
A Unified Framework for Robots that Influence Humans over Long-Term Interaction

Shahabedin Sagheb¹, Sagar Parekh¹, Ravi Pandya², Ye-Ji Mun³, Katherine Driggs-Campbell³, Andrea Bajcsy², and Dylan P. Losey¹

Abstract

Robot actions influence the decisions of nearby humans. Here *influence* refers to intentional change: robots influence humans when they shift the human’s behavior in a way that helps the robot complete its task. Imagine an autonomous car trying to merge; by proactively nudging into the human’s lane, the robot causes human drivers to yield and provide space. Influence is often necessary for seamless interaction. However, if influence is left unregulated and uncontrolled, robots will negatively impact the humans around them (e.g., autonomous cars that repeatedly merge in front of humans may cause human drivers to become more aggressive). Prior works have begun to address this problem by creating a variety of control algorithms that seek to influence humans. Although these methods are effective in the short-term, they fail to maintain influence over time as the human adapts to the robot’s behaviors. In this paper we therefore present an optimization framework that enables robots to purposely regulate their influence over humans across both *short-term* and *long-term* interactions. Here the robot maintains its influence by reasoning over a dynamic human model which captures how the robot’s current choices will impact the human’s future behavior. Our resulting framework serves to *unify* current approaches: we demonstrate that state-of-the-art methods are simplifications of our underlying formalism. Our framework also provides a principled way to *generate* influential policies: in the best case the robot exactly solves our framework to find optimal, influential behavior. But when solving this optimization problem becomes impractical, designers can introduce their own simplifications to reach tractable approximations. We experimentally compare our unified framework to state-of-the-art baselines and ablations, and demonstrate across simulations and user studies that this framework is able to successfully influence humans over repeated interactions. See videos of our experiments here: <https://youtu.be/nPekTUFUEbo>

Keywords

Human-Robot Interaction, Control Theory, Influence, Optimization

1 Introduction

When robots and humans interact, it is inevitable that robot behavior will influence human decisions. Consider an autonomous car moving near a human-driven vehicle. If the autonomous car drives in an aggressive way — e.g., merging directly in front of the human — the human may start to drive more defensively to give the robot space. Conversely, if the autonomous car is consistently defensive — e.g., always yielding to the human — then the human may change their behavior to drive aggressively around the robot. We refer to these robots as *influential* because they choose actions that i) cause the human’s behavior to shift in a way that ii) increases the robot’s overall performance.

On one hand, the ability to influence other agents is what enables autonomous robots to seamlessly coordinate with humans. On the other hand, if the influential behaviors are haphazardly designed, they can have significant consequences. In non-embodied domains, we have already observed the effects of influential artificial intelligence (AI) agents on elections, consumer habits, and popular trends (Nicas 2018; Hagen et al. 2022; Stella et al. 2018; Tran et al. 2021; Khaund et al. 2021). Embodied robots are not nearly as pervasive as AI agents; but as robots increasingly integrate into human lives, the scale to which they influence human

behaviors will continually grow. Consider our example of shared roads — as the number of autonomous vehicles increases, humans have new expectations for how closely they can follow other vehicles, and place more responsibility on autonomous cars for avoiding incidents (Schneble and Shaw 2021; Marshall 2018; Eliot 2020). Prior works have accordingly proposed algorithms to control and regulate influential robots. Over *short-term* interactions current approaches can discourage or encourage human behaviors (Hong et al. 2023; Lazar et al. 2018), change leader and follower roles within a team (Li et al. 2021; Chari et al. 2024), and manipulate the human’s actions to increase the robot’s reward (Xie et al. 2021; Tian et al. 2023).

Despite success in short-term and one-off interactions, robots still struggle to regulate the *long-term* impacts their

¹Virginia Tech, Department of Mechanical Engineering

²Carnegie Mellon University, Robotics Institute

³University of Illinois Urbana-Champaign, Department of Electrical and Computer Engineering

Corresponding author:

Dylan P. Losey, Department of Mechanical Engineering, Virginia Tech, Blacksburg, VA, 24060, USA.

Email: losey@vt.edu

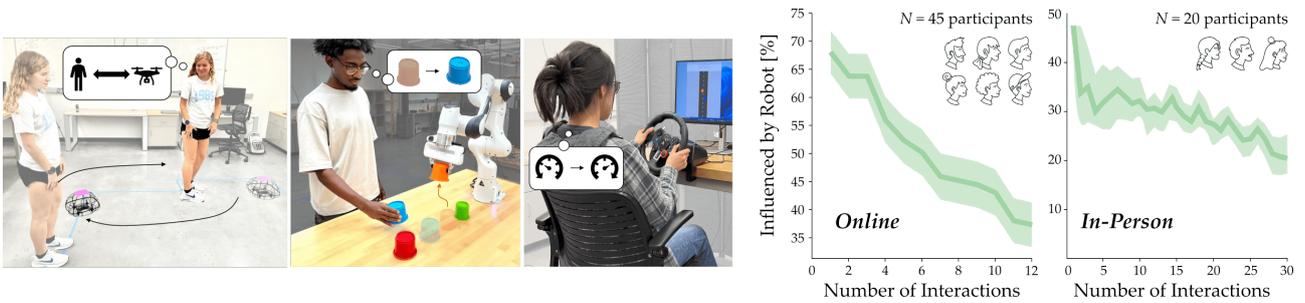


Figure 1. (Left) Human interacting with a drone, a robot arm, and an autonomous car. In each scenario, the robot uses state-of-the-art algorithms to influence the human and change their actions. (Right) Results from online and in-person user studies. The state-of-the-art approaches work in the *short-term*, i.e., human behavior is influenced by the robot in the first few interactions, but not in the *long-term*, i.e., over time the participants adapt to avoid or ignore the influential robot.

behaviors have on humans (see Figure 1). This stems from a fundamental issue with the way today’s robots model human partners. Prior work on influence typically treats the human as a static agent; i.e., it assumes that the human will always respond to a given robot behavior in a fixed way. Returning to our motivating example, each time the autonomous car merges in front of the human, it predicts that the human will slow down. Accordingly — to influence the speed of the human driver — the autonomous car simply needs to move into the human’s lane. This approach works well the first few times the human interacts with the robot. But gradually the human adapts and their responses shift; for instance, after the autonomous car merges into the human’s lane several times, the human driver may quickly accelerate to pass the autonomous car or prevent it from merging.

In this paper, we introduce a unifying formalism for influential robots. Our approach connects prior works that focus on short-term interactions, while also enabling robots to regulate their long-term influence on nearby humans. To achieve this unification, we recognize that static human models are insufficient: robots must account for how humans continually adapt over time. Indeed, relying on any single human model is bound to fail — so we instead propose a control-theoretic framework which accounts for a variety of human responses. Our core hypothesis is that:

When robots interact with humans, the robot’s current actions will affect not only the human’s current response, but also the human’s future latent strategy.

By applying this hypothesis, we are able to formulate influence as an underactuated dynamical system with an unknown and evolving latent state which parameterizes the human’s strategy. The robot is uncertain about how the human’s latent strategy will change over time, i.e., how the human will adapt to robot actions. However, during interaction, the robot gathers information about the dynamic human model, and updates not only its estimate of the current human, but also how that human will respond in the future. Robots can then leverage this dynamic understanding of the human to select actions that guide the underactuated system towards increased performance. In practice, applying our control framework results in robots that account for how their current behavior affects the human in the long-term, without relying on a single, pre-specified human model.

Our work is a step towards embodied systems that account for their potentially unintended effects on human decision making. Overall, we make the following contributions*:

Formulating Long-Term Influence. We present a formal definition for influence in human-robot interaction (Section 3). When robots interact with humans, the humans’ behavior will change in multiple ways — and not all of these changes are purposeful or meaningful. We restrict influence to only include robot behaviors that alter the human’s strategy in ways that increase the robot’s long-term reward.

Optimizing for Influential Behavior. Leveraging our definition of influence, we derive a dynamical system composed of the robot, human, and environment (Section 5). From the robot’s perspective, these overall dynamics are an underactuated system, where the immediate and future behavior of the human is indirectly controlled by the robot’s actions. The advantage of this framework is that the robot is not tied to any specific human model; instead, the robot can maintain a distribution of models, and learn how the adaptive human shifts between these models in response to the robot’s actions. More formally, we write the underactuated system as a mixed-observability Markov decision process (MOMDP), where the parameters of the human’s short- and long-term dynamics are latent states. From this MOMDP we can extract an optimal robot policy that maximizes cumulative reward while maintaining long-term influence over the human’s latent strategy. In practice, our framework enables both precise solutions (for low-dimensional settings) and tractable approximations (for more realistic problems).

Prior Works as Instances of Our Unifying Framework. Next, we theoretically and empirically demonstrate that current approaches to influence (such as Sadigh et al. (2016); Fisac et al. (2019); Schwarting et al. (2019); Xie et al. (2021); Parekh and Losey (2023)) are actually instances of our unifying framework (Sections 4 and 5.2). This includes robots that influence humans by treating the interaction as a turn-based game, and robots that learn influential policies by reasoning over latent parameters. For each of these approaches to influence, we demonstrate that we can reach the same methods by applying different simplifications or approximations to our unified MOMDP framework.

*Note that a preliminary version of this work was published at the IEEE International Conference on Robotics and Automation (Sagheb et al. 2023). Results from the preliminary conference paper are included in Section 4.2 and Section 7.1. However, this journal version is significantly different because i) we develop a unified framework for long-term influence, ii) we demonstrate how prior approaches are approximations of this unified framework, and iii) we conduct new simulations and user studies to demonstrate how the unified framework maintains long-term influence.

Designers can also apply new approximations to our unifying framework to obtain principled but tractable policies that influence humans in the long-term.

