts are required to make initial progress. To address this challenge, we introduce hybrid experience learning module, which consists of curious exploration and expert demonstration.

In curious exploration, we inspire the robot agent to explore the state space of socially aware navigation thoroughly to provide a better state coverage by using a k -NN-based state entropy estimator, adapted from [14], [15], [19], as the incipient reward function. The state entropy estimator $H_{state}(s)$, evaluates the sparsity and randomness of the state distribution by measuring the space distance between each state and its K^{th} nearest neighbors as:

$$H_{state}(s) := \sum_{t=1}^n \log \left(p + \frac{1}{k} \sum_{s_t^{(k)} \in N_k(s_t)} \|s_t - s_t^{(k)}\| \right). \quad (5)$$

where p is a parameter for numerical stability (usually be fixed to 1), and $s_t^{(k)}$ are the k neighbors of s_t in state space.

Correspondingly, the incipient exploration reward and objective function can be defined as:

$$R_{exploration}(s_t) = \log \left(p + \frac{1}{k} \sum_{s_t^{(k)} \in N_k(s_t)} \|s_t - s_t^{(k)}\| \right), \quad (6)$$

$$\phi_\pi^* = \underset{\phi}{\operatorname{argmax}} \sum_{t=1}^n R_{exploration}(s_t). \quad (7)$$

Maximizing the objective function in Eq. (7), we train an initial policy π_ϕ , by which the robot agent can explore a broader range of states. However, the socially aware navigation task not only requires the robot to reach the final goal without collisions, but values more on human comfort, which is hidden in the state space and is, thus, quite difficult to cover by the robot itself. As a result, a lot of iteration rounds and therefore a high volume of human feedback is needed to optimize a desired policy. To accelerate learning process and further improve feedback efficiency, we add human demonstrations, e.g., controlling the robot with a keyboard to demonstrate the navigation task, to introduce expert experiences to cover the hidden aspects of social compliance and provide naturalness to robot movements. The human teachers can regard expert demonstrations as benchmarks to show preferences, offering critical feedback more easily.

Meanwhile, it must be mentioned that human demonstrations are always accompanied with noise [9]. To avoid the influence of inaccurate demonstrations, we only store a triple, $(s_t, a_{\text{expert}}, s_{t+1})$, without a reward value in the replay buffer EB , which is different from traditional IRL. Then the human teacher can express preference on good trajectories from demonstration or dislike on inaccurate ones to add reward labels on these demonstration samples. Namely, the expert experiences are not regarded as oracles, they will be judged by human teachers like the exploration experiences, under which circumstance, we can make advantages of good demonstrations without being affected by noisy ones. The full process of hybrid experience learning is concluded in Algorithm 1.

C. Active Reward Learning

The objective of active reward learning is to learn a reward function \hat{R}_μ , which is a neural network with parameters μ , to encode human expectation and preference of how a socially compliant robot should act in spatial HRI. It has been shown that people feel much easier to make relative judgements than direct rating [20]. Therefore, instead of asking human teacher to give a rate value for a single set of robot trajectories, we provide two segments for human to express preferences at a time, e.g., which segment is better or worse. We follow the framework of [14], [16], to optimize our reward model \hat{R}_μ . Robot trajectories stored in the experience replay buffer EB will be divided into several segments, each segment σ contains a sequence of states and actions in robot movements: $\sigma = (s_t, a_t, \dots, s_{t+n}, a_{t+n})$. In each feedback step, the agent will query human preference Υ , which is one of (1,0), (0,1), (0.5,0.5), on two segments σ_1, σ_2 . The judgment with two segments will be stored in a database D as $(\sigma_1, \sigma_2, \Upsilon)$.

