

# MIRROR: Differentiable Deep Social Projection for Assistive Human-Robot Communication

Kaiqi Chen, Jeffrey Fong, and Harold Soh  
 Dept. of Computer Science, National University of Singapore.  
 {kaiqi, jfong, harold}@comp.nus.edu.sg

**Abstract**—Communication is a hallmark of intelligence. In this work, we present MIRROR, an approach to (i) quickly learn human models from human demonstrations, and (ii) use the models for subsequent communication planning in assistive shared-control settings. MIRROR is inspired by social projection theory, which hypothesizes that humans use self-models to understand others. Likewise, MIRROR leverages self-models learned using reinforcement learning to bootstrap human modeling. Experiments with simulated humans show that this approach leads to rapid learning and more robust models compared to existing behavioral cloning and state-of-the-art imitation learning methods. We also present a human-subject study using the CARLA simulator which shows that (i) MIRROR is able to scale to complex domains with high-dimensional observations and complicated world physics and (ii) provides effective assistive communication that enabled participants to drive more safely in adverse weather conditions.

## I. INTRODUCTION

Communication is an essential skill for intelligent agents; it facilitates cooperation and coordination, and enables teamwork and joint problem-solving. However, effective communication is challenging for robots; given a multitude of information that can be relayed in different ways, how should the robot decide *what*, *when*, and *how* to communicate?

In this paper, we focus on assistive shared-control or teleoperation scenarios. As an example, consider the scenario in Fig. 1 where a robot assistant is tasked to provide helpful information to a human driving the blue car in fog (or to explain the robot’s own driving behavior). There are other cars in the scene, which may not be visible to the human. The robot, however, has access to sensor readings that reveal the environment and surrounding cars. To prevent potential collisions, the robot needs to communicate relevant information — either using a heads-up display or verbally — that intuitively, should take into account what the human driver currently believes, what they can perceive, and the actions they may take.

Prior work on planning communication methods in HRI typically rely on human models, which are typically handcrafted using prior knowledge (e.g., [1], [2]) or learned from collected human demonstrations (e.g., [3]). Unfortunately, handcrafted models do not easily scale to complex real-world environments with high-dimensional observations, and data-driven models typically require a large number of demonstrations to generalize well. In this work, we seek to combine prior knowledge with data in a manner that reduces *both* manual specification and sample complexity.

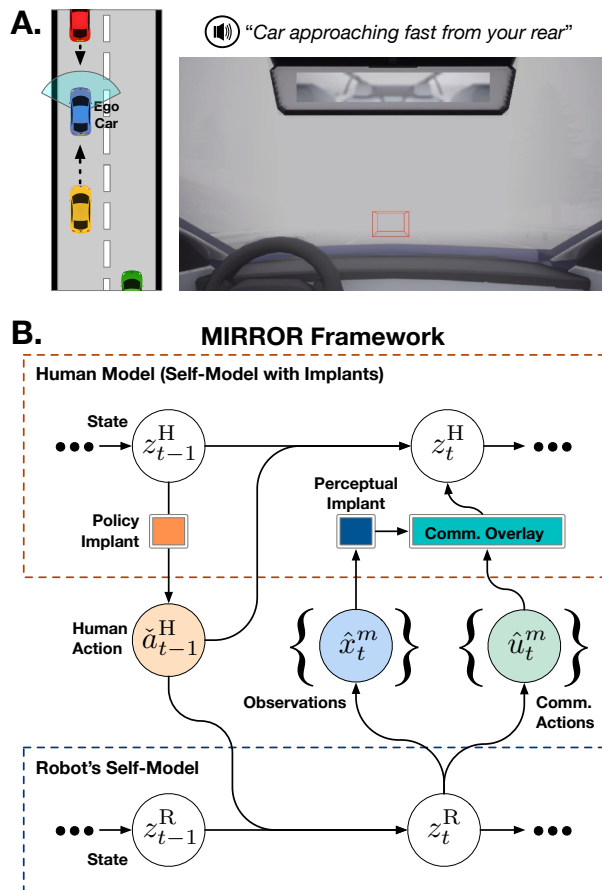


Fig. 1: Human-Robot Communication Example. (A) A Robot Assistant needs to provide information to help a human driving the blue car in dense fog. The human has limited visibility and the assistant can highlight cars on a heads-up-display or provide verbal cues (as the human is unable to see highlighted cars that are in the rear). In this scenario, a collision is imminent—the red car in front is slowing down and a yellow car speeding up from the rear. How can the robot determine *what*, *when*, and *how* to inform the user about? Our proposed MIRROR assistant immediately highlights the red car in front and verbally tells the user about yellow car at the rear. It chooses *not* to tell the human about the green car, whose location has little impact on the human’s decision. (B) Our MIRROR framework is inspired by *social projection*; the robot reasons using a human model that is constructed from its own internal self-model. Strategically-placed learnable “implants” capture how the human is different from the robot. MIRROR plans communicative actions by forward simulating possible futures by coupling its own internal model (to simulate the environment) and the human model (to simulate their actions).

Our key insight is that learning *differences* from a suitable reference model is more data-efficient than learning an entire human model from scratch. We take inspiration from *social projection theory* [4], which suggests that humans have a tendency to expect others to be similar to ourselves, i.e., a person understands other individuals using one’s self as a reference. This inductive bias can be effective when the agents are similar and may promote cooperation [5], [6]. From a cognitive perspective, social projection is a heuristic by which we evaluate and predict another agent’s behavior. Likewise, a robot can use its self-model to reason about a human; in our setup, the robot first learns how to perform the task on its own then uses this model to reason about other agents.

A natural concern is that the human is unlikely to be similar to the robot; indeed, recent work has shown that (model-free) RL agents that are trained to respond to self-policies do not work well with actual humans [7]. The ways in which the human and robot perceive the world and make decisions are likely to differ. Here, we aim to isolate and learn these differences in a sample-efficient manner.

We argue that latent state-space models obtained via deep reinforcement learning are well-suited for this purpose — modern deep variants [8] can handle multiple high-dimensional sensory modalities, capture complex dynamics, and are useful for robot decision-making, yet are sufficiently modular to permit structural interpretability [9], [10]. We exploit this modularity and strategically place *learnable implants* that capture differences in perception and/or policy (Fig. 1.B.). These implants can be small and relatively unstructured (e.g., neural networks) or specified using prior knowledge about the human (e.g., known cognitive biases or perceptual limitations) with associated parameters that can be quickly optimized via gradient-based learning.

We call our framework **Model Implants for Rapid Reflective Other-agent Reasoning/Learning (MIRROR)**; an allusion to the mirror neurons in human brains that are hypothesized to play a role in social projection [11], [12]. Returning to our example, our MIRROR-enabled robot highlights the red car in front and verbally informs the user about yellow car at the rear. To avoid distracting the user, it chooses *not* to tell the human about the green car, whose location and velocity has little impact on the human’s decision and safety.

These communicative actions are the result of *planning* over the implant-augmented self-model and the robot’s internal model (Fig. 4). Specifically, perceptual implants in the MIRROR model were used to capture that the human had limited visibility, and could not see highlighted cars in the rear (but could be notified about them via auditory means). To plan forward, the robot samples possible future trajectories using *both* its self-model and the human model, which are coupled together via generated observations and predicted actions (Fig. 1.B). Unlike recent work (e.g., [1], [3], [13], [14]), planning with structured deep *multi-modal* self-models allows the robot to take into account longer-term behavior and choose among rich communication modalities, without having to hard-code environmental or human properties (e.g., belief dynamics).