Influencing Humans in the Long-Term. We perform multiple simulations (Section 6) and user studies (Section 7) to explore how our unified approach influences human over repeated interactions. Our simulations align with our theoretical analysis and suggest that state-of-the-art algorithms are approximations of our unified approach. Across pursuit-evasion, driving, and manipulation environments, robots that apply our unified framework influence simulated humans more consistently than existing methods. We then conduct two user studies: in the first, $N = 11$ participants interact with an aerial drone, and in the second, $N = 20$ participants drive a simulated car alongside an autonomous vehicle. Across more than 25 repeated interactions, we find that our unified approach — and tractable algorithms generated from that approach — regulate long-term influence on the human more effectively than state-of-the-art baselines.

2 Related Works

Influence in HRI. As robots enter human environments, influencing humans is inevitable. Robots on factory floors have already been influencing how human workers complete tasks in ways that reduce idle time and increase team efficiency (Sanneman et al. 2021; Unhelkar et al. 2018; Mainprice and Berenson 2013; Pearce et al. 2018; Liu et al. 2018). Prior works have explored two fundamental ways in which robots can purposely influence humans: *influence by design* and *influence by action*. Influence by design uses factors such as gender (Siegel et al. 2009; Bryant et al. 2020), posture (Obaid et al. 2016), transparency (Wright et al. 2019; Hellström and Bensch 2018), or appearance (Breazeal and Scassellati 1999; Natarajan and Gombolay 2020) to alter the human’s willingness to collaborate. Although we do not focus on influence by design here, we recognize the importance of social factors for influential robots. In this work, we focus on influence by action. Here, robots select their behavior (e.g., their actions or policy) to guide humans towards advantageous states (Newman et al. 2020; Xie et al. 2021; Parekh et al. 2022; Tian et al. 2023), increase the team’s reward (Fisac et al. 2019; Sadigh et al. 2016; Tian et al. 2022; Pandya et al. 2024), or change underlying leader and follower roles (Reily et al. 2020; Li et al. 2021). For instance, the way a robot arm passes objects to a human can influence how the human grasps and holds these objects (Bestick et al. 2017; Kedia et al. 2024). Robots can also explicitly measure the influence that they have on humans they interact with by computing the divergence between the two agents’ future trajectories (Tolstaya et al. 2021), seeing if removing the robot changes the human’s behavior (Hsu et al. 2023; Schaefer et al. 2021), or measuring the deviation from the human’s nominal behavior (Ding et al. 2025). We instead recognize that not all deviation from nominal behavior is necessarily meaningful, so we formally define influential robot behavior as actions that both change the human’s underlying policy and result in *higher reward*. This enables us to explicitly optimize for the *long-term* behavior of robots to have positive influence on humans.

Influence via Games. One way for robots to choose influential actions is to pose and solve a multi-agent game. Prior works have utilized game theoretic approaches to identify control policies that maximize the robot’s total reward (Nikolaidis et al. 2017; Hadfield-Menell et al. 2016; Fridovich-Keil et al. 2020; Peters et al. 2020, 2021; Le Cleac’h et al. 2021; Mehr et al. 2023; Schwarting et al. 2019). Influential actions arise *naturally* as part of this optimization. This is particularly evident in leader-follower games, known as Stackelberg games, where the leader *acts* (e.g., the robot chooses its actions first) and then the follower *reacts* (e.g., the human selects their response given the robot’s chosen behavior). Works such as Sadigh et al. (2016); Lazar et al. (2018); Fisac et al. (2019); Schwarting et al. (2019) treat human-robot interaction as a Stackelberg game to solve for influential policies and Li et al. (2021); Tian et al. (2022) additionally consider changing the leader and follower roles online. Optimal robots in these Stackelberg games intentionally take actions that maximize the robot’s reward by shaping the human’s response (Ratliff et al. 2019). Other work has used a zero-sum game formulation for learning a robot policy that can influence the human towards safer outcomes (Hu et al. 2023; Pandya et al. 2025). Outside human-robot interaction, work on multi-agent games with reinforcement learning has shown that it can be beneficial for agents to influence how the other agents will *learn* in the environment (Foerster et al. 2018; Letcher et al. 2019; Kim et al. 2022; Lu et al. 2022) and that including a positive reward for influencing other agents can encourage exploration and cooperation (Jaques et al. 2019; Yang et al. 2020). While all these game-theoretic approaches generate influential actions, they assume both that the robot knows the human’s reward function (or it can be estimated *a priori*), and that humans are static agents and will always react the same way. In practice, human behavior shifts over time as they learn from and adapt to the robot’s behavior. We accordingly present an optimization framework that can account for such shifts by allowing the robot to *learn* the human’s policy and how it evolves online as a function of the robot’s current behavior in long-term interactions.

Influence via Latent Representations. Another line of work has shown success in learning influential robot policies by learning a *latent representation* of the human’s strategy or intention. Specifically, some work has learned the dynamics of this latent strategy between interactions (Xie et al. 2021; Wang et al. 2022b; Parekh et al. 2022) or within a single interaction (Bajcsy et al. 2024; Li et al. 2021). Other work has learned a latent strategy space offline from a fixed dataset for playing collaborative games using offline reinforcement learning (Hong et al. 2023) and imitation learning (Wang et al. 2022a). Similarly, Carroll et al. (2024) proposes a formalism called “Dynamic Reward MDPs” where an AI model may influence a human user’s reward function (i.e. their latent strategy) during an interaction. While these approaches do account for how the human’s latent intention may change over time, they still assume that the long-term dynamics are stationary — meaning that they do not capture the full dynamics of the interaction. We show that these latent representation formulations are instances of our unified framework which ignore the long-term dynamics of the human’s intentions that depend on the robot’s behavior.

Short-Term vs Long-Term Influence. Prior work has been very successful optimizing for influential behavior for robots by predicting how their actions will affect a human’s trajectory (Tolstaya et al. 2021; Tang et al. 2022; Huang et al. 2023), solving games (Sadigh et al. 2016; Fisac et al. 2019), or learning the dynamics of a latent representation (Xie et al. 2021; Bajcsy et al. 2024). However, these prior works all focus on *short-term* influence — while the human may react to the robot (e.g. yielding to an autonomous car), they do not consider how the human’s *long-term* behavior will be influenced (e.g. shifting from defensive to aggressive driving). We know from literature in AI that algorithms and recommendation systems have the ability to influence the long-term preferences of humans that interact with them, often to their detriment (Nicas 2018; Hagen et al. 2022; Adomavicius et al. 2019; Franklin et al. 2022; Bezou-Vrakatseli et al. 2023; Carroll et al. 2023). As embodied agents are starting to be integrated into human life (Waymo 2024), this will inevitably be true for robots as well. Prior work has shown that robots quickly lose their ability to influence humans in a simple repeated driving scenario (Cooper et al. 2019) and that human trust over repeated interactions with robots breaks down (Ayub et al. 2025). As a result, it is very important to understand and optimize for robot behavior that has positive *long-term influence* on humans. While some have studied repeated Stackelberg games inspired by human-robot interaction (Cooper et al. 2019; Zhao et al. 2023; Donahue et al. 2024), prior work lacks any general framework for modeling how robots can continue to exert both short-term and long-term influence on humans. In this work, we thus formally define long-term influence then introduce a unifying framework that ultimately allows the robot to find optimal, influential behavior across both the short and long term.

3 Problem Setting

In this section, we formulate the class of settings where robots can influence humans. For ease of explanation, we will focus on interactions between one robot and one human; however, our formulation can be extended to one robot working alongside multiple people. We find that influence is not present in all human-robot interactions. Instead, influence is restricted to settings where the human and robot behaviors are *interdependent*. Specifically, the robot’s performance must depend upon the human’s actions, and the human must be willing to change their actions in response to the robot.

Dynamics. Let $s_{\mathcal{R}}$ be the robot’s state, let $s_{\mathcal{H}}$ be the human’s state, and let $s = (s_{\mathcal{R}}, s_{\mathcal{H}}) \in \mathcal{S}$ be the overall system state. The robot takes actions $a_{\mathcal{R}} \in \mathcal{A}_{\mathcal{R}}$ and the human takes actions $a_{\mathcal{H}} \in \mathcal{A}_{\mathcal{H}}$. For instance, within our driving example s contains the heading and position of both cars, $a_{\mathcal{R}}$ is the steering angle and acceleration of the autonomous vehicle, and $a_{\mathcal{H}}$ is the steering angle and acceleration of the human’s car. Using t to index the current timestep, the system state transitions according to its dynamics:

$$s^{t+1} = f(s^t, a_{\mathcal{R}}^t, a_{\mathcal{H}}^t) \quad (1)$$

We emphasize that both the robot and human actions affect the system state s . Returning to the driving example, the

steering and acceleration of the human and robot change the position and heading of both vehicles. We will assume that this system state s is fully observable, and that the human and robot can observe each other’s previous actions. In other words, at timestep t the robot senses s^t and $a_{\mathcal{H}}^{t-1}$, and the human perceives s^t and $a_{\mathcal{R}}^{t-1}$.

Rewards. Both agents have a task they want to accomplish. We formulate these tasks using reward functions that map the system state s to scalar values. The robot’s reward function is $r_{\mathcal{R}}(s, \theta_{\mathcal{R}})$, and the human’s reward function is $r_{\mathcal{H}}(s, \theta_{\mathcal{H}})$. Here $\theta_{\mathcal{R}} \in \Theta_{\mathcal{R}}$ and $\theta_{\mathcal{H}} \in \Theta_{\mathcal{H}}$ are the parameters of the robot and human reward functions. For instance, in our driving example the robot’s reward function includes the distance the autonomous vehicle travels (i.e., speed) and the distance between the autonomous vehicle and the human car (i.e., safety). The parameters $\theta_{\mathcal{R}}$ determine the relative importance of the speed and safety terms. More generally, $r_{\mathcal{R}}(s, \theta_{\mathcal{R}})$ and $r_{\mathcal{H}}(s, \theta_{\mathcal{H}})$ could be neural networks, and $\theta_{\mathcal{R}}$ and $\theta_{\mathcal{H}}$ could be the weights of those networks.