Then, a preference predictor \mathcal{P}_μ is built to train the reward model \hat{R}_μ , where $\sigma_1 \succ \sigma_2$ means that σ_1 will be preferred:

$$\mathcal{P}_\mu [\sigma_1 \succ \sigma_2] = \frac{\exp \left(\sum_{s_t, a_t \in \sigma_1} \hat{R}_\mu(s_t, a_t) \right)}{\exp \left(\sum_{s_t, a_t \in \sigma_1 \cup \sigma_2} \hat{R}_\mu(s_t, a_t) \right)}. \quad (8)$$

Algorithm 1 Hybrid Experience Learning.

```

1: Initialize a critic  $Q_\theta$  and a policy  $\pi_\phi$ 
2: Initialize two replay buffers  $B \leftarrow \emptyset$  and  $EB \leftarrow \emptyset$ 
3: while Expert Demonstration do
4:   Store expert experiences
5:    $EB \leftarrow EB \cup \{(s_t, a_{\text{expert}}, s_{t+1})\}$ 
6: end while
7: while Curious Exploration do
8:   Given  $t_{\text{max}}$  and  $t_{\text{memory}}$ 
9:    $t \leftarrow 1$ 
10:  while  $t < t_{\text{max}}$  do
11:    Take  $a_t \sim \pi_\phi(a_t | s_t)$  in  $s_t$  and reach  $s_{t+1}$ 
12:    Compute reward  $r_t \leftarrow R_{\text{exploration}}^{s_t}$  in (6)
13:    if  $t < t_{\text{memory}}$  then
14:      Store temporary experiences
15:       $B \leftarrow B \cup \{(s_t, a_{\text{temporary}}, s_{t+1}, r_t)\}$ 
16:    else
17:      Store exploration experiences
18:       $EB \leftarrow EB \cup \{(s_t, a_{\text{exploration}}, s_{t+1})\}$ 
19:    end if
20:    Sample minibatch transitions  $\sim B$ 
21:    Optimize  $\theta, \phi$  by following  $\mathcal{L}_Q(\theta)$  in (2)
22:    and  $\mathcal{L}_\pi(\phi)$  in (4)
23:     $t=t+1$ 
24:  end while
25: end while
26: return  $B, EB, \pi_\phi$ 

```

The assumption behind is that the probability of a human teacher's preference on one segment σ_i in a pair of segments (σ_i, σ_j) depends exponentially on the accumulated reward value of σ_i gained from \hat{R}_μ over the whole reward value of both segments.

Based on the preference predictor \mathcal{P}_μ , we optimize the reward model \hat{R}_μ by minimizing the cross-entropy loss function $\mathcal{L}(\hat{R}_\mu)$, which evaluates the difference between the prediction of \hat{R}_μ and ground-truth human preference:

$$\mathcal{L}(\hat{R}_\mu) = - \sum_{(\sigma_1, \sigma_2, \mu) \in D} \Upsilon(1) \log \mathcal{P}_\mu [\sigma_1 \succ \sigma_2] + \Upsilon(2) \log \mathcal{P}_\mu [\sigma_2 \succ \sigma_1]. \quad (9)$$

In addition, we adapt the uncertainty-based query selection method proposed by [21] to determine which pair of segments (σ_i, σ_j) of robot behaviors stored in EB are selected to query the human teacher for preference Υ each time. Such an informative query selection can pick up behaviors with maximum entropy, which are typically "uncertain samples on the decision boundary", leading to a significant decrease of uncertainty in unlabelled behaviors and informative feedback from humans.

D. Off-policy learning

To further improve sample efficiency, we will adapt an off-policy RL framework, SAC, for following training. However, compared with on-policy RL frameworks, it is less stable when utilizing the reward model \hat{R}_μ gained from active

learning, for such a reward function is always updated and thus non-static during training process, under which circumstance off-policy RL will reuse all samples in the replay buffer that contain inaccurate reward values provided by previous in-process and non-optimized reward models. To address this issue, all previous samples stored in both EB and B will be updated with a new reward label every time when a new reward model is distilled from human preference, and temporary experiences in B will be transferred to EB . Then the Q-function Q_θ and pre-trained policy π_ϕ gained from curious exploration will rely on the reward model \hat{R}_μ combined with updated samples in EB for optimization. The process of active reward learning and FAPL is presented in Algorithm 2 and 3 respectively.