Experiments in three simulated domain show that MIRROR is able to learn better human models faster compared to behavioral cloning (BC) and a state-of-the-art imitation learning method [15]. We found the MIRROR models were more robust when planning during transfer (e.g., from clear to foggy weather). In addition, we report on a human-subject study ( $n = 21$ ) using the CARLA simulator [16], which reveals that MIRROR provides useful assistive information, enabling participants to complete a driving task with fewer collisions in adverse visibility conditions.

In summary, this paper makes the following contributions:

- MIRROR, a sample-efficient framework for learning human models using deep self-models for initial structure;
- A planning-based communication approach that leverages learned world dynamics and human models;
- Findings from a human-subject study in the assistive driving domain, showing that MIRROR provides useful communication that improves task performance.

We believe MIRROR is a step towards data-efficient human models for human-robot interaction; MIRROR demonstrates how deep representation learning during RL can be combined with demonstrations for human modeling and planning. MIRROR can be used for human-robot communication during robot tele-operation and shared-control settings, and it opens up an alternative path to human-robot collaboration with deep models. To motivate future work, we discuss current limitations and potential research avenues in the concluding section of this paper.

## II. BACKGROUND & RELATED WORK

MIRROR builds upon the existing literature on human modeling and human-robot communication. Due to its importance, the field of agent communication is large; here, we briefly summarize closely-related work that learn and use human models for human-robot/AI interaction and communication [7], [17]–[19].

**Human Models for HRI.** In this work, we focus on model-based methods that explicitly model human behavior. Compared to model-free approaches to HRI, model-based methods tend to make reasonable predictions with far less data [18]. Model-based methods can be “black-box” in that they make few assumptions about the human and focus on learning a policy function. For example, recent work learns a human policy via imitation learning, followed by a residual policy for shared control [20]. In contrast, Theory of Mind (ToM) models incorporate (possibly strong) assumptions about how humans perceive the world and make decisions. For example, a ToM model may assume people are rational and learn in a Bayesian manner [21]–[23], which is generally not true [22].

MIRROR can be seen as a hybrid approach that scaffolds human model learning using the robot’s own internal model (obtained using RL). Compared to standard black-box human models, MIRROR provides additional structure that can ease data requirements. Compared to handcrafted ToM approaches [2], MIRROR is able to handle high-dimensional

observations. MIRROR is related to recent approaches that focus on capturing human traits, e.g., biases under risk and uncertainty [24] or action errors due to misunderstood environmental dynamics [25]. However, these approaches typically build on top of hand-crafted ToM models.

**Assistance via Human-Robot Communication.** Enabling robots to communicate with humans has had a long history; early robots in the 1990s (e.g., Polly [26] and RHINO [27]) were simple stimulus-response systems. In contrast, modern day robots leverage learning and planning to generate a variety of communication patterns, e.g., legible motion [28]–[30], and natural language [31], [32].

Recent work has shown that human-robot communication can improve human task performance [1], [33], explain robot errors [13] and calibrate human-robot trust [34]. However, these approaches typically use hand-specified human models and known environment models. In contrast, MIRROR plans communication actions using *learned* models. MIRROR is related to prior work on personalized assistive navigation [35], but the mechanism differs: MIRROR adapts implant parameters whilst [35] uses a mixture of expert models.

MIRROR is closely-related to Assistive State Estimation (ASE) [3] in that both approaches augment user observations to communicate state information. However, there are crucial differences: ASE assumes known dynamics and perceptual models to compute the near-optimal human policies, while MIRROR uses learned dynamics and implant models. In addition to differences in the human model, ASE generates observations that minimize the KL divergence between the (predicted) user’s beliefs and the assistant’s beliefs. Instead, MIRROR forward simulates possible futures using its internal models, and plans communication to maximize task rewards while minimizing communication costs.

### III. MODEL IMPLANTS FOR RAPID REFLECTIVE OTHER-AGENT REASONING/LEARNING (MIRROR)

Our problem setting is one of assistance: a (human) user is acting in a partially-observable environment to maximize rewards. The assistive robot’s goal is to help the user achieve their objective. The robot may receive different observations from the environment and can modify the user’s observations to provide additional information. We seek to derive an effective assistant. At a high level, we will imbue the robot with a structured model of the human that can be adapted with data. After learning, the robot plans using its own internal (self) model and the human model to communicate valuable information. We first detail the robot’s underlying self-model, then describe the human model (specifically the implants), and finally, how communication can be achieved using both the self and human models.

#### A. Self Model: Multi-Modal Latent State-Space Model

**Model Structure.** In MIRROR, the robot’s self-model is a multi-modal state-space model (MSSM) [10] (Fig. 2). Intuitively, the MSSM models an agent that is sequentially taking

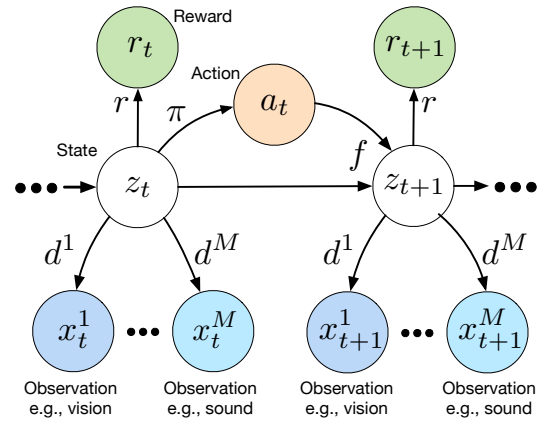


Fig. 2: MIRROR’s self-model is a multi-modal latent state-space model (MSSM). In the above, circle nodes represent random variables and shaded nodes are observed during learning.

actions in a world and receiving rewards and multi-modal observations. The observation  $x_t^m$  at time  $t$  for  $m = 1, \dots, M$  sensory modalities are generated from the latent state  $z_t$ . The model assumes Markovian transitions where the next state is conditioned upon the current state and the action  $a_t$  taken by the agent. Upon taking an action, the agent receives reward  $r_t$ . In general, the reward can also be conditioned upon the action and next state. Given the graphical structure in Fig. 2, the MSSM’s joint distribution factorizes as:

$$p_\theta(x_{1:T}^{1:M}, r_{1:T}, z_{0:T}, a_{1:T-1}) = \prod_{t=1}^T \left[ \prod_{m=1}^M p_\theta(x_t^m | z_t) \right] p_\theta(r_t | z_t) p_\theta(z_t | z_{t-1}, a_{t-1}) p_\psi(a_{t-1} | z_{t-1}) p(z_0) \quad (1)$$

where  $\theta$  are the world model parameters and  $\psi$  are the policy parameters,  $x_{1:T}^{1:M}$  denotes all observations from  $t = 1, \dots, T$ , and similarly for  $z_{0:T}$ ,  $a_{1:T-1}$ , and  $r_{1:T}$ . In this work, each of the factorized distributions are modelled using nonlinear functions:

- Transitions:  $p_\theta(z_t | z_{t-1}, a_{t-1}) = p(z_t | f_\theta(z_{t-1}, a_{t-1}))$
- Observations:  $p_\theta(x_t^m | z_t) = p(x_t^m | d_\theta^m(z_t))$
- Rewards:  $p_\theta(r_t | z_t) = p(r_t | r_\theta(z_t))$
- Policy:  $p_\psi(a_t | z_t) = p(a_t | \pi_\psi(z_t))$

where  $f_\theta$ ,  $d_\theta^m$ ,  $r_\theta$ , and  $\pi_\psi$  are neural networks. Note the slight abuse of notation; we denote both the reward function and the reward random variable as  $r$ , where the meaning should be clear from context. Depending on the application, the policy  $\pi_\psi$  may be deterministic or stochastic. We write  $\pi_\psi(a_t | z_t) = p_\psi(a_t | z_t)$  to refer to the latter case.