Each agent knows their own reward function, but they do not know what the other agent is optimizing for. Hence, the robot does not know $\theta_{\mathcal{H}}$ and the human does not know $\theta_{\mathcal{R}}$. We define $P_0(\theta_{\mathcal{R}}, \theta_{\mathcal{H}})$ as the prior distribution over reward parameters; at the start of the interaction, the robot and human sample their reward parameters from this prior. Depending on the robot and human rewards, the agents could be collaborating (i.e., sharing a common reward function), competing (i.e., attempting to complete different tasks), or a mix of the two. We assume that the robot’s reward parameters remain constant throughout the interaction.

Interdependence. To study influence, we explore problem settings where the agents’ decisions are interdependent. More formally, we focus on scenarios where the actions of the robot impact the human’s reward, and — conversely — the human’s actions impact the robot’s reward. This interconnection is captured by the reward functions of both agents: we assume that $r_{\mathcal{R}}$ and $r_{\mathcal{H}}$ have a non-trivial dependence on the overall system state $s = (s_{\mathcal{R}}, s_{\mathcal{H}})$. For example, consider rewards that penalize collisions between the autonomous and human-driven vehicles. These collisions rely on the position of both vehicles, and thus the rewards depend on $s_{\mathcal{R}}$ and $s_{\mathcal{H}}$. In practice, this restricts our analysis to scenarios where the robot’s task performance is affected by how the human updates their state $s_{\mathcal{H}}$, and vice versa.

Interaction. The human and robot interact over a finite horizon that lasts H timesteps. In some cases this overall interaction is broken into multiple segments. For instance, human drivers will interact with autonomous cars at many intersections, and at each intersection the human driver and autonomous car need to decide who crosses that intersection first. Within these cases there are N repeated interactions of length T , such that $H = N \cdot T$, and the overall, long-term interaction is composed of the repeated segments.

In summary, our problem setting occurs during human-robot interaction when i) both agents’ actions affect the overall system dynamics, and ii) the reward of each agent depends on the entire system state. Many real-world situations fall under this problem setting: consider our motivating example of autonomous and human-driven cars, or an assistive robot working alongside a human user. These conditions result in interconnected agents whose behaviors

affect each other’s task performance. The interplay between the human and robot could play out over a single short-term game, or over long-term and repeated interactions.

Influence. Within this interdependent problem setting we can formally define influential robot behavior. Let $a_{\mathcal{R}}^t = \pi_{\mathcal{R}}(s^{0:t}, a_{\mathcal{R}}^{0:t-1}, a_{\mathcal{H}}^{0:t-1}, \theta_{\mathcal{R}})$ be a robot policy that maps its history of observations to robot actions $a_{\mathcal{R}}$, and similarly let $a_{\mathcal{H}}^t = \pi_{\mathcal{H}}(s^{0:t}, a_{\mathcal{R}}^{0:t-1}, a_{\mathcal{H}}^{0:t-1}, \theta_{\mathcal{H}})$ be a human policy that maps their history of observations to human actions $a_{\mathcal{H}}$. Without loss of generality, the robot begins with some initial policy $\pi_{\mathcal{R}}^0$, and the human has a baseline policy $\pi_{\mathcal{H}}^0$. For instance, within our motivating example the robot’s objective is to ensure the human driver maintains a safe speed. The autonomous car might initially plan to stay in the right lane, and the human is initially an aggressive driver speeding along the left lane. We say that the robot policy $\pi_{\mathcal{R}}^t$ is *influential* if i) $\pi_{\mathcal{R}}^t$ has a higher expected cumulative reward than $\pi_{\mathcal{R}}^0$, and ii) the human follows a policy different from $\pi_{\mathcal{H}}^0$ when interacting with $\pi_{\mathcal{R}}^t$. Put another way, robots influence humans when they take actions that increase their own performance while manipulating the human’s behavior. Returning to our example, an influential autonomous car might merge into the human’s lane — causing the human to switch from aggressive to defensive driving while increasing the robot’s reward for maintaining safe driving speeds.

4 Existing Approaches and Long-Term Influence

Recent works have created learning and control algorithms that regulate influence within our problem setting. Although these methods can successfully influence humans in the short-term, here we explore how humans respond to the robot over repeated, long-term interactions. We begin by providing a general formulation for existing approaches (see Section 4.1). As we will show, these approaches rely upon a fixed model of the human, i.e., they assume that the human consistently responds to the same robot behaviors in the same way. We hypothesize that this is a fundamental limitation that will cause unregulated influence in the long-term — because humans adapt, actions that were once influential may later be avoided or ignored. To test our hypothesis, we perform online and in-person user studies that measure how effectively today’s robots influence humans as a function of time (see Section 4.2). Our results suggest that — although existing frameworks can correctly influence human’s over a handful of interactions — these approaches fail as humans adapt to the robot’s behaviors.

4.1 Formulating Existing Methods

We first describe how robots currently choose influential actions. Existing approaches are based on optimization: the robot seeks to maximize its cumulative reward, while recognizing that both the robot’s and human’s behaviors contribute to that interdependent reward. The robot *directly* controls its own behaviors. In order to *indirectly* influence the human’s behaviors, the robot is equipped with a human model. This model — which can be learned or pre-programmed — predicts how the human will react to the robot’s actions. For instance, within our driving example this

human model might predict that the human will yield if the autonomous car merges in front of the human driver, and accelerate otherwise. Robots then reason over the human model to select actions (e.g., merging in front of the human) that will manipulate the user and optimize the robot’s overall reward (e.g., minimizing the human driver’s speed).

Trajectories. We formulate the underlying optimization over a finite horizon of H timesteps. Let the sequence of robot actions over H timesteps be $\mathbf{a}_{\mathcal{R}} = (a_{\mathcal{R}}^0, \dots, a_{\mathcal{R}}^H)$, and similarly let $\mathbf{a}_{\mathcal{H}} = (a_{\mathcal{H}}^0, \dots, a_{\mathcal{H}}^H)$ be the human’s action sequence. If we execute these action sequences starting at state s^0 and following the dynamics from Equation (1), we obtain a trajectory of system states. This trajectory can be written as: $\xi(s^0, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{H}}) = (s^0, \dots, s^{H+1})$. Note that the robot directly controls $\mathbf{a}_{\mathcal{R}}$, but the human’s action sequence $\mathbf{a}_{\mathcal{H}}$ relies on the person the robot is interacting with.

Total Reward. The robot seeks to maximize its cumulative reward across the interaction. Given a trajectory ξ , the robot’s total reward is the sum of state rewards:

$$R_{\mathcal{R}}(s^0, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{H}}) = \sum_{s \in \xi(s^0, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{H}})} r_{\mathcal{R}}(s, \theta_{\mathcal{R}}) \quad (2)$$

Likewise, the human’s cumulative reward is the sum of state rewards $r_{\mathcal{H}}(s, \theta_{\mathcal{H}})$ along the trajectory ξ . Because these rewards encode the desired tasks (e.g., ensuring the human driver maintains a safe speed), maximizing total reward corresponds to efficient task performance.

Optimization. Existing approaches achieve influential behaviors by maximizing Equation (2) in interdependent problem settings. We can generally break down these approaches into two components: i) an optimizer that seeks to maximize the robot’s cumulative reward, and ii) a model that predicts how the human will react to robot actions. Both components are formulated below:

$$\begin{aligned} \mathbf{a}_{\mathcal{R}}^* &= \arg \max_{\mathbf{a}_{\mathcal{R}}} R_{\mathcal{R}}(s^0, \mathbf{a}_{\mathcal{R}}, g(s^0, \mathbf{a}_{\mathcal{R}})) \\ \text{s.t. } s^{t+1} &= f(s^t, a_{\mathcal{R}}^t, a_{\mathcal{H}}^t) \end{aligned} \quad (3)$$

where $\mathbf{a}_{\mathcal{R}}^*$ is the optimal sequence of H robot actions, and $\mathbf{a}_{\mathcal{H}} = g(s^0, \mathbf{a}_{\mathcal{R}})$ is the human model. Influential behaviors naturally emerge by solving Equation (3) because the robot considers how its own actions guide the human’s response. If the robot can identify behaviors that cause an advantageous human reaction, the robot will select those behaviors and attempt to influence the human. Take our running example where the autonomous car seeks to regulate the speed of a human driver: the human model predicts that merging into the human’s lane will slow that driver. The robot therefore performs a merge — and influences the human to slow down — in order to maximize its overall reward. In theory, Equation (3) can extend across an arbitrary time horizon H . In other words, this framework could find behaviors that influence humans in the short-term and in the long-term.

Human Models. Existing works generally agree on the framework in Equation (3). However, these approaches differ in their choice of the human model; i.e., how they instantiate the function g . Below we summarize two common themes from the literature.

Methods such as Sadigh et al. (2016); Lazar et al. (2018); Fisac et al. (2019); Schwarting et al. (2019); Losey and

Sadigh (2019); Fridovich-Keil et al. (2020); Tian et al. (2022) model the human as a game-theoretic agent that seeks to maximize their own reward. If the human and robot act simultaneously, the human model solves for the Nash Equilibrium; alternatively, if the human and robot take turns, this can be treated as a Stackelberg game. In either case, here the human model becomes:

$$\begin{aligned} g(s^0, \mathbf{a}_{\mathcal{R}}) &= \arg \max_{\mathbf{a}_{\mathcal{H}}} R_{\mathcal{H}}(s^0, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{H}}) \\ \text{s.t. } s^{t+1} &= f(s^t, a_{\mathcal{R}}^t, a_{\mathcal{H}}^t) \end{aligned} \quad (4)$$

where $R_{\mathcal{H}}$ is the human’s cumulative reward. Existing methods assume that the robot knows the human’s reward function, or the robot can learn that reward function from prior data. Combining Equation (4) with Equation (3) results in nested optimization, where the robot reasons over how the optimal human will respond to its actions $\mathbf{a}_{\mathcal{R}}$.