Algorithm 2 Active Reward Learning

```

1: Initialize  $\hat{R}_\mu$  and a database  $D \leftarrow \emptyset$ 
2: Given number of feedback  $M$ 
3:  $m \leftarrow 1$ 
4:  $B, EB \leftarrow$  Algorithm 1
5: while  $m < M$  do
6:   Select segments  $(\sigma_0, \sigma_1) \sim EB$ 
7:   Query human feedback  $\Upsilon$ 
8:   Store preference  $D \leftarrow D \cup \{(\sigma_0, \sigma_1, \Upsilon)\}$ 
9:   Sample minibatch preference  $\sim D$ 
10:  Optimize  $\mu$  by  $\mathcal{L}(\hat{R}_\mu)$  in (9)
11:  Update entire  $EB$  and  $B$  using  $\hat{R}_\mu$ 
12:   $m \leftarrow m + 1$ 
13: end while
14: Store  $EB \leftarrow EB \cup B$ 
15: return  $\hat{R}_\mu$  and  $EB$ 

```

Algorithm 3 FAPL

```

1: Initialize a critic  $Q_\theta$ 
2: // Hybrid Experience Learning
3:  $\pi_\phi \leftarrow$  Algorithm 1
4: // Agent Learning
5: for each time step do
6:   // Active Reward Learning
7:    $EB, \hat{R}_\mu \leftarrow$  Algorithm 2
8:   Take  $a_t \sim \pi_\phi(a_t | s_t)$  and reach  $s_{t+1}$ 
9:   Compute reward  $\hat{R}_\mu(s_t, a_t) \leftarrow \hat{R}_\mu$ 
10:  Store transitions
11:   $EB \leftarrow EB \cup \{(s_t, a_t, s_{t+1}, \hat{R}_\mu(s_t, a_t))\}$ 
12:  // Policy Optimization
13:  for each gradient step do
14:    Sample minibatch transitions  $\sim EB$ 
15:    Optimize  $\theta, \phi$  by  $\mathcal{L}_Q(\theta)$  in (2) and  $\mathcal{L}_\pi(\phi)$  in (4)
16:  end for
17: end for
18: return  $\pi_\phi, Q_\theta$ 

```

E. Implementation Details

We pretrain the policy π_ϕ in curious exploration for 3,000 episodes, the samples gained in first 1,000 episodes are stored in B and the rest is stored in EB . In expert demonstration part, human trainers (authors) use the keyboard to

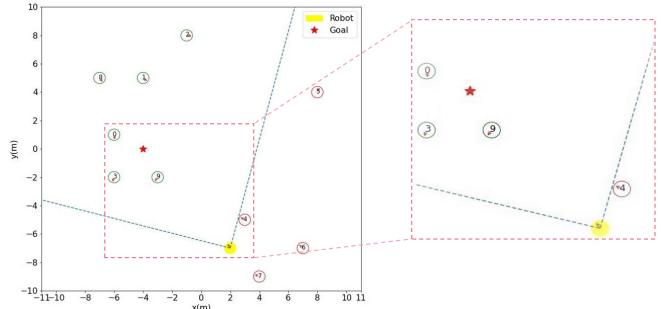


Fig. 3. Simulation environment illustration. In a $22 m \times 20 m$ two-dimensional space, the yellow circle indicates the robot. The blue dotted line illustrates the robot FoV, humans that can be detected by the robot are green circles while those out of robot view are red circles. The red star is the robot goal, and the orientation and number of each agent are presented as a red arrow and a black number respectively.

control the robot velocity along the direction of x and y axis: (v_x, v_y) to provide 500 rounds of demonstrations in simulation. Once the robot is controlled to reach the goal, we say one round of demonstration is done. The reward model in active reward learning is a 3-layer neural network with 265 hidden units in each layer and the activation function of Leaky ReLUs. We utilize Adam to train the reward model with an original learning rate of 4×10^{-4} . The segment length is set to 35 (about 15 s), which means each segment contains 35 tuples of (s_t, a_t) . And we recruited 10 human raters, who are students or professionals in robotics, to provide 1,500 rounds of preferences on robot behaviors for our model and ablation model respectively (IV-A.3). Once the human rater gives a judgement Υ on one pair of robot trajectory segments σ_1, σ_2 , one round of preference is done.