**Dynamics Learning.** Given trajectories of the form  $\tau = \{(x_t^{1:M}, a_t, r_t)\}_{t=1}^T$ , we seek to learn the parameters  $\theta$ . Because maximum likelihood estimation is intractable in this setting, we optimize the evidence lower bound (ELBO) under the data distribution  $p_d$  using a variational distribution  $q_\phi$  over the latent state variables  $z_t$ ,

$$\mathbb{E}_{p_d}[\mathcal{L}_e] \leq \mathbb{E}_{p_d}[\log p_\theta(x_{1:T}^{1:M}, r_{1:T} | a_{1:T-1})] \quad (2)$$

where

$$\mathcal{L}_e = \sum_{t=1}^T \left( \mathbb{E}_{q_\phi(z_t)} \left[ \sum_{m=1}^M \log p_\theta(x_t^m | z_t) \right] + \mathbb{E}_{q_\phi(z_t)} [\log p_\theta(r_t | z_t)] - \mathbb{E}_{q_\phi(z_{t-1})} [\mathbb{D}_{\text{KL}}[q_\phi(z_t) \| p_\theta(z_t | z_{t-1}, a_{t-1})]] \right) \quad (3)$$

The first two terms in the ELBO are reconstruction terms and the Kullback-Leibler (KL) divergence term encourages consistency between the variational distribution and the transition dynamics. In this work,  $q_\phi$  is an inference network,

$$q_\phi(z_t | z_{t-1}, a_{t-1}, x_t^{1:M}) = \mathcal{N}(g_\phi^\mu(z_{t-1}, a_{t-1}, x_t^{1:M}), g_\phi^\Sigma(z_{t-1}, a_{t-1}, x_t^{1:M})) \quad (4)$$

where  $g_\phi^\mu$  and  $g_\phi^\Sigma$  are neural networks that give the mean and the covariance of the Gaussian latent state variable, respectively. When observations are missing, we apply a simple masking operation similar to prior work [36], but alternative approaches such as a product-of-experts [37] can be used as well (with corresponding changes to the inference network structure). The inference network  $q_\phi$  will continue to play an important role in MIRROR to facilitate fast inference and planning during communication.

**Policy Learning.** Given the latent states  $z_t$  sampled from the inference network, we leverage RL to learn optimal policies. Our approach is based on Stochastic Latent Actor Critic [8]. In our human-subject study involving the assistive driving task, we train a critic network  $Q_\psi(z_t, a_t)$  and an actor network  $\pi_\psi(a_t | z_t)$  using Soft-Actor Critic (SAC) [38]. In our simulated experiments involving gridworld environments, we train a Q-network  $Q_\psi(z_t, a_t)$  using Deep Q-Learning (DQN) [39].

**Practical Aspects.** We find that eventual human model performance is significantly improved when the self-model is trained with data-augmentation [40] and random dropouts across the sensory modalities [41], [42]. For example, we randomly drop LIDAR range readings in each training batch when training our robot in CARLA. We hypothesize that to generalize appropriately, the model should be trained with data that approximates the different ways observations can appear to a human. As an added bonus, this training method often resulted in more robust internal world models and policies.

### B. Learning Human Models via MIRROR

In MIRROR, the human model is identical to the robot’s self-model *except* for implants that are injected to change the model’s behavior. To distinguish variables in the human and robot models, we will use the superscript H and R to refer to the human or robot, respectively. For example, the human action at time  $t$  is  $a_t^H$ . Furthermore, we will denote generated/predicted variables by a hat accent, e.g.,  $\hat{a}_t^H$ , and variables that are changed by an implant with a check accent, e.g.,  $\check{a}_t^H$ .

**Model Implants.** Once the self-model is trained, we augment it with implanted functions  $h(\cdot)$ . In this work, we examine two implant classes:

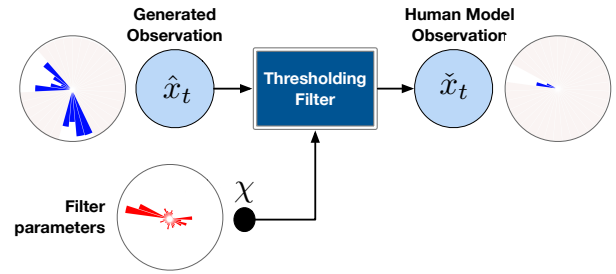


Fig. 3: An Example Thresholding Filter as a Perceptual Implant. The (generated) range observations (blue bars) are passed through a thresholding filter to eliminate observations that are beyond a parameterized distance (red bars) in each segment from the human.

- **Perceptual implants** model how humans perceive the world by changing the observation  $x_t$ . In several of our experiments, we use a threshold mask/filter implant that models that the human can only see objects within a certain range (Fig. 3)
- **Policy implants** model how the human acts differently from the robot by changing the policy. Simple implants can simply add noise to increase policy entropy. In this work, we add a state-dependent residual  $\delta(z_t^H)$  term to the parameters of the policy distribution  $\pi$ . As an example, for the Gaussian policies  $\mathcal{N}(\mu(z_t^H), \sigma^2(z_t^H))$  used in our CARLA experiments, we add  $\delta(z_t^H)$  to the mean of the Gaussian such that the augmented policy distribution follows  $\mathcal{N}(\mu(z_t^H) + \delta(z_t^H), \sigma^2(z_t^H))$ . For our discrete policy  $\text{Cat}(K, \mathbf{p})$ , we add  $\delta(z_t^H)$  to the distribution parameters  $\hat{\mathbf{p}} = \sigma(\mathbf{p} + \delta(z_t^H))$ , where  $\sigma$  is softmax function. The residual is parameterized by a small neural network.

The above provides a flavor of what is possible and is not exhaustive. For example, to model humans who are slow to change their beliefs, we could implant a low-pass filter  $z_{t+1}^H = \alpha z_t + (1 - \alpha)f(z_t, x_t, a_t)$  where  $\alpha$  is a learnable parameter. We leave exploration of other implants to future work.

**Learning Implant Parameters.** Given an implant  $h_\chi$  parameterized by  $\chi$ , we can learn  $\chi$  by minimizing the following loss given data:

$$\arg \min_{\chi} \mathcal{L}(\chi) = - \sum_{t=1}^T \mathbb{E}_{z_t \sim p(z_t | x_{1:t}^{1:M}, a_{1:t-1}, \chi)} [\log \pi(a_t | z_t)] - \lambda \log p(\chi) \quad (5)$$

where  $\lambda$  is a regularization hyperparameter that controls the strength of an optional prior. Intuitively, this loss optimizes the likelihood of observing a human’s actions given the self-model and the implant parameters. In our work, we approximate  $p(z_t | x_{1:t}^{1:M}, a_{1:t-1})$  using the learnt inference network and perform stochastic gradient descent by sampling  $z_t$ . Note that *only* the implant parameters are modified when learning from human data.