Alternatively, methods such as Xie et al. (2021); Bajcsy et al. (2024); Jaques et al. (2019); Li et al. (2021); Chen et al. (2020) model the human as an agent whose decisions are parameterized by some latent variable z . Here the human’s policy $\pi_{\mathcal{H}}(s, z) \rightarrow a_{\mathcal{H}}$ maps system states and the latent z to human actions. This latent parameter could be the human’s willingness to adapt, the human’s strategy for completing the task, or the human’s reward or preferences when interacting with the robot. Importantly, z updates over time in response to the robot’s behavior. Consider a setting where the human and robot repeatedly interact N times, and let i index the current interaction. In this case the human model from Equation (3) becomes:

$$\begin{aligned} g(s^0, \mathbf{a}_{\mathcal{R}}) &= \pi(s, z^i), \quad i = 1, \dots, N \\ \text{s.t. } z^i &= g(z^0, s^0, a_{\mathcal{R}}^{0:i}) \end{aligned} \quad (5)$$

where $z^i = g(z^0, s^0, a_{\mathcal{R}}^{0:i})$ captures how the latent parameter updates between interactions, and $a_{\mathcal{R}}^{0:i}$ includes all robot actions up to the start of interaction i . In practice, z could express whether the human is an aggressive or defensive driver. If the human is initially aggressive, and the autonomous car merges into its lane, the human may become wary of the autonomous car and switch to more defensive driving. Under this model the robot influences the human by guiding z towards values that align with the robot’s underlying objective (e.g., manipulating the human to be a defensive driver so that the human slows down).

4.2 Testing Long-Term Influence

In Section 4.1, we outlined the general optimization framework existing robots use to influence human actions. At a high-level, this framework seeks to maximize the robot’s reward, and at a low-level, the robot achieves reward maximization by reasoning over a human model. But the two types of human models currently employed are *static*. Both Equation (4) and Equation (5) state that the human will change behaviors in response to the robot’s actions. However, this change follows a consistent pattern; e.g., if the autonomous car merges, the human driver slows down. In practice, real humans adopt and change — making robot behaviors that were once influential now ineffective.

In Section 4.2, we now explore this limitation by measuring how an existing method influences humans over

long-term interaction. Across online and in-person user studies, participants drive a simulated vehicle while sharing the road with an autonomous car. The autonomous car selects actions according to Equation (3) while using the static human model in Equation (4) (i.e., the robot applies the game-theoretic method from Sadigh et al. (2016); Fisac et al. (2019); Tian et al. (2022); Schwarting et al. (2019)). Each participant interacts with the autonomous car across highway, intersection, and roundabout environments for 36 total trials. Our results from 45 online users and 10 in-person drivers show that the robot is initially able to influence humans to yield, but over time human drivers adapt to ignore or avoid the autonomous car.

Experimental Setup. Participants shared the road with an autonomous car in three settings: highway, intersection, and roundabout (see Figure 2). These settings are consistent with related works (Sadigh et al. 2016; Tian et al. 2022; Fisac et al. 2019; Fridovich-Keil et al. 2020; Schwarting et al. 2019). To simulate the driving environment and vehicle dynamics in real-time we used CARLO (Cao et al. 2020). *In-person* participants controlled their simulated vehicle using a Logitech G29 steering wheel and responsive pedals. Each interaction ended after a fixed number of timesteps. *Online* participants first watched an animated video of the start of the interaction, and then selected their behavior from a multiple choice menu (i.e., participants could choose to yield or try and pass the autonomous car). Both in-person and online participants earned points for avoiding a collision, staying on the road, and making lane progress. We displayed the participant’s current score throughout the experiment. All participants interacted within the highway, intersection, and roundabout settings 12 times each for a total of 36 repeated interactions. The presentation order for each road setting was randomized and balanced across users.

Independent Variables. The autonomous car solved the optimization problem in Equation (3) and Equation (4) to select its actions $\mathbf{u}_{\mathcal{R}}$. More specifically, the robot treated each interaction as a separate Stackelberg game where the autonomous car played first and the human model predicted the optimal response to the robot’s actions. We rewarded the robot for avoiding collisions and minimizing the human’s lane progress (i.e., slowing down the human driver). More specifically, we selected:

$$r_{\mathcal{R}}(s, \theta_{\mathcal{R}}) = -\dot{s}_{\mathcal{H}} - 10 \cdot \mathbb{1}\{\text{collision in } s\} \quad (6)$$

where $\dot{s}_{\mathcal{H}} \subset s$ is the velocity of the participant’s vehicle. The robot assumed that the human’s reward matched the score that was displayed on the screen:

$$\begin{aligned} r_{\mathcal{H}}(s, \theta_{\mathcal{H}}) &= \dot{s}_{\mathcal{H}} - 10 \cdot \mathbb{1}\{\text{on road in } s\} \\ &\quad - 100 \cdot \mathbb{1}\{\text{collision in } s\} \end{aligned} \quad (7)$$

Positive values for $\dot{s}_{\mathcal{H}}$ indicate that the human’s car is moving forward along the road (i.e., making lane progress), while negative values mean the human’s car is in reverse. In order to maximize the state reward from Equation (6), the robot attempts to influence humans to yield, slow down, or reduce their overall lane progress.

Dependent Variables. For *online* participants, we recorded whether the human chose to yield or pass the autonomous

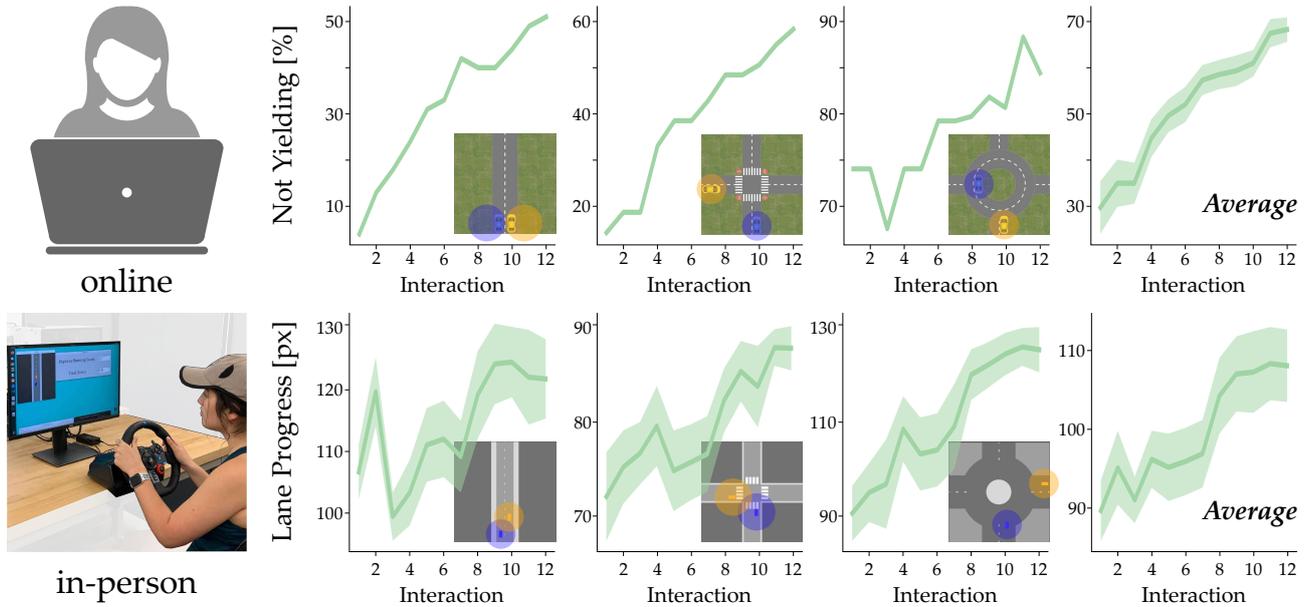


Figure 2. Exploring how the existing approaches from Section 4.1 influence humans over short- and long-term interactions. Participants repeatedly interact with an autonomous car that uses Equation (3) and the static human model from Equation (4) to influence their behavior. The autonomous car selects actions $a_{\mathcal{R}}$ by treating each interaction as an independent Stackelberg game; this is consistent with prior works (Sadigh et al. 2016; Fisac et al. 2019; Tian et al. 2022; Schwarting et al. 2019). The robot is rewarded for influencing the human to slow down, yield, and reduce lane progress. For both **online** (Top) and **in-person** (Bottom) participants, *the robot’s influence decreases over time* (i.e., human yields less or makes more lane progress). In the last column (Right) we display the average behavior across all three environments. Shaded regions show standard error.

car. For *in-person* subjects we measured their lane progress, i.e., the vertical distance they traveled. In each environment, the human’s car started at the bottom of the screen and drove towards the top of the screen; a driver that never yields to the autonomous car would maximize their lane progress.

Participants. For the *online* component of the user study we recruited 63 anonymous participants. At the start of the experiment these participants read our instructions and then answered qualifying questions to check that they understood the experimental procedure. A total of 45 users passed these questions and continued on to the survey.

For the *in-person* component we recruited 12 participants from the Virginia Tech community. Of these, 10 answered the qualifying questions correctly and completed the experiment (5 female, ages 24.7 ± 5.2 years). All participants provided informed written consent consistent with university guidelines (IRB #20-755). We recognize that users may adapt to become better drivers as they continue to interact in our simulated environment. To account for this confounding factor we had participants practice driving without any autonomous cars in each environment until they were able to consistently reach expert-level scores.

Hypothesis. We hypothesized that:

H1. *As the number of interactions increases, existing approaches will fail to regulate how they influence the human driver.*

Results. Our results from this initial user study are summarized in Figure 2. The top row shows how frequently the online users chose not to yield to the autonomous car as a function of interaction number; the bottom row displays lane progress for in-person drivers over repeated interactions.