IV. EXPERIMENTS AND RESULTS

A. Simulation Experiment

1) *Simulation Environment*: We adapted a simulation environment from [6], [7] for training and experiments as is shown in Fig. 3. Each agent in the environment follows holonomic kinematics, where agents can move to any direction at any time. The agent's action at time t is the preferred velocity in x-axis and y-axis direction: $a_t = (v_t^x, v_t^y)$. Such a velocity is assumed to be immediately achievable. All human agents are controlled by ORCA [3], and the parameters of their policy are generated by a Gaussian distribution to provide random pedestrian behaviors, e.g., different velocity and goals. To avoid learning the exceedingly aggressive behavior that the robot compels all humans to yield, we follow previous research [7] to set an invisible robot setting, where the human agent is required only to do reaction, e.g., yield, to other human agents. Meanwhile, different from previous setting, we set the degree of robot field-of-view (FoV) as 90° rather than 360° to better mimic real-world scenarios and reduce sim-to-real domain-mismatch issues, since sorting multiple sensors on a physical robot to obtain a global FoV is quite impractical and expensive.

2) *Baselines*: We compare the proposed FAPL with other four state-of-the-art methods mentioned before: ORCA [3] is regarded as the baseline of model-based approaches; CADRL

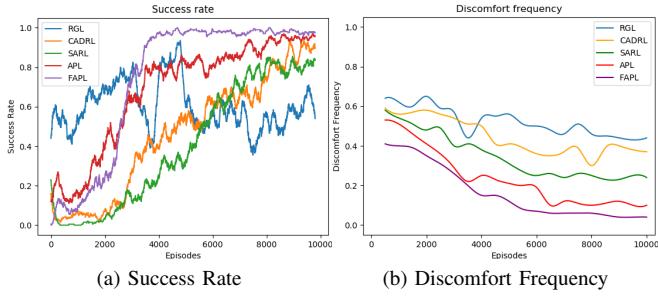


Fig. 4. Learning curves of CADRL, RGL, SARN, APL and FAPL as measured, on (a) success rate (each episode) and (b) discomfort frequency (the mean across each five hundred episodes) under the setting of 10 humans.

[6], SARN [7] and RGL [8] are regarded as the baselines of learning-based approaches.

3) *Ablation Study*: To evaluate the feedback efficiency of FAPL, we also implement another ablation model APL, which removes the hybrid experience learning module in FAPL, as the baseline of active preference learning.

4) *Training Details*: We utilize the same reward function in 1 for CADRL, SARN and RGL. The architectures of all networks stay the same in each experiment. All baseline networks are trained by following the implementing details in original papers for 1×10^4 episodes. We train APL and FAPL with 1,500 pieces of human feedback respectively (III-E) for 1×10^4 episodes using a learning rate of 2×10^{-4} with same parameters.

5) *Evaluation*: We compare the learning curves of all learning-based methods in terms of success rate and discomfort frequency, which refers to the percentage of duration where the robot is too close with a human pedestrian. Three experiment setting: 5, 10, 15 motional human agents are involved respectively, are set to evaluate the performance of all methods. In each setting, 500 unseen testing scenarios are tested with five indicators recorded: success rate, time-out rate, collision rate, discomfort frequency and time to reach the goal. The time-out threshold is set as 30, 35 and 40 seconds in each setting separately, and the margin of discomfort frequency is set as 0.3 m based upon [1].