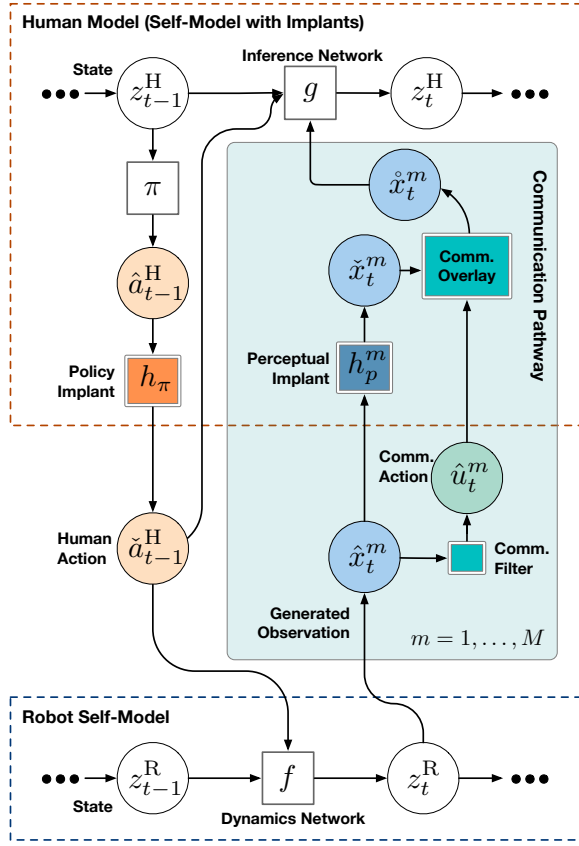


Fig. 4: Human-Robot Communication via Forward Simulation. In brief, the robot’s self-model is used to simulate the environment and the human model is used to simulate actions. By leveraging the learnt dynamics network  $f$  and inference network  $g$ , the two models are coupled and used to “imagine” possible futures. The communication pathways (in teal) serves to link the robot’s communication actions  $\hat{u}_t^m$  (filtered generated observations) to the human model’s observations. The multi-modal models support  $M$  potential communication pathways and the robot may choose one or more of these pathways at any given time step. Please see text for additional details.

### C. Human-Robot Communication with MIRROR

Next, we turn our attention to how the implanted self-model can be used for human-robot communication. The key idea is to plan using both learned models; we couple the robot’s self-model and the human model together via generated robot observations and communication actions, and predicted human actions (Fig. 4).

**Communication Pathways.** To enable the robot to influence the human, we must first establish a means of communication between the two agents. Similar to [3], we assume the robot can change the human’s observations of the world, e.g., by providing some visual cues or alerting the user verbally. A communication pathway for modality  $m$  is illustrated in Fig. 4 (teal block) where the human’s observation  $\tilde{x}_t^m$  and the robot’s communication action  $\hat{u}_t^m$  are combined via an overlay function to yield a “combined” observation  $\hat{x}_t^m = v(\tilde{x}_t^m, \hat{u}_t^m)$ . As an example, for images, we can replace selected pixels of  $\tilde{x}_t^m$  with corresponding pixels from  $\hat{u}_t^m$ .

Communication actions  $\hat{u}_t^m$  are actually observations that are generated from the robot belief state  $\hat{x}_t^m \sim p(\hat{x}_t^m | d^m(z_t^R))$  and passed through a communication filter; for example, we can use simple masks filters (parameterized by  $\omega_t$ ) to filter away irrelevant information about the robot’s belief. Note that a communication action can be generated for each modality  $m$ , which offers the robot multiple ways to communicate with the human.

**Forward Imagination.** Given our robot model, human model, and communication pathway, we now seek to sample future trajectories. We begin at time-step  $t-1$  where we have samples from the robot’s belief over the world state  $z_{t-1}^R$  and the (predicted) human belief  $z_{t-1}^H$ . We take a step from  $t-1$  to  $t$  to obtain  $z_t^R$  and  $z_t^H$ :

- 1) Sample the human’s action,  $\hat{a}_{t-1}^H \sim h_\pi(\pi(z_{t-1}^H))$ .
- 2) Forward simulate using the world dynamics  $z_t^R \sim p(z_t^R | f(z_{t-1}^R, \hat{a}_{t-1}^H))$ ;
- 3) Generate the observations  $\hat{x}_t^{1:M} \sim p(\hat{x}_t^{1:M} | d^m(z_t^R))$ ;
- 4) Obtain the human observation using the perceptual implants  $\tilde{x}_t^m = h_p^m(\hat{x}_t^m)$  across the modalities;
- 5) Obtain the communication action using the communication filter  $\hat{u}_t^m = w_\omega(\hat{x}_t^m)$  for each modality  $m$ ;
- 6) Combine the communication action and the observation to obtain  $\hat{x}_t^m = v(\tilde{x}_t^m, \hat{u}_t^m)$ ;
- 7) Finally, sample the human’s belief state using the inference network  $z_t^H \sim q(z_t^H | z_{t-1}^H, \hat{a}_{t-1}^H, \hat{x}_t^{1:M})$ .

Given the human’s action (the observed action if available or generated by the human policy), the forward dynamics function provides the next state distribution  $z_t^R \sim p(z_t^R | f(z_{t-1}^R, a_{t-1}^H))$ . Given  $z_t^R$ , we generate observations  $\hat{x}_t^m$  for modalities  $m = 1, \dots, M$ , which are modified by the perceptual implants to yield  $\tilde{x}_t^m$ . The modified observations  $\hat{x}_t^{1:M}$  are used to update the human’s beliefs using the inference network:  $z_t^H \sim q(z_t^H | z_{t-1}^H, \hat{a}_{t-1}^H, \hat{x}_t^{1:M})$ . This process can be iterated to forward propagate the belief states up to a future horizon  $T$ .

**Planning for Communication.** Given the forward simulation and communication pathways above, the robot can optimize communication to maximize task rewards while minimizing costs:

$$\arg \max_{\omega_{0:T}} J = \mathbb{E}_{z_{0:T} | \omega_{0:T}} \left[ \sum_{t=0}^T \gamma^t (r(z_t) - C(\omega_t)) \right] \quad (6)$$

where  $\omega_{0:T}$  is the parameters of our communication filter for time-steps 0 to  $T$ ,  $\gamma$  is the discount factor,  $r$  is the task reward function<sup>1</sup>, and  $C$  is the cost function. The expectation is taken with respect to the trajectories  $\tau$  under the models and filter parameters  $\omega_{0:T}$ . Various methods can be used to optimize  $J$ . In our work, we use the cross-entropy (CE) method [43], and re-plan at each time-step. Note that real observations are obtained after each step, which are used to update the beliefs of the robot and human models.

<sup>1</sup>In principle, this task reward function may differ from the reward function that the model was trained with, but we leave such experiments to future work.