Online users chose to either yield or pass the autonomous car during each interaction. We performed Wilcoxon signed-rank tests to see how the human’s choice evolved between the first interaction and the final interaction (i.e., to see how the human’s response to the robot changed over time). Our results averaged across all three driving scenarios reveal that humans passed the autonomous car more frequently by the end of experiment ($Z = -5.798$, $p < .001$). This change was also statistically significant for the highway ($Z = -4.583$, $p < .001$) and intersection environments ($Z = -3.838$, $p < .001$), but not for the roundabout environment ($Z = -1.155$, $p = .248$). Within the roundabout, humans rarely yielded to the robot, perhaps because they already perceived their own car as having the right of way.

For *in-person* drivers we measured their lane progress. Remember that the autonomous car is trying to influence humans to reduce their speed; as such, higher lane progress is correlated with less robot influence. Paired t-tests show the human’s average lane progress was significantly higher at the final interaction as compared to their first interaction ($t(29) = -5.952$, $p < 0.001$). This trend is consistent across highway ($t(9) = -2.4$, $p < .05$), intersection ($t(9) = -3.1$, $p < .05$), and roundabout ($t(9) = -5.6$, $p < .001$). The result suggests that robots were able to influence humans to slow at the start of the experiment, but as the number of interactions increased the robot was no longer able to influence the human effectively.

Summary. Our results from this study support **H1**. Autonomous cars that leverage an existing game-theoretic framework to generate influential behaviors are effective in the *short-term*, but do not maintain the same influence across the *long-term*. Over repeated interactions, human drivers adapt to ignore the robot’s influential actions (in the online study) or anticipate the robot and avoid its behavior (in the

in-person study). This suggests that existing approaches to influence are insufficient; if we want robots that can regulate how they manipulate the human’s actions in the long term, we need to re-examine our underlying framework.

5 A Unified Framework for Long-Term Influence

Existing robot policies fail to maintain long-term influence during human interactions. Based on our experiments, we hypothesize that this failure occurs because existing methods rely on *static* human models. Under static models, the robot assumes the human has a fixed strategy for reacting to the robot’s actions. But — as we observed in Section 4.2 — the human’s response often evolves over time (e.g., accelerating to prevent the autonomous car from merging). In Section 5.1, we now capture dynamic humans by reasoning over history-aware human models. These models predict both short-term human behaviors (e.g., yielding to the autonomous car) as well as long-term shifts (e.g., the human adapting from defensive to aggressive driving). Optimizing over *dynamic* human models leads to our unified framework for short- and long-term influence. We ultimately formulate influence as a mixed-observability Markov decision process (MOMDP). As shown in Section 5.2, casting the problem of long-term influence as a MOMDP enables two things. First, we can compute robot policies that consistently influence the human while accounting for how their current behaviors affect the user in later interactions. Second, we can demonstrate that existing approaches are actually simplifications of our framework, revealing why these methods fail to control influence across long-term interactions.

5.1 Optimizing over Dynamic Human Models

Here we introduce our unified framework for long-term influence. Similar to existing methods, our approach is based on the high-level goal of maximizing the robot’s cumulative reward. However, the key difference is how we model the human agent — particularly in the long-term. An overview of our unified approach is shown in Figure 3.

History-Aware Human Model. Let $\pi_{\mathcal{H}}^*$ denote the human’s actual policy. By definition, this policy maps some subset of all available human observations to human actions $a_{\mathcal{H}} \in \mathcal{A}_{\mathcal{H}}$. At timestep t the human can reason over $s^{0:t}$ (the system states up to and including the current state), $a_{\mathcal{R}}^{0:t-1}$ and $a_{\mathcal{H}}^{0:t-1}$ (the previous robot and human actions), and $\theta_{\mathcal{H}}$ (the human’s own reward parameters). Accordingly, the human’s true policy is of the form:

$$a_{\mathcal{H}}^t = \pi_{\mathcal{H}}^*(s^{0:t}, a_{\mathcal{R}}^{0:t-1}, a_{\mathcal{H}}^{0:t-1}, \theta_{\mathcal{H}}) \quad (8)$$

Intuitively, Equation (8) means that the human’s decisions are based on more than the robot’s most recent actions — the human may reason over the history of interactions between themselves, the robot, and the environment. This implies that the robot’s current behaviors could affect downstream human responses (e.g., the robot actions at interaction i could affect the human’s decisions at interaction $i + j$).

Dynamic Human Model. The robot does not have access to $\pi_{\mathcal{H}}^*$. Instead, the robot maintains its own human model. We have already seen examples of fixed human models in

Equation (4) and Equation (5). Here we extend these models to capture the history-aware aspects of the human’s time-varying decision making. We rewrite Equation (8) into:

$$a_{\mathcal{H}}^t = \pi_{\mathcal{H}}(s^t, z^t) \quad (9)$$

$$z^{t+1} = g_s(s^t, a_{\mathcal{R}}^t, a_{\mathcal{H}}^t, z^t, \phi^t) \quad (10)$$

$$\phi^{t+1} = g_l(s^t, a_{\mathcal{R}}^t, a_{\mathcal{H}}^t, \phi^t) \quad (11)$$

where $\pi_{\mathcal{H}}$ is the robot’s current estimate of the human’s policy. This policy is parameterized by the latent representation $z \in \mathbb{R}^d$. For instance, within our motivating example z could capture whether the human is an aggressive or defensive driver. The latent state z updates at two levels: with *short-term* dynamics g_s and *long-term* dynamics g_l , with the latent vector $\phi \in \mathbb{R}^k$ parameterizing how the human chooses their next latent strategy z^{t+1} .

Imagine our autonomous car on the highway with a human-driven vehicle. If the autonomous car merges in front of the human at timestep (or interaction) t , that could cause the human to slow and drive defensively at timestep (or interaction) $t + 1$. This immediate response is formulated by short-term dynamics g_s , which expresses how the human reacts to the robot. Here human’s reaction strategy for selecting z^{t+1} is parameterized by vector $\phi \in \mathbb{R}^k$. In practice, ϕ could be the weights of a neural network g_s , or the set of rules which govern the human’s response strategy (e.g., slow down if the autonomous car merges, accelerate otherwise). Prior works in Section 4 stop at this point: they leave ϕ as a constant value. But — to capture how humans adapt over time — we propose the long-term dynamics g_l that update ϕ in Equation (11). Returning to our example, if the autonomous car repeatedly merges in front of the human, at some timestep (or interaction) τ the human may shift their response strategy. Instead of yielding to the robot with their original rules ϕ^0 , now the human uses a new set of rules ϕ^τ to react to the robot and update z . The long-term dynamics over ϕ enable the robot to model how its current actions could affect the human’s downstream behavior (e.g., causing the human to shift from defensive to aggressive driving).

Dynamical System: State and Human Evolution. Now that we have a human model with short-term and long-term components, we next will integrate this model within a larger representation of the system. From the robot’s perspective, its actions $a_{\mathcal{R}}$ affect the interaction in two ways. The robot’s behavior causes the environment state s to transition according to the state dynamics from Equation (1). In addition, the robot’s actions cause the human’s state with parameters z and ϕ to evolve according to Equations (10) and (11). To capture the robot’s overall effect on the environment and human we therefore introduce the *augmented system state* $x = (s, z, \phi)$. This augmented state has dynamics:

$$x^{t+1} = \begin{bmatrix} s^{t+1} \\ z^{t+1} \\ \phi^{t+1} \end{bmatrix} = \begin{bmatrix} f(s^t, a_{\mathcal{R}}^t, \pi_{\mathcal{H}}(s^t, z^t)) \\ g_s(s^t, a_{\mathcal{R}}^t, \pi_{\mathcal{H}}(s^t, z^t), z^t, \phi^t) \\ g_l(s^t, a_{\mathcal{R}}^t, \pi_{\mathcal{H}}(s^t, z^t), \phi^t) \end{bmatrix} \quad (12)$$

$$= \begin{bmatrix} F(x^t, a_{\mathcal{R}}^t) \\ G_S(x^t, a_{\mathcal{R}}^t) \\ G_L(x^t, a_{\mathcal{R}}^t) \end{bmatrix} \quad (13)$$

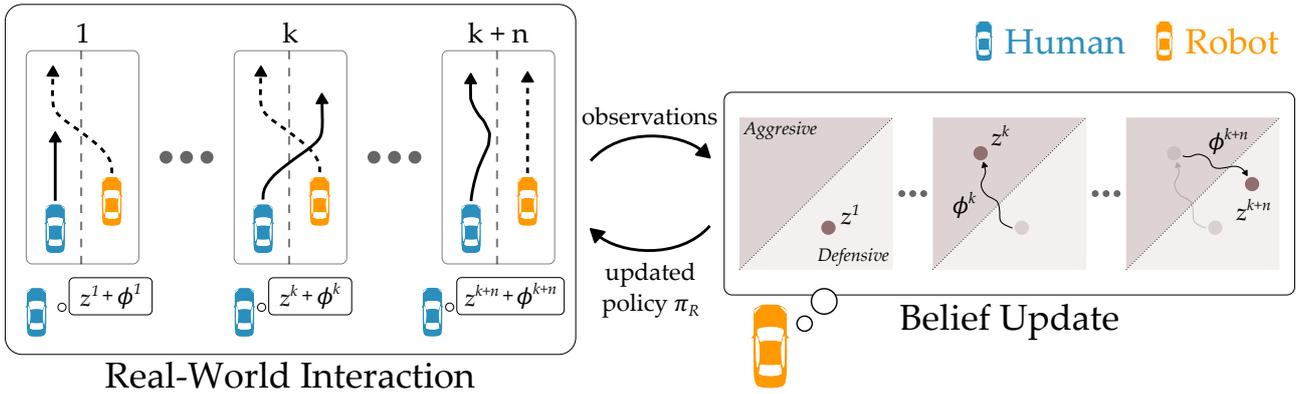


Figure 3. Our unified framework for influence. When interacting with humans, the robot models the human’s short-term as well as long-term dynamics. During interaction 1, the robot merges into the human’s lane to influence them. Upon observing the human slow down, the robot models their short term dynamics and infers that the human is a defensive driver z^1 . However, over k repeated interactions the human might change their response strategy and start driving aggressively according to their set of rules ϕ^k . Modeling the human’s long-term dynamics enables the robot to anticipate the change in the human’s response strategy. This enables the robot to optimize for a policy which can influence the human over long-term interaction. During interaction $k + n$, the robot does not merge into the human’s lane, anticipating how the human will change lanes to try and avoid the autonomous car.