6) *Results*: Fig. 4 shows the learning curves of all learning-based methods in terms of success rate and discomfort frequency. FAPL (purple) can improve the success rate and reduce discomfort frequency continuously, and achieve the best performance in the end of training, outperforming all baselines significantly. Much as RGL (blue) achieves a higher success rate than ours in first 3,000 episodes, it is very unstable and finally only reaches a success rate of about 0.6 in the end. Compared with APL (red), FAPL has a lower success rate and higher discomfort frequency in the beginning, for its initial policy comes from curious exploration where the robot is only encouraged to visit different states. However, FAPL surpasses APL after around 3,000 episodes and achieves the best performance much more quickly, which shows that FAPL can converge significantly faster with the same number of human feedback, quantitatively proving the benefits of the hybrid experience learning module.

Fig. 5 presents the percentage of success, time-out, and

TABLE I
OUTCOME: TIME (SECOND) AND DISCOMFORT FREQUENCY (DISFQ.).

Methods	Time			DisFq.		
	Human Number			Human Number		
	5	10	15	5	10	15
CADRL	29.4	34.7	41.3	0.203	0.258	0.435
SARN	27.3	32.5	38.7	0.137	0.172	0.251
RGL	22.9	27.4	32.7	0.288	0.354	0.474
APL(Ours)	24.2	28.7	31.9	0.029	0.043	0.061
FAPL(Ours)	25.3	30.1	34.1	0.018	0.025	0.048

collision rates of each model. ORCA performs quite badly as expected, for its model is based on the assumption that all agents are always reciprocal, violating the invisible robot setting. RGL also performs badly since it is a model-based RL built based upon global robot FoV and precise perception of human velocity which are hard to obtain in our setting. Meanwhile, the learning based-baselines are competitive to our FAPL and APL in the setting of 5 humans. However, with the increase of the complexity of environment, i.e., the number of humans, our methods perform better and better than the baselines. This shows that APL and FAPL are more efficient and robust in complex task scenarios. And FAPL outperforms APL apparently in the most difficult scenario involving 15 humans.

Table I demonstrates the discomfort frequency and navigation time of each model except ORCA due to its extremely high collision rate. RGL enjoys the best overall performance in navigation time, however, it has the highest discomfort frequency, which means, to achieve a shorter time, the robot takes a lot of aggressive action, hurting human comfort badly. FAPL obtains the lowest discomfort frequency with shorter time than CADRL and SARN, and is competitive to RGL and APL in terms of navigation time. Combining the indicators of navigation time and discomfort frequency shows that our methods lead to robot behaviors that respect human comfort distance more than all baselines and ablation model, without sacrificing too much time due to lazy behaviors, e.g., waiting still for all pedestrians to pass or excessive detouring.

B. Real-world Experiment

Since the social compliance is non-quantifiable and broader than the comfortable distance, it is unsatisfied and insufficient to evaluate it only via the indicator of discomfort frequency in simulation experiments. To intuitively and further evaluate the social compliance of learned robot trajectories from our methods, we recruited human participants to conduct real-world experiments, collecting their feedback from experiences of walking with a robot controlled by different models as another indicator. This real-world experiment has been reviewed by Institutional Review Board (IRB) of BUCT.

1) *Robot and Environment Setup*: We utilized an Enlighten mobile robot [22] (Fig. 6, left) as the platform for real-world experiments. A Kinect v2 camera with an approximately field-of-view of 84.1° was set on the top of the robot to capture human positions. YOLOv5 [23] combined with DeepSORT [24], and Monoloco Library [25], [26] were used for robot-centered human tracking and localization (Fig. 7). To provide sufficient real-time computing power for the