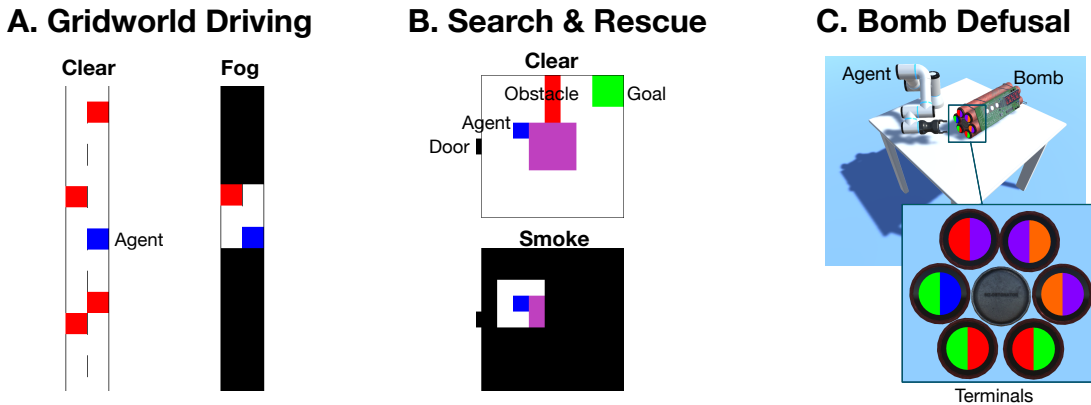


Fig. 5: Experimental domains. (A) Gridworld Driving where the blue vehicle is moving on a road at constant speed and has to avoid the other red vehicles. In the Fog setting, visibility is reduced; the black region indicates areas not visible to the human. (B) A Search-&Rescue task where the agent (blue box) starts at the door and is tasked to rescue a victim at the green goal and bring them back to the door. The obstacle in red can appear in either the top or the bottom path, and the victim’s position is randomly initialized in one of three potential positions. In the Smoke variant, visibility is reduced to a small region around the human (C) A Bomb Defusal game where a teleoperated robot has 15 seconds to disarm the bomb by pressing three buttons (one in each stage). The correct button at each stage depends on six visible “terminals” (which change after each button press), the bomb type (not visible to the human, but detectable by the robot) and game rules. The rules differ slightly between the robot training environment and the test environment. The human, who has access to the updated rules, has to confirm the robot’s selection. This domain features asymmetric sensing and knowledge/policy between the human and robot. In all three domains, the assistive agent/robot has to communicate relevant information to help the human agent complete the task successfully.

#### IV. EXPERIMENTS

In the following, we describe experiments with simulated humans (created using real human data) in a simplified driving task, a search-&-rescue, and a bomb defusal game (Fig. 5). Using simulated humans allowed us to compare the learned models against ground truth models with varying amounts of training data, and enabled us to test communication. Our primary hypothesis was that leveraging the self-model and learning the implant parameters allows MIRROR to better model human behavior using fewer samples. In particular, we posit that using MIRROR results in better communication during transfer (e.g., from clear weather to dense fog/smoke where visibility is reduced). We focus on the main results and relegate details (e.g., domain parameters, network architecture) to the Appendix. Source code is available at <https://github.com/clear-nus/mirror>.

##### A. Experimental Setup

**Domains and Communication Modalities.** Our setup comprised a primary human agent who is attempting to complete a task, and an assistive robot who is able to reveal information via communication. Figure 5 summarizes the three domains used in our experiments. The Driving task is similar to those used in prior work in HRI (e.g., [18]) where state information is available to the agent. The Search-&-Rescue domain is a more complex gridworld environment, where the robot can only observe raw image data and textual descriptions of the position of the victim and obstacles. In these two domains, there is a “transfer” setting (fog/smoke) where the human’s visual perception is degraded, i.e., it doesn’t receive any state information beyond its field of view. For the Bomb Defusal

game, the transfer setting is different in that the robot’s performance is diminished; the rules dictating which button to press is different from the rules it was trained with. As such, the communication explains the robot’s choice and the human has to verify.

In all three domains, there are two communication modalities: visual and verbal. The total cost for each modality is quadratically related to the number of items communicated. Verbal communication is more costly, but can reveal information that is not visually apparent (e.g., vehicles in the rear or the bomb type). For more information, please see Table I in the Appendix.

**Simulated Humans.** For each domain, we created simulated human agents using human data; we collected data from 10 real humans (recruited from our university community) playing multiple rounds (24 to 36 rounds depending on the domain). The data from each person was used to train an agent model for that individual. For example, in the Driving domain, we used Maximum Entropy IRL [44] to learn a reward function from the collected trajectories. Similar to prior work [18], we use as features (i) the distance to the center of the road, (ii) distances to the other cars, and (iii) the driver’s action. Given the learnt reward function, the simulated human agent perceives state information (up to its perceptual capabilities) and plans at each time-step to maximize its discounted cumulative reward up to a specified horizon. Further details about the other simulated agents are in Sec. VII-B of the Appendix.

**Compared Methods.** In total, we compare five different human modeling methods:

- Ideal Model (IM): This baseline uses the same model as

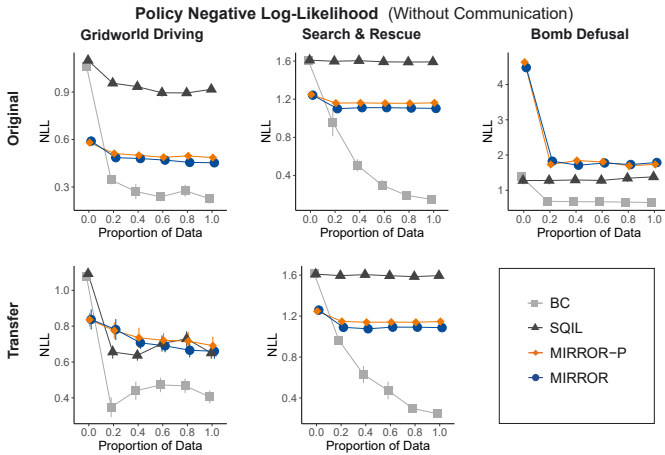


Fig. 6: Policy Negative Log-Likelihood (NLL) in the different domains after training with different proportions of data. Lower scores indicate better fit to the data. In both the original and transfer settings, BC achieved the best scores. However, this did not translate into good performance during communication. See main text for details.

the simulated human and thus, represent “ideal” models that are typically unavailable in practice. These serve as an upper-bound on performance.

- Behavioral cloning (BC): a black-box policy learnt via supervised learning on observed trajectories.
- Soft Q-Imitation Learning (SQIL) [15]: A state-of-the-art method that trains a policy via RL to match the human demonstrations. SQIL represents a class of imitation learning methods that have access to the environment<sup>2</sup>.
- MIRROR: With a perceptual implant and a policy implant. See Appendix (Sec. VII-C) for implant details.
- MIRROR-P: A MIRROR variant with the perceptual implant only.

It is important to note that BC, SQIL, and IM are *aware* of what the human can perceive; we provided the ground truth human observations. This is *not* the case for the MIRROR models that perceive what the robot observes and had to learn what the human could observe. The MIRROR self-models were trained to perform the task in the original domain, and only the implants were adapted using the simulated human data (in both the original and transfer settings).

To train BC and SQIL in the transfer domains, we combined both the original and transfer data, e.g., if 10 trajectories were used to train MIRROR, we used 20 trajectories (10 original, 10 transfer) for BC and SQIL. This was necessary to obtain reasonable performance for these methods. Each method was used to train a neural network policy (comprising GRUs with fully-connected layers) with early-stopping on a validation set. We varied the number of layers and report best results for the baseline models.

<sup>2</sup>In preliminary experiments, we also tested GAIL [45] but SQIL outperformed GAIL on all our tests.