Note that in the dynamical system above we have substituted the robot’s estimated human policy $\pi_{\mathcal{H}}(s, z)$ in place of the human actions $a_{\mathcal{H}}$. Incorporating this human model results in the augmented dynamics F , G_S , and G_L , which only depend on the augmented state x and the robot’s action $a_{\mathcal{R}}$.

The resulting dynamics in Equation (13) formulate how the robot can influence the human. By taking actions $a_{\mathcal{R}}$ the robot not only i) changes the physical state s , but the robot also ii) impacts the human’s short-term and long-term latent states z and ϕ . If the robot had access to Equation (13), it could use these dynamics to select actions that would drive the human towards a beneficial latent state (e.g., causing the human to remain as a safe, defensive driver).

MOMDP. The augmented state x has mixed observability: the robot can measure s , but the human’s latent states z and ϕ are unknown. Despite this uncertainty, the robot seeks to maximize its expected cumulative reward subject to the overall, single-agent dynamics presented in Equation (13). We can formulate this optimization problem as an instance of a mixed-observability Markov decision process (MOMDP) (Ong et al. 2010). The key advantage of this formulation is that — if we can solve the MOMDP — we obtain the optimal robot policy that maximizes expected reward while reasoning over the short-term and long-term human model.

Our influence MOMDP is defined by the tuple $\mathcal{M} = \langle \mathcal{X}, \mathcal{A}_{\mathcal{R}}, (F, G_S, G_L), \mathcal{O}, r_{\mathcal{R}}, H \rangle$. Here $x \in \mathcal{X}$ is the augmented system state, where s is observed by the robot and z and ϕ are unknown and must be inferred during the interaction. The robot takes actions $a_{\mathcal{R}} \in \mathcal{A}_{\mathcal{R}}$. These actions cause the augmented state to transition according to the underactuated dynamics F , G_S , and G_L from Equation (13). At each timestep the robot makes observation $o^t = (s^t, a_{\mathcal{H}}^{t-1}) \in \mathcal{O}$ with two parts: the robot measures the current system state s^t and the previous human action $a_{\mathcal{H}}^{t-1}$. The robot then uses these observations to update its probability distribution over the human’s latent parameters z and ϕ . We will refer to this probability distribution as the robot’s *belief*. The robot updates its belief using the human model, where $a_{\mathcal{H}}^t = \pi_{\mathcal{H}}(x^t, z^t)$ is the likelihood of action $a_{\mathcal{H}}^t$ given representation z^t , and $z^{t+1} = g_s(s^t, a_{\mathcal{R}}^t, a_{\mathcal{H}}^t, z^t, \phi^t)$ is the likelihood of representation z given representation dynamics

g_s with parameters ϕ . Hence, the human model presented in Equations (9)–(11) provides the observation model within the MOMDP. The final two components of the MOMDP tuple are the same as our original problem statement. The robot is trying to optimize its reward function $r_{\mathcal{R}}(x, \theta_{\mathcal{R}}) = r_{\mathcal{R}}(s, \theta_{\mathcal{R}})$, and the interaction lasts a total of H timesteps.

5.2 Existing Approaches are Approximations

In Section 5.1 we introduced a general human model with short- and long-term dynamics, and then incorporated this human model within the augmented system dynamics. From the robot’s perspective, the resulting dynamics convert the optimization problem into an instance of an MOMDP. Formulating our problem as a MOMDP enables us to apply existing tools to identify the optimal robot policy $\pi_{\mathcal{R}}$.

Solving for a Long-Term Influential Policy. A variety of online solvers have been developed to obtain near-exact estimates of $\pi_{\mathcal{R}}$. For example, methods such as POMCP (Silver and Veness 2010), DESPOT (Somani et al. 2013), and their variants (Sunberg and Kochenderfer 2018; Kurniawati 2022) input the MOMDP \mathcal{M} , the current system state s , and the current belief over z and ϕ , and then output the optimal robot action $a_{\mathcal{R}}$. Because our MOMDP includes both a short-term and long-term human model, the actions output by these solvers estimate not just how they will immediately influence the human, but also how they will guide the human’s internal state for future interactions. We apply these solvers in Sections 6 and 7 to compute influential robot’s policies. Our results empirically demonstrate how the robot policies obtained via this MOMDP maintain influential robot control over humans, maximizing robot reward while still enabling effective coordination.

Generating Tractable Approximations. The underlying MOMDP structure developed in Section 5.1 provides the *gold standard* for influential robot behaviors. The policies obtained from this framework optimize the robot’s expected performance while reasoning over short- and long-term human dynamics. However, computing exact or near-exact solutions to the MOMDP may not always be feasible. As the dimension of the continuous states $x \in \mathcal{X}$ and actions

$a_{\mathcal{R}} \in \mathcal{A}_{\mathcal{R}}$ increase, it becomes intractable to identify $\pi_{\mathcal{R}}$ with state-of-the-art MOMDP solvers.

A fundamental advantage to our unifying framework is that it provides a principled starting point for influence. When the exact solution is infeasible, we can i) derive approximations that maintain influence while being more computationally efficient. We can also ii) fairly compare different methods as simplifications of our overarching MOMDP framework. For example — instead of maintaining a belief distribution — the robot can approximate the MOMDP by assuming a point estimate over z and ϕ . Alternatively, the robot might assume that z and ϕ will be revealed at a later interaction, and so the robot does not need to take information gathering actions.

To demonstrate the value of our unifying framework, below we show how the current approaches for influence described in Section 4 are approximations of our MOMDP.

5.2.1 To Game-Theoretic Approaches Here we derive the game-theoretic methods from Equation (3) and Equation (4) (also tested in Section 4.2) as approximations of our unified approach. First, we slightly modify the human’s observations: by default, our unified approach assumes that the human can only observe the robot’s previous actions. But for this approximation we extend the observations so that the human has access to the robot’s next H actions. Let Equation (10) in our human model be:

$$z^{t+1} = g_s(s^t, \mathbf{a}_{\mathcal{R}}, a_{\mathcal{H}}^t, z^t, \phi^t) = \mathbf{a}_{\mathcal{R}} \quad (14)$$

so that the human’s short term dynamics are a prediction $\mathbf{a}_{\mathcal{R}}$ of what the robot will do. For this simplification the robot does not maintain any long-term dynamics, and $\phi = \emptyset$.

Because the robot also has access to its future actions (i.e., the robot knows what actions it is planning to take) here z is fully observable. This means that the augmented state $x = (s, z, \phi)$ is observable, and accordingly the MOMDP reduces to a Markov decision process (MDP). Within this MDP where the robot searches for actions $\mathbf{a}_{\mathcal{R}}$ that will maximize its cumulative reward $R_{\mathcal{R}}$. Simultaneously, the human observes $z = \mathbf{a}_{\mathcal{R}}$ and responds with actions $\mathbf{a}_{\mathcal{H}}$ that maximize their cumulative reward $R_{\mathcal{H}}$. This exactly corresponds to the nested optimization problem outlined in Equation (3) and Equation (4).

5.2.2 To Latent Representation Approaches Next, we derive the latent representation methods from Equation (3) and Equation (5) as a simplification of our unified MOMDP framework. As before, we begin by ignoring the long-term human dynamics. Here the human follows some rules ϕ when responding to the robot, but these rules do not change over time; hence, Equation (11) reduces to a constant value $\phi^{t+1} = \phi^t$. The key step for this derivation is to approximate our unified MOMDP as a QMPD (Littman et al. 1995). Under the QMPD approximation the robot assumes that it will uncover the augmented state (s, z, ϕ) at the next timestep. Accordingly, the robot does not need to take information gathering actions, and the QMPD simplifies into two subproblems: i) estimating the augmented state x and ii) finding optimal actions given the current estimate of x .

To perform step i), existing approaches such as Xie et al. (2021); Bajcsy et al. (2024); Parekh and Losey (2023) directly observe system state s and obtain a point estimate

of z and ϕ . Hence, the robot believes the augmented state is $\hat{x} = (s, \hat{z}, \hat{\phi})$. Treating this point estimate as the current state, in step ii) the robot solves an MDP to obtain the policy $\pi_{\mathcal{R}}$ that maximizes the cumulative reward starting at \hat{x} and evolving according to Equation (13). Unwinding our definitions, the optimal actions for this MDP correspond to the solutions of Equation (3) and Equation (5).

Interestingly, we notice that both existing approaches for influence simplify our unified framework by ignoring the long-term human dynamics g_l . This is consistent with our experimental findings from Section 4.2, and suggests that these methods fall short over repeated interactions because they fail to explicitly account for how the human will adapt to the influential robot over time.

6 Simulating Influential Robots

Here we perform two experiments with simulated humans under controlled conditions. The purpose of these simulations is to compare our unified framework from Section 5.1 with the two state-of-the-art simplifications discussed in Section 5.2. First, we implement our **Unified** and the game-theoretic **Stackelberg** approach in a simulated highway environment (see Section 6.1). Second, we test the performance of our **Unified** framework against a **Latent** representation method within pursuit-evasion, driving, and robot manipulation environments (see Section 6.2). Across both experiments we find that robots using our proposed formalism are more successful in influencing simulated humans, particularly over long-term interactions where the simulated agents completed the same tasks multiple times. Code for our simulations and experiments is available here: <https://github.com/VT-Collab/influence>

6.1 Comparing to Game-Theoretic Approaches

Our first simulation can be seen as a continuation of the user study from Section 4.2. Here simulated humans drive alongside an autonomous car on a highway environment (see Figure 4); as before, the autonomous car’s objective is to safely slow down the human driver. The simulated human attempts to get around the autonomous car and maximize their lane progress by optimizing Equation (7). Both agents have continuous state-action spaces.