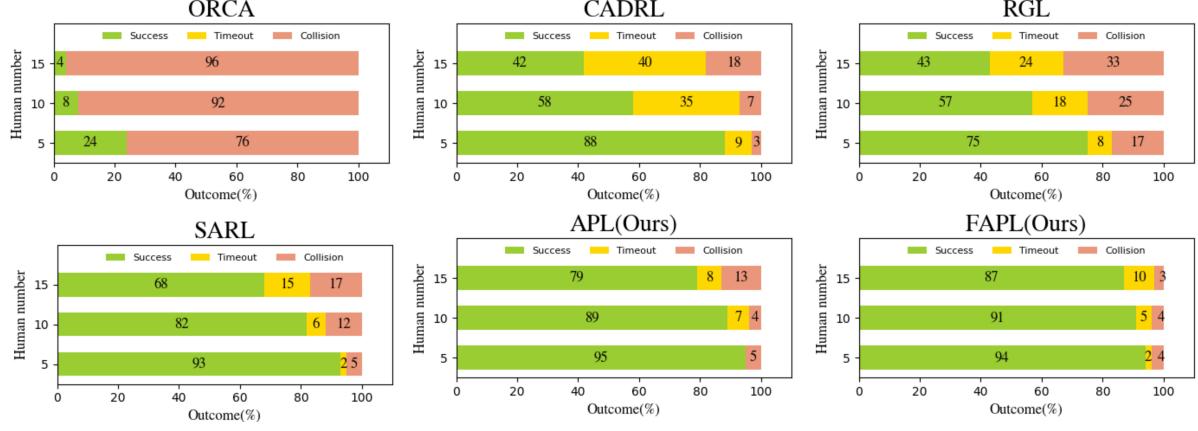


Fig. 5. Outcome: rates of success, collision and time-out under three experimental settings with five, ten and fifteen human agents involved respectively.

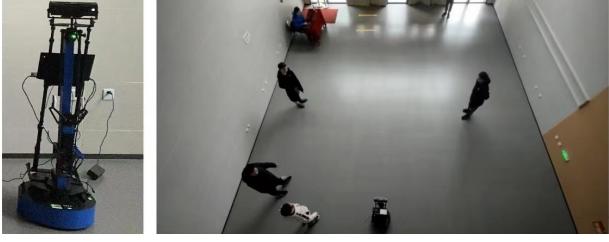


Fig. 6. Illustration of the robot platform (left) and environment (right) employed for the real-world experiment.

perception algorithms and RL controllers, we connected the robot to a host with a RTX 3090 GPU via ROS. To avoid potential damage to the robot hardware, we set up an action restriction: the robot is enforced to stop when its action contains a turning degree greater than 90° . The experimental environment was an approximate $23\text{ m} \times 16\text{ m}$ controlled open space as shown in Fig. 6, right.

2) *Baselines*: The ablation model APL and SARN, the best performing baseline in terms of discomfort frequency in simulation, were selected as the baselines.

3) *Experimental Design*: We recruited 10 participants (2 females and 8 males), all aged over 18 years old ($\mu = 21.9$; $\sigma = 1.92$), and divided them into 2 groups. During each experiment, the robot and participants as pedestrians were required to perform point-to-point navigation tasks in the same environment. The human pedestrians were encouraged to behave naturally without any action restrictions during tasks. We set up three scenarios, where the robot and human pedestrians were assigned with different starting positions and goals. We asked each group to spatially interact with the robot controlled by three models: SARN, APL, and FAPL respectively in the three different scenarios, resulting total 18 tests. The order of three models was randomly set in each scenario for each group, and all participants were unaware of which model was active during each test.

4) *Evaluation*: Once each test was done, the participants were asked to fill a questionnaire to rate robot behaviors of three models in terms of comfort and naturalness using Likert Scale [27] with 1 being strongly disagree and 5 being strongly agree, based on their experiences of spatial interaction with the robot. The participant's responses to the questionnaire were used for quantitative measurement.

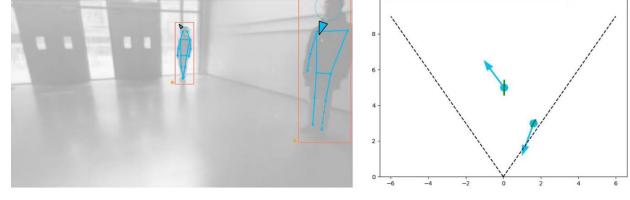


Fig. 7. Illustration of perception algorithms for human tracking (left) and localization (right), where the blue circle and arrow present the position and orientation of a human respectively.