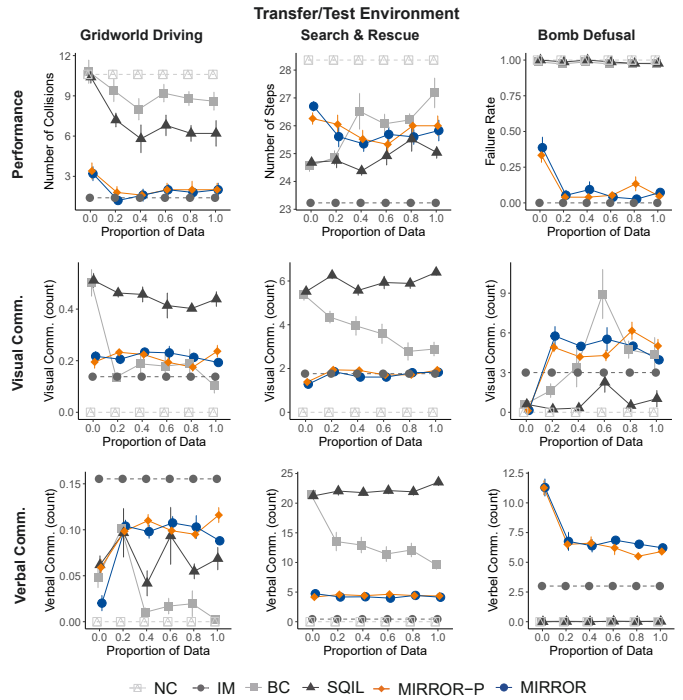


Fig. 7: Task Performance and Communication Amount in the Transfer/Test Environment. NC indicates agents that received no communication. In all domains, the simulated human’s performance improved given communication. For the Search & Rescue Task, the SQIL agents slightly outperform the MIRROR agents but at far higher communication cost. In the Driving and Bomb Defusal tasks, the best performing simulated agents were those that were paired with the assistive MIRROR agents. The performance was achieved with reasonable amounts of communication, close to that of the ideal models.

## B. Results and Analysis

Overall, we find that MIRROR is able to learn better human models with fewer samples compared to BC and SQIL. We initially compared the negative log-likelihood (NLL) of the trained policies (on a test set) with different amounts of training data (Fig. 6). Interestingly, BC achieved the best NLL scores, indicating that the BC models best fit the data and may be good proxy human models for planning. However, these scores are misleading as they fail to capture a model’s robustness to input changes that invariably arise during planning.

The MIRROR models were far more effective than BC for communication planning. Fig. 7 summarizes the performance of the simulated human agents with communication from the compared methods, along with the amount of communication provided. In all domains, the agents that received communication outperformed the agents that did not (NC). For the Gridworld Driving and Bomb Defusal tasks, the simulated humans achieved significantly better performance scores with reasonable amounts of communication from the MIRROR assistants. MIRROR learned rapidly—good models could be obtained with 20% of the data (4 trajectories). For the Search-&-Rescue task, SQIL obtained the best performance scores (about 1-2 steps better than MIRROR), but only by using large amounts of verbal communication. Qualitatively,



Fig. 8: CARLA Experiment Setup. (A) The stretch of highway that participants drove along. (B) Participants drove the simulated car using a steering wheel with accelerator and brake pedals. (C) and (D) show the difference in visibility in clear and foggy weather. Both cars are visible in the clear setting. In the fog setting, the car on the left is visible, but the car in the front can barely be seen.

the SQL agent would reveal almost the entire map, but with no discernible pattern to the communication. Moreover, providing more training data did not reduce the verbal communication provided. On the other hand, the MIRROR models revealed information more selectively, often focusing on the locations of the goal item and the obstacle (See Appendix Sec. VII-F for more information).

## V. HUMAN-SUBJECT EXPERIMENTS

In this section, we report on a second set of experiments designed to test whether MIRROR is able to provide useful information to human users in a more realistic setting. We use CARLA [16], a modern driving simulator (Fig 8). Unlike the previous gridworld experiments, the CARLA environment has continuous state and action spaces, with realistic dynamics and visuals. Our main hypotheses were that (H1) planning with implanted self-models would yield more helpful communication than planning with behavioral cloning models and (H2) planning to optimize task rewards and communication costs would lead to less redundant communication compared to belief matching [3]. Our study was approved by our institution’s ethical review board.

### A. Experimental Setup

**Task Description.** Participants interacted with our assistive driving agents in a highway driving task under dense fog (see Fig. 8.C and 8.D for a comparison between clear and adverse weather conditions). The goal was to drive along a stretch of CARLA’s Town04’s two-lane carriageway from the starting position to the destination, while navigating through a normal highway traffic. The other vehicles may slow down or speed up, and due to poor visibility, would not be able to avoid

the participant’s vehicle. The participant had to actively avoid other vehicles along the way.

**Assistive Communication.** The car is equipped with a semantic LIDAR and a driving assistant that can provide both visual and verbal cues. Specifically, the agent could highlight selected vehicles through visual bounding boxes and/or provide informative speech (as previously shown in Fig. 1.A). The visual bounding boxes can be generated when the other cars enter the range of LIDAR detection range. Verbal communication comprised of free text generated by OpenAI’s GPT-2 language model [46] (fine-tuned for our task). The system is capable of informing participants when a car approaches or slows down, as well as when no cars are detected in a specific direction. The robot is able to observe 36 LIDAR beams along three angles (108 beams in total), the velocity of the ego car, distances to the center of the lane, and the relative curvature of the road 10 meters ahead.

**Compared Methods.** Based on our previous experiments, we trained MIRROR with both perceptual and policy implants. We compared four conditions:

- No Communication (NC): the human receives no assistance.
- Behavioral Cloning (BC): the human model is trained using behavioral cloning. The model is provided with both expert demonstrations in clear weather, and participant data in the foggy weather.
- MIRROR: our MIRROR method that first trains a self-model via RL, then learns the implants from demonstrations, and plans communication.
- MIRROR-KL: a MIRROR variant that does not plan but minimizes the KL-Divergence between the human’s mental state and the robot’s belief, similar to [3]. Unlike [3], MIRROR-KL uses a learnt implant model and is capable of multi-modal communication.

We were unable to use SQL due to experimental time constraints; training on the demonstrations with SQL required several hours.

**Participants.** A total of 21 participants (mean age = 23.2, 11 males and 10 females) were recruited from the university community. The experiment was designed to be within-subjects with all 21 participants in each condition.

**Procedure.** Participants entered the lab and were briefed about the task. They then engaged in two practice trials; the first trial involved driving freely along the highway in clear weather conditions, and the second trial involved three rounds of the driving task under dense fog conditions without any assistive communication. Thereafter, they performed a total of 24 rounds, with the first 6 rounds without any assistive communication, followed by 18 rounds with three different agents (six rounds per agent). The data from the first 6 rounds were used to train the models and the order of the three agents was counterbalanced. Participants could choose to take a one minute rest after every 6-th round to reduce fatigue.






Fig. 9: Objective Measures. Error bars indicate one standard error. (A) Number of Collisions. (B) Amount of Visual Communication. (C) Amount of Verbal Communication (D). MIRROR results in fewer collisions compared to behavioral cloning (BC) and no communication (NC) conditions, but with far less communication required compared to belief matching (MIRROR-KL).

**Dependent Measures.** We use both objective and subjective measures to evaluate agent performance. Objective measures comprised the number of collisions with other vehicles/environment, and the amount of communication (e.g., the number of speech utterances). We also collected a range of subjective measures to ascertain cognitive load, communication properties (helpfulness, redundancy, timeliness, modality selection), and trust after interaction with each agent (See Table III in the Appendix).

## B. Results and Analysis

In brief, the results support both hypotheses; we give a summary of key findings below and place additional details (e.g., breakdown across NASA-TLX dimensions) in the Appendix. We compared the methods along both objective and subjective dimensions using a repeated measures one-way ANOVA, followed by selected pairwise  $t$ -tests with adjusted- $\alpha = 0.0167$  using Bonferroni correction.