Simulated Human. To simulate the human, we implemented an agent that treats the interaction as a Stackelberg game: this game-theoretic human approximates the interaction as a turn-based sequence and takes actions according to Equation (3) and Equation (4). The human initially plays second (i.e., the human responds to the robot), but based on the behavior of the autonomous car the simulated human can shift between playing first or second in the Stackelberg game. When playing first the human takes a more proactive role — i.e., the human expects the robot to respond to its actions. If the autonomous car merges in front of the human in more than 3 of the past 6 interactions, the human plays first (i.e., the human becomes more aggressive). Alternatively, if the human and autonomous car crash, the human reverts to playing second (i.e., the human becomes more defensive). The autonomous car does not know how humans will behave

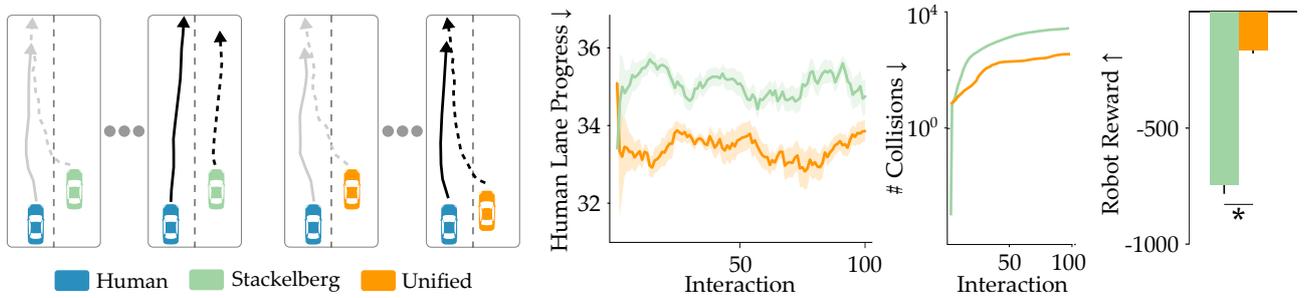


Figure 4. Experiment from Section 6.1. (Left) Simulated humans drive alongside robots using **Unified** and **Stackelberg** methods. The robot car takes actions to slow down the human without crashing. We show examples of the trajectories that the cars followed initially (i.e., first interactions) vs later on (i.e., final interactions). (Right) We plot average lane progress and total number of collisions per interaction. A higher lane progress indicates that the robot is unable to successfully influence the human to slow down. The shaded regions and error bars indicate standard error. We compare robot reward averaged across all interactions for each method. Fewer collisions and a higher reward are indicative of better robot performance. Asterisks * denote statistical significance ($p < .05$).

a priori. Simulations were performed in Julia (Egorov et al. 2017) based on the CARLO environment (Cao et al. 2020).

Independent Variables. We compared two algorithms for influencing the simulated human. As a baseline we used **Stackelberg**, the game-theoretic approximation discussed in Section 5.2 and previously implemented in works such as Sadigh et al. (2016); Lazar et al. (2018); Fisac et al. (2019); Schwarting et al. (2019). We compared this baseline to our **Unified** framework. As a reminder, our approach is based on a MOMDP formulation: to solve this MOMDP for a tractable approximation of the influential robot policy, we used the POMCPPO algorithm (Sunberg and Kochenderfer 2018).

Dependent Measures. We simulated every human for 100 interactions. Each individual interaction lasted $T = 120$ timesteps. The initial state of the cars were chosen at random, but with the condition that the autonomous car had an initial y -value greater than the simulated car (i.e., the robot started in front of the human). We measured total lane progress for both cars (in pixels), total of number of collisions between the cars, and robot’s total reward using Equation (6).

Results. Our results from the first simulation are presented in Figure 4. On the left we show qualitative examples of the trajectories that the cars followed initially (within the first 10 interactions) vs. later on (within the final 10 interactions). In the middle we plot average lane progress and the total number of collisions per interaction. Finally, on the right we compare the robot’s reward averaged across all interactions for each method. To analyze these results we performed paired t -tests and found that with **Unified** the human made significantly less lane progress by the end of the interaction ($t(99) = -10.73, p < 0.001$), the human and autonomous car had fewer collisions ($t(99) = -16.01, p < 0.001$), and the autonomous car achieved a higher average reward ($t(99) = 14.42, p < 0.001$).

Overall, these experimental results are aligned with our theoretical analysis. **Stackelberg** is a simplification of **Unified** that does not take into consideration how the human’s behavior might change over repeated interactions. Here the human shifted their role within the Stackelberg game: switching from passively playing second to more proactively playing first. With **Stackelberg** the autonomous car could not predict or account for these changes, resulting in failed influence and increased collisions. But with our **Unified** framework the autonomous car could reason over

how the human’s responses would change (i.e., the robot modeled ϕ as playing first or playing second). The **Unified** robot leveraged these long-term dynamics to maintain safe influence over 100 interactions.

6.2 Comparing to Latent Representation Approaches

In our second simulation we compare the performance of our **Unified** formalism against a state-of-the-art **Latent** representation (Parekh and Losey 2023). This **Latent** algorithm is consistent with the human model in Equation (5): the robot maintains a low-dimensional representation of the human’s strategy z , and learns how to update this strategy based on the human’s behavior from repeated interactions. More formally, the robot learns latent representations $z \in \mathcal{Z}$ and short-term dynamics g_s directly from interaction data. We test how **Latent** and **Unified** influence simulated humans across three environments: *Circle*, *Driving*, and *Robot*. In each environment the simulated human has different rules governing how they update their short-term and long-term behavior in response to the robot’s actions. We detail these rules and the environments below.

Simulated Environments & Humans. Our three environments for the second simulation are shown in Figure 5. We selected these environments to be consistent with prior work (Parekh and Losey 2023). All environments are composed of continuous state-action spaces. Within each environment the robot repeatedly interacts with a simulated human 100 times; each interaction has $T = 10$ timesteps. We repeat this process 50 times (to simulate 50 different humans) and report the averaged results. Overall, the robot’s objective is to maintain its influence across long-term interaction.

Circle. The *Circle* environment is an instance of pursuit-evasion games (Vidal et al. 2002) with two-dimensional states and actions. The robot agent (i.e., the pursuer) tries to reach the simulated human agent (i.e., the evader). To avoid the robot the human moves along the circumference of the circle. Here z encodes where the human hides from the robot along the circle, and ϕ governs how the human moves along the circle to avoid the robot. For example, some human evaders move away from the robot’s previous position. The robot’s reward is its negative distance from the human.

Driving. In *Driving* an autonomous car is trying to pass a human vehicle. At every interaction the human starts out