We also conducted a semi-structured interview to collect oral responses of participant's experience in experiments for qualitative analysis.

5) *Quantitative Measurement*: During each test, participants were divided into two categories: close interaction participant (CI) and non-close interaction participant (NCI), to collect fair questionnaire responses, where CI refers to pedestrians who approach the robot within 0.6 m for more than 2 seconds during the test while NCI refers to those who do not. The results are summarized in Fig. 8. The benefits of robot behaviors of FAPL in terms of comfort and naturalness can be obviously seen from the responses of CI. The reason why FAPL does not outperform the other two much from the responses of NCI for comfort is that NCI did not have enough direct experiences to judge, and therefore, they tended to give all three models a middle score. And the reason why responses for naturalness of CI and NCI are similar is that the robot unnatural behaviors, e.g., suddenly stopping, were accompanied with noises that attracted NCI's attention.

6) *Qualitative Analysis*: After each test, all participants were asked an open-ended interview question that: "Could you share your experiences and explain reasons behind?". The insights thematically coded from participants' responses are presented along with supporting quotes as follows.

Long-sighted Decision. One reason why the robot behavior in FAPL is preferred by pedestrians is that it takes more long-sighted decisions, avoiding potential risky situations where the robot is hard to maintain socially compliant interaction with pedestrians. We believe this is because we introduce human intelligence and insights via expert demonstration and feedback in the learning process, enabling the robot to handle complex situations by valuing more on long-term benefits than short-term gains. Some notable quotes are as follows:

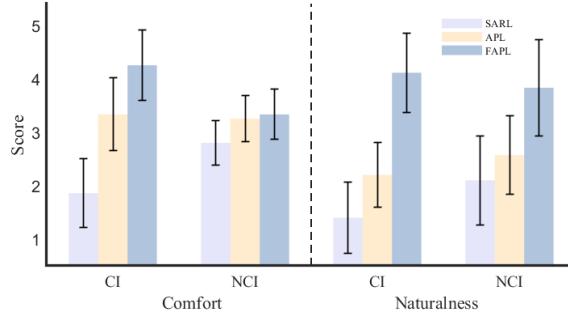


Fig. 8. Summary of participant responses to evaluate robot navigation behaviors in terms of comfort and naturalness.

(P5) "The reason why I like that one (FAPL) is it turned right to give me enough space before I got close to it."

(P9) "Because the third one (FAPL) went left to detour at the start instead of getting involved in the center of the crowd directly, like the first one (SARL)."

Behavior Naturalness. Another reason why our FAPL is preferred is that it leads to more natural behaviors, e.g., does not stop suddenly or turn left and right frequently. We owe this to the human-imitative trajectory samples introduced by expert demonstration in hybrid experience learning, and the avoidance of reward exploitation problems by replacing handcrafted reward functions with reward models distilled from human preference. Some notable quotes are as follows:

(P1) "The second one (SARL) always stopped immediately, and the first (APL) also did that sometime, these made me a little nervous and annoyed."

(P9) "For me, that one (FAPL) moved more like a human, I mean it hardly stopped suddenly and turned frequently, which is preferable to me."

V. CONCLUSION

In this work, we present FAPL, a feedback-efficient active preference learning approach for socially aware robot navigation that distills human comfort and expectation into a reward model to guide the robot to explore the latent space of social compliance. We demonstrate via both simulation and real-world experiments that our method outperforms existing state-of-the-art approaches, leading to more desirable and natural robot behaviors. We also show that by introducing hybrid experience learning, the efficiency of human feedback can be improved.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant No. IIS-1846221.