Fig. 9.A shows that participants experienced significantly fewer collisions when interacting with MIRROR ( $F_{3,60} = 47.977$ ,  $p < 0.001$ ; MIRROR vs BC:  $t(9.5) = 7.683$ ,  $p < 0.001$ ). Subjectively, participants found MIRROR provided information that was more helpful and timely compared to BC, and were also more comfortable with the communication modality chosen ( $p < 0.005$  across the measures and pairwise tests, Fig. 14 in the appendix). Participants also trusted the MIRROR agent more than BC ( $F_{3,60} = 47.730$ ,  $p < 0.001$ ; MIRROR vs BC:  $t(9.5) = -12.526$ ,  $p < 0.001$ ). The overall Raw-TLX scores indicated that the participants felt less mentally burdened when they interacted with MIRROR ( $F_{3,60} = 10.071$ ,  $p < 0.001$ ; MIRROR vs BC:  $t(9.5) = 6.658$ ,  $p < 0.001$ ).




Fig. 10: Samples of Learned Perceptual Implants (threshold filters) for the CARLA experiment. Top is the front of the vehicle. Length of red bars indicate visibility distance. The implants indicate the human was not able to see far ahead or in the rear but could see cars at the side. Compare against Fig. 8.C and 8.D.

Taken together, both objective and subjective evidence strongly support hypothesis **H1**. This finding is corroborated by participant survey responses; they shared that MIRROR “conveys critical information at a good timing” in contrast to BC, which they felt “is not helping me at all” and “doesn’t inform me about the cars that are approaching from the rear”. Qualitatively, we found planning with the BC human model to be inaccurate; the BC model would quickly overfit, which led to poor communication. In contrast, the MIRROR implant models resulted in better communication; Fig. 10 shows the perceptual implants learned by MIRROR to well approximate what the human could see, even with a small amount of training data (six demonstrations).

Next, we turn our attention to **H2**, i.e., whether planning with task rewards and communication costs reduced redundant communication relative to belief matching. Figs. 9.B and 9.C show how long each agent highlighted cars and the number of times they verbally alerted the driver, respectively. MIRROR-KL tended to be overly communicative; it provided more than 5 times more visual and verbal communication compared to MIRROR without significant benefits in terms of task performance. Subjectively, participants rated MIRROR-KL to provide more redundant information compared to MIRROR ( $t(9.5) = -11.411$ ,  $p < 0.001$ ). In their survey responses, participants wrote that the MIRROR-KL agent “is too talkative”, “told me a lot of useless information” and “is distracting and annoying”. In comparison, they found MIRROR to be “straight to the point.” Recall that the only difference between the two is the objective function; MIRROR-KL tries to align belief distributions, regardless of whether the alignment leads to better task accomplishment. We observed that MIRROR-KL would communicate verbally even when there was no need to, e.g., it would repeatedly tell participants that “there is no car in the rear”.

When asked which agent they were most comfortable with, a majority of the participants (14 out of 21) selected MIRROR. The remaining participants picked MIRROR-KL. When asked about which agent they would prefer for long-distance driving (above 1 hour), the number of participants selecting MIRROR increased to 18. No participant selected BC. Interestingly, some participants preferred MIRROR-KL’s talkative nature, with one participant stating that they “felt annoyed and safe at the same time”. A few participants found MIRROR to be too quiet: “I have some confidence that it works but

*I'm not entirely sure because it is quieter*". These responses suggest individual preferences for information/reassurance and differing trust in the system. Incorporating these aspects within MIRROR would make for interesting future work.

## VI. CONCLUSION

In summary, we present MIRROR, a framework for learning human models using deep self-models for initial structure, along with a planning-based communication approach that couples the human model with learned world model. Experiments show MIRROR to be effective, outperforming existing behavioral cloning and imitation learning methods. The results show that bootstrapping human model learning with latent-variable models learnt during reinforcement learning leads to generalizable models that are more useful for interaction planning.

More broadly, we consider MIRROR to be a step towards more data-efficient human models for human-robot interaction. The key idea examined in this work—learning differences from the robot self-model—can be potentially be applied to general human-robot collaboration. Compared to existing work on human models, MIRROR embodies an alternative paradigm that “front-loads” the learning the environmental dynamics and task structure, and thus, offers savings both in terms of sample complexity and computation when learning from human demonstrations. Below, we highlight potential areas for future work:

- In our current setup, the agent and human are not co-present in the environment. We believe MIRROR can be adapted for these situations by incorporating observations/models of other agents acting in the world. Related work involves crossing the Sim2Real gap, which remains a challenging problem for deep RL.
- Intuitively, MIRROR’s effectiveness depends on similarity between the human and the robot self-model; drastic differences will render the self-model ineffective as a reference point. We are working on how to characterize this intuition theoretically and estimating how effective MIRROR is given varying amounts of data.
- A related open question is whether the human model is identifiable within our framework. Similar to IRL, we posit that the differences in the policy and perception cannot be completely ascertained given a dataset. Future work can look into methods that attempt to address this ambiguity, e.g., by maintaining a distribution over parameters [47] or via Bayesian Optimization [48]. A related issue is how to handle alternative policies and norms that achieve the same objective.
- As our human subject experiments reveal, some human traits cannot be initially captured within a RL self-model. How we can incorporate elements such as human-robot trust [23], [49], [50], emotions, and social norms remains a key open problem.

We believe that building upon MIRROR forms a compelling pathway towards robots that can interact fluently with humans.

## ACKNOWLEDGEMENTS

This research is supported by the National Research Foundation Singapore under its AI Singapore Programme (AISG Award No: AISG-RP-2019-011).

## REFERENCES

- [1] A. Tabrez, S. Agrawal, and B. Hayes, “Explanation-based reward coaching to improve human performance via reinforcement learning,” in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 249–257.
- [2] M. C. Buehler and T. H. Weisswange, “Theory of mind based communication for human agent cooperation,” in *2020 IEEE International Conference on Human-Machine Systems (ICHMS)*, 2020, pp. 1–6.
- [3] S. Reddy, S. Levine, and A. D. Dragan, “Assisted perception: optimizing observations to communicate state,” in *Conference on Robot Learning (CoRL 2020)*, 2020.
- [4] F. H. Allport, *Chapter 13: Social Attitudes and Social Consciousness*, ser. Social Psychology. Houghton Mifflin Company, 1924.
- [5] J. I. Krueger, T. E. DiDonato, and D. Freestone, “Social projection can solve social dilemmas,” *Psychological Inquiry*, vol. 23, no. 1, pp. 1–27, 2012.
- [6] J. I. Krueger, “Social projection as a source of cooperation,” *Current Directions in Psychological Science*, vol. 22, no. 4, pp. 289–294, 2013.
- [7] M. Carroll, R. Shah, M. K. Ho, T. Griffiths, S. Seshia, P. Abbeel, and A. Dragan, “On the utility of learning about humans for human-ai coordination,” *Advances in Neural Information Processing Systems*, vol. 32, pp. 5174–5185, 2019.
- [8] A. Lee, A. Nagabandi, P. Abbeel, and S. Levine, “Stochastic latent actor-critic: Deep reinforcement learning with a latent variable model,” *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [9] D. Hafner, T. Lillicrap, J. Ba, and M. Norouzi, “Dream to control: Learning behaviors by latent imagination,” *arXiv preprint arXiv:1912.01603*, 2019.
- [10] K. Chen, Y. Lee, and H. Soh, “Multi-modal mutual information (mummi) training for robust self-supervised deep reinforcement learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2021.
- [11] S. Bekkali, G. J. Youssef, P. H. Donaldson, N. Albein-Urios, C. Hyde, and P. G. Enticott, “Is the putative mirror neuron system associated with empathy? a systematic review and meta-analysis,” *Neuropsychology review*, vol. 31, no. 1, pp. 14–57, 2021.
- [12] G. Rizzolatti and L. Craighero, “The mirror-neuron system,” *Annu. Rev. Neurosci.*, vol. 27, pp. 169–192, 2004.
- [13] D. Das, S. Banerjee, and S. Chernova, “Explainable ai for robot failures: Generating explanations that improve user assistance in fault recovery,” in *Proceedings of the 2021 ACM/IEEE International Conference on Human-Robot Interaction*, 2021, pp. 351–360.
- [14] M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg, “Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks,” in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8943–8950.
- [15] S. Reddy, A. D. Dragan, and S. Levine, “Sqil: Imitation learning via reinforcement learning with sparse rewards,” in *International Conference on Learning Representations*, 2019.
- [16] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, “CARLA: An open urban driving simulator,” in *Proceedings of the 1st Annual Conference on Robot Learning*, 2017, pp. 1–16.
- [17] A. Tabrez, M. B. Luebbbers, and B. Hayes, “A Survey of Mental Modeling Techniques in Human–Robot Teaming,” *Current Robotics Reports*, vol. 1, no. 4, pp. 259–267, 2020.
- [18] R. Choudhury, G. Swamy, D. Hadfield-Menell, and A. D. Dragan, “On the utility of model learning in hri,” in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 317–325.
- [19] N. Mavridis, “A review of verbal and non-verbal human-robot interactive communication,” *Robotics and Autonomous Systems*, vol. 63, no. P1, pp. 22–35, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.robot.2014.09.031>
- [20] C. Schaff and M. R. Walter, “Residual policy learning for shared autonomy,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2020.

- [21] C. L. Baker, J. Jara-Ettinger, R. Saxe, and J. B. Tenenbaum, "Rational quantitative attribution of beliefs, desires and percepts in human mentalizing," *Nature Human Behaviour*, vol. 1, no. 4, pp. 1–10, 2017.
- [22] M. K. Ho and T. L. Griffiths, "Cognitive science as a source of forward and inverse models of human decisions for robotics and control," *arXiv preprint arXiv:2109.00127*, 2021.
- [23] H. Soh, Y. Xie, M. Chen, and D. Hsu, "Multi-task trust transfer for human-robot interaction," *The International Journal of Robotics Research*, vol. 39, no. 2-3, pp. 233–249, 2020.
- [24] M. Kwon, E. Biyik, A. Talati, K. Bhasin, D. P. Losey, and D. Sadigh, "When humans aren't optimal: Robots that collaborate with risk-aware humans," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 2020, pp. 43–52.
- [25] S. Reddy, A. D. Dragan, and S. Levine, "Where do you think you're going?: Inferring beliefs about dynamics from behavior," in *NeurIPS*, 2018.
- [26] I. Horswill, "Polly: A vision-based artificial agent," in *AAAI*, 1993, pp. 824–829.
- [27] J. Buhmann, W. Burgard, A. B. Cremers, D. Fox, T. Hofmann, F. E. Schneider, J. Strikos, and S. Thrun, "The mobile robot rhino," *Ai Magazine*, vol. 16, no. 2, pp. 31–31, 1995.
- [28] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2013, pp. 301–308.
- [29] A. D. Dragan, S. Bauman, J. Forlizzi, and S. S. Srinivasa, "Effects of robot motion on human-robot collaboration," in *2015 10th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2015, pp. 51–58.
- [30] B. Busch, J. Grizou, M. Lopes, and F. Stulp, "Learning legible motion from human-robot interactions," *International Journal of Social Robotics*, vol. 9, no. 5, pp. 765–779, 2017.
- [31] T. Kollar, S. Tellex, D. Roy, and N. Roy, "Toward understanding natural language directions," in *2010 5th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2010, pp. 259–266.
- [32] J. Thomason, M. Murray, M. Cakmak, and L. Zettlemoyer, "Vision-and-dialog navigation," in *Conference on Robot Learning*. PMLR, 2020, pp. 394–406.
- [33] V. V. Unhelkar, S. Li, and J. A. Shah, "Decision-making for bidirectional communication in sequential human-robot collaborative tasks," in *Proceedings of the 2020 ACM/IEEE International Conference on Human-Robot Interaction*, 2020, pp. 329–341.
- [34] J. Lee, J. Fong, B. C. Kok, and H. Soh, "Getting to know one another: Calibrating intent, capabilities and trust for human-robot collaboration," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 6296–6303.
- [35] E. OhnBar, K. Kitani, and C. Asakawa, "Personalized dynamics models for adaptive assistive navigation systems," in *Conference on Robot Learning*. PMLR, 2018, pp. 16–39.
- [36] Z. C. Lipton, D. Kale, and R. Wetzel, "Directly modeling missing data in sequences with rnns: Improved classification of clinical time series," in *Machine learning for healthcare conference*. PMLR, 2016, pp. 253–270.
- [37] G. E. Hinton, "Training products of experts by minimizing contrastive divergence," *Neural computation*, vol. 14, no. 8, pp. 1771–1800, 2002.
- [38] 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.
- [39] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
- [40] M. Laskin, K. Lee, A. Stooke, L. Pinto, P. Abbeel, and A. Srinivas, "Reinforcement learning with augmented data," *Advances in Neural Information Processing Systems*, vol. 33, 2020.
- [41] M. Wu and N. Goodman, "Multimodal generative models for scalable weakly-supervised learning," in *Advances in Neural Information Processing Systems*, 2018, pp. 5575–5585.
- [42] T. Zhi-Xuan, H. Soh, and D. C. Ong, "Factorized inference in deep markov models for incomplete multimodal time series," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [43] R. Rubinstein, "The cross-entropy method for combinatorial and continuous optimization," *Methodology and computing in applied probability*, vol. 1, no. 2, pp. 127–190, 1999.
- [44] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *Proceedings of the 23rd national conference on Artificial intelligence-Volume 3*, 2008, pp. 1433–1438.
- [45] J. Ho and S. Ermon, "Generative adversarial imitation learning," *Advances in neural information processing systems*, vol. 29, pp. 4565–4573, 2016.
- [46] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever, et al., "Language models are unsupervised multitask learners," *OpenAI blog*, vol. 1, no. 8, p. 9, 2019.
- [47] D. Ramachandran and E. Amir, "Bayesian inverse reinforcement learning," in *IJCAI*, vol. 7, 2007, pp. 2586–2591.
- [48] S. Balakrishnan, Q. P. Nguyen, B. K. H. Low, and H. Soh, "Efficient exploration of reward functions in inverse reinforcement learning via bayesian optimization," in *Advances in Neural Information Processing Systems*, 2020.
- [49] Y. Xie, I. P. Bodala, D. C. Ong, D. Hsu, and H. Soh, "Robot capability and intention in trust-based decisions across tasks," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 39–47.
- [50] B. C. Kok, , and H. Soh, "Trust in robots: Challenges and opportunities," *Current Robotics Reports*, 2020.
- [51] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*. Elsevier, 1988, vol. 52, pp. 139–183.