REFERENCES

- [1] J. Rios-Martinez, A. Spalanzani, and C. Laugier, "From proxemics theory to socially-aware navigation: A survey," *International Journal of Social Robotics*, vol. 7, no. 2, pp. 137–153, 2015.
- [2] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, "Human-aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726–1743, 2013.
- [3] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*. Springer, 2011.
- [4] A. Vemula, K. Muelling, and J. Oh, "Social attention: Modeling attention in human crowds," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 4601–4607.
- [5] G. Ferrer and A. Sanfeliu, "Behavior estimation for a complete framework for human motion prediction in crowded environments," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 5940–5945.
- [6] 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.
- [7] 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.
- [8] C. Chen, S. Hu, P. Nikdel, G. Mori, and M. Savva, "Relational graph learning for crowd navigation," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 10007–10013.
- [9] H. Kretzschmar, M. Spies, C. Sprunk, and W. Burgard, "Socially compliant mobile robot navigation via inverse reinforcement learning," *The International Journal of Robotics Research*, vol. 35, no. 11, pp. 1289–1307, 2016.
- [10] B. Okal and K. O. Arras, "Learning socially normative robot navigation behaviors with bayesian inverse reinforcement learning," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 2889–2895.
- [11] B. Kim and J. Pineau, "Socially adaptive path planning in human environments using inverse reinforcement learning," *International Journal of Social Robotics*, vol. 8, no. 1, pp. 51–66, 2016.
- [12] C.-E. Tsai and J. Oh, "A generative approach for socially compliant navigation," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 2160–2166.
- [13] B. Eysenbach, A. Gupta, J. Ibarz, and S. Levine, "Diversity is all you need: Learning skills without a reward function," in *International Conference on Learning Representations*, 2018.
- [14] H. Liu and P. Abbeel, "Behavior from the void: Unsupervised active pre-training," *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [15] K. Lee, L. M. Smith, and P. Abbeel, "Pebble: Feedback-efficient interactive reinforcement learning via relabeling experience and unsupervised pre-training," in *International Conference on Machine Learning*. PMLR, 2021, pp. 6152–6163.
- [16] P. F. Christiano, J. Leike, T. B. Brown, M. Martic, S. Legg, and D. Amodei, "Deep reinforcement learning from human preferences," in *NIPS*, 2017.
- [17] S. Jamieson, J. P. How, and Y. Girdhar, "Active reward learning for co-robotic vision based exploration in bandwidth limited environments," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 1806–1812.
- [18] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International conference on machine learning*. PMLR, 2018, pp. 1861–1870.
- [19] H. Singh, N. Misra, V. Hnizdo, A. Fedorowicz, and E. Demchuk, "Nearest neighbor estimates of entropy," *American journal of mathematical and management sciences*, vol. 23, no. 3-4, pp. 301–321, 2003.
- [20] N. Wilde, A. Blidaru, S. L. Smith, and D. Kulić, "Improving user specifications for robot behavior through active preference learning: Framework and evaluation," *The International Journal of Robotics Research*, vol. 39, no. 6, pp. 651–667, 2020.
- [21] Y. Yang and M. Loog, "A benchmark and comparison of active learning for logistic regression," *Pattern Recognition*, vol. 83, pp. 401–415, 2018.
- [22] "Enlighten mobile robot." [Online]. Available: <http://www.6-robot.com>
- [23] G. Jocher, "Yolov5." [Online]. Available: <https://github.com/ultralytics/yolov5>
- [24] Z. Pei, "Deepsort." [Online]. Available: https://github.com/ZQPei/deep_sort_pytorch
- [25] L. Bertoni, "Monoloco library." [Online]. Available: <https://github.com/vita-epfl/monoloco>
- [26] L. Bertoni, S. Kreiss, and A. Alahi, "Perceiving humans: from monocular 3d localization to social distancing," *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [27] A. Joshi, S. Kale, S. Chandel, and D. K. Pal, "Likert scale: Explored and explained," *British journal of applied science and technology*, vol. 7, no. 4, p. 396, 2015.