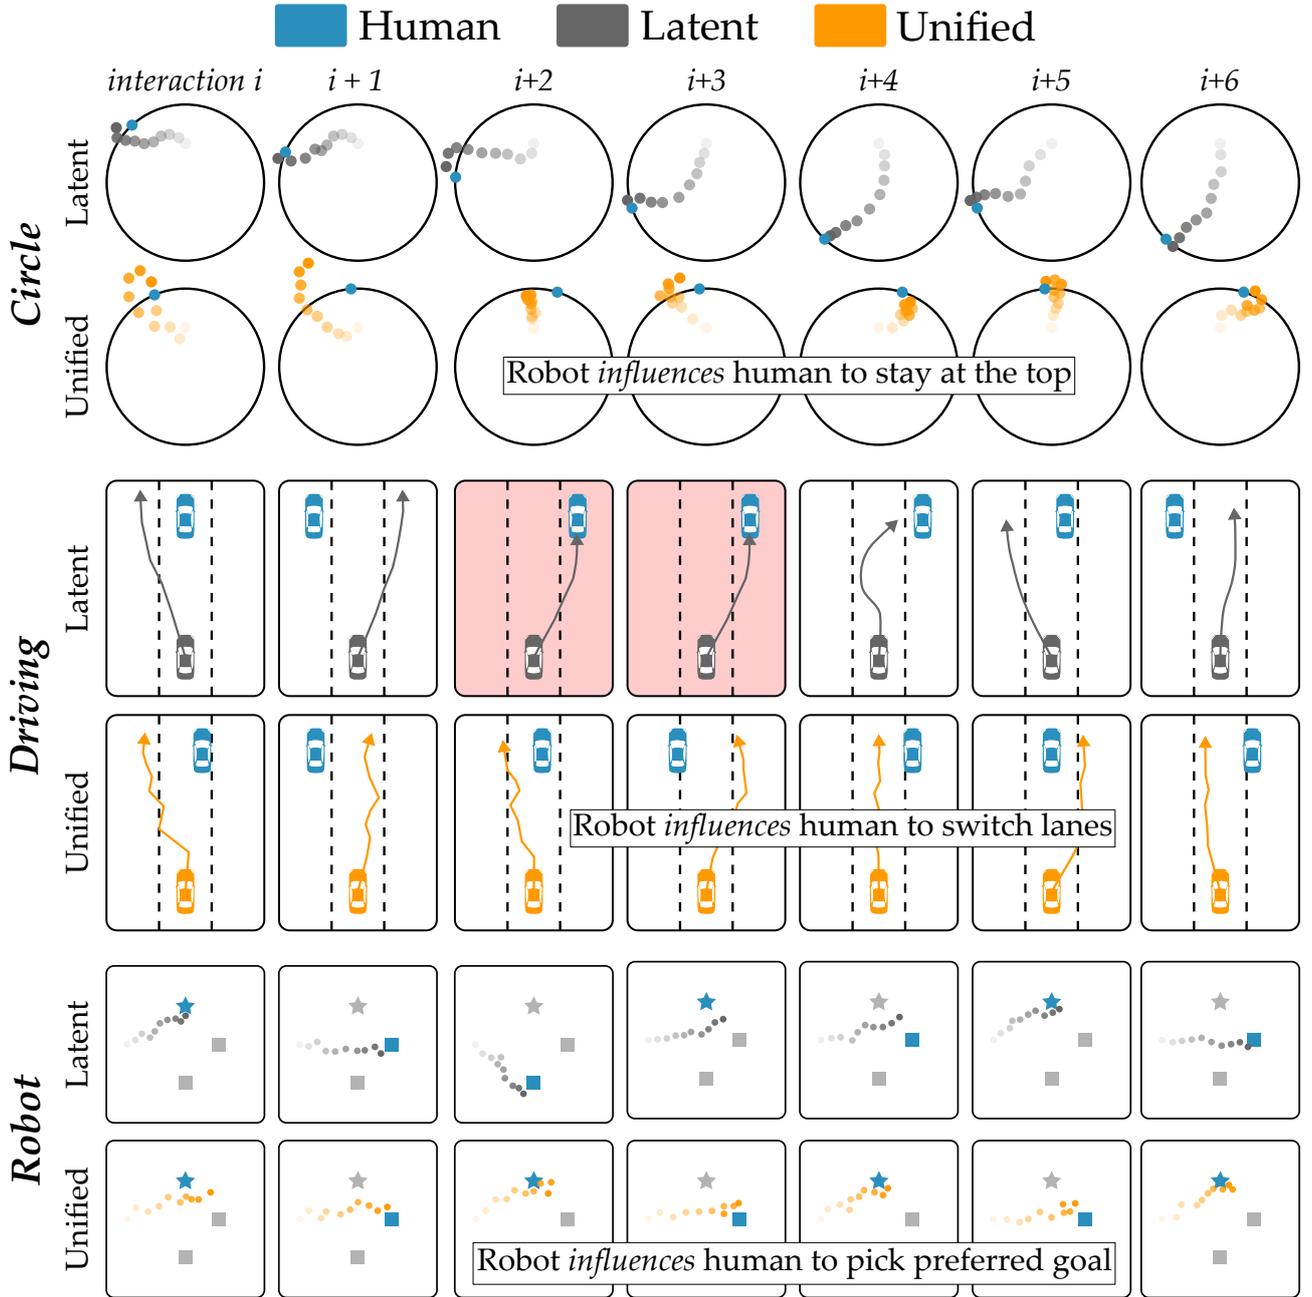


Figure 5. Example interactions from Section 6.2 between the robot and simulated human in the Circle (Top), Driving (Middle), and Robot (Bottom) environments. The first row for each environment shows interactions with a robot using the **Latent** method, while the second row shows interactions with a robot using the **Unified** method. In Circle, **Unified** is able to successfully influence the human by trapping them at the top of the circle while **Latent** is not. In Driving, **Latent** leads to higher number of collisions since it is not able to influence the human to stay out of the robot’s lane. Finally, in the Robot environment while **Latent** cycles through each of the three goals, **Unified** influences the human to frequently pick the robot’s preferred goal (shown with a star).

in front of the autonomous car, and changes lanes as the autonomous car attempts to pass. Here z encodes the lane that the human merges into, and ϕ determines how the human selects that lane. For instance, the some simulated humans merge into the lane where the autonomous car passed at the previous interaction. The robot is rewarded for passing the human, and penalized for crashing with the human. Similar to the *Circle* environment, the robot does not know which lane the human will select; to safely pass the human the robot must anticipate and influence the driver’s behavior.

Robot. Finally, in *Robot* the autonomous agent reaches for goals within a shared workspace. The robot is rewarded if it reaches for the same goal as the simulated human (i.e., if both agents move to the far left object). The robot’s

action space is its 3-DoF end-effector velocity, the latent representation z encodes the target the human wants to reach, and ϕ determines how the human updates their target object. For instance, at interaction $i + 1$ some humans select the object to the left of its target at the previous interaction i .

Implementation. To ensure a fair comparison between the baseline **Latent** and our proposed **Unified** approach, we provided the system state $s = (s_{\mathcal{R}}, s_{\mathcal{H}})$ to both algorithms. This system state included the position of the robot and the human. Here providing the robot with s alleviated the need to learn the latent strategy z , since the human’s current strategy was often captured by their state (i.e., z in *Circle* encodes the human’s location). However, we emphasize that the robot does not know how the human will update this z

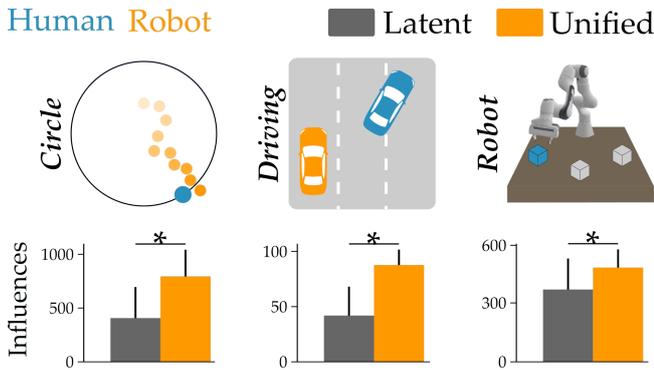


Figure 6. Experimental results from Section 6.2 comparing **Latent** and **Unified** with simulated humans. The results are shown for the three environments Circle (Left), Driving (Middle), and Robot (Right). Plots show the number of interactions where the robot manages to successfully influence the human. Asterisks * denote statistical significance ($p < .05$).

between interactions, i.e., ϕ is still a latent parameter that must be learned. Similar to Section 6.1, we implemented our MOMDP framework using the Julia programming language (Egorov et al. 2017) and the POMCPOW approximation (Sunberg and Kochenderfer 2018).

Results. Our results from the second simulation are shown in Figure 6. In the middle we provide quantitative examples of how each method interacted with simulated humans, and on the right we plot the number of times the robot successfully influenced the human. The definition of successful influence varied based on the robot’s reward. For *Circle* this was the number of times the robot pursuer trapped the human evader; for *Driving* this was the number of times the human driver moved out of the autonomous car’s way; and for *Robot* this was the number of times the human selected the object closest to the robot. Robots that applied our **Unified** framework were able to influence simulated human more frequently over long-term interactions than the **Latent** baseline. Paired t -tests revealed that the number of successful influences for **Unified** was significantly higher than **Latent** in *Circle* ($t(49) = -6.543, p < .001$), *Driving* ($t(49) = -10.933, p < .001$), and *Robot* ($t(49) = -4.411, p < .001$).

Overall, we again found that our experimental results aligned with our theoretical analysis. The **Latent** simplification was able to learn the human’s latent representation z and the patterns the human leveraged to update z in the short-term. However, when humans changed the rules that underpinned z (i.e., when the long-term dynamics caused ϕ to shift) **Latent** was confused, and the robot needed to re-learn the human’s latent strategy and response pattern. By contrast, **Unified** constantly reasoned about both short-term and long-term dynamics, ultimately learning not just how the human would update z for the next interaction, but also how the robot’s actions could impact ϕ . The robot’s ability with **Unified** to account for long-term impacts enabled the robot to more frequently influence the simulated human.

7 Influencing Users in the Long-Term

In Section 6 we demonstrated how our unified formulation can influence simulated humans over repeated interactions. Equipped with this formulation, we now return to our core challenge: controlling how robots influence actual users

over long-term interaction. In Section 4 we highlighted that existing approaches fail to regulate this long-term influence — within a span of 12 interactions, humans quickly learned to disregard the robot’s actions. Now we return to similar experimental setups, and seek to influence participants over the span of **25+ interactions**. We demonstrate two contributions of our unified approach. First, in Section 7.1 we show how designers can leverage our unified framework to generate principled but tractable simplifications. We make a greedy approximation (i.e., a one-step look-ahead) to arrive at a new algorithm for influencing humans, and test how that method influences users during human-drone interactions. Second, in Section 7.2 we leverage particle-based solvers to reach a near-exact solution to our unified framework, and compare that approach to state-of-the-art baselines. Here users drive a simulated vehicle alongside an autonomous car: we again test whether the robot is able to consistency regulate how aggressively the human drives. Videos of our user studies are available online at: <https://youtu.be/nPekTUfUEbo>.

7.1 Interacting with an Aerial Drone

In this experiment we test how our unified framework can be leveraged to reach principled but tractable simplifications. In general, computing near-exact solutions to the MOMDP may not always be feasible — but designers can introduce approximations into this framework to derive real-time methods. Game-theoretic and latent approaches to influence are two examples of these approximations. Here we generate another novel approximation: the **One-Step** method. We then compare that approach to the **Stackelberg** baseline in a user study where participants are walking near an aerial drone. The humans and drone have orthogonal goals, and often need to intersect each other’s path. At these intersections the drone takes actions to attempt to influence the humans to yield (so that the drone passes first and completes its task more efficiently). We measure how successfully the drone influences the human over 25 repeated interactions.

One-Step Approximation. On one hand, fully solving the MOMDP requires the robot to reason about all possible evolutions of its belief. On the other hand, if the robot does not consider how its belief can change, then it will be unable to reason about how the person could be influenced and thus fail to maintain control. Our **One-Step** approximation to the MOMDP in Section 5 attempts to strike a balance between these two extremes by having the robot reason about one belief update it will perform after the current timestep. Let b be the robot’s belief over the augmented state $x = (s, z, \phi)$, where z and ϕ are both unknown parameters of the human model. We instantiate z as the human’s reward function, such that $z = r_{\mathcal{H}}$, and we instantiate ϕ as the human’s estimate of the robot’s reward, such that $\phi = b(r_{\mathcal{R}})$. Within this instantiation the human is trying to determine what the robot is optimizing for: as the human uncovers $r_{\mathcal{R}}$, the human can accurately anticipate or ignore the robot’s influential actions.

To tractably solve this instantiation for the optimal robot policy $\pi_{\mathcal{R}}$, we assume that the robot will fully observe x *two timesteps in the future*. This is an improvement over the QMDP approximation (where the robot assumes it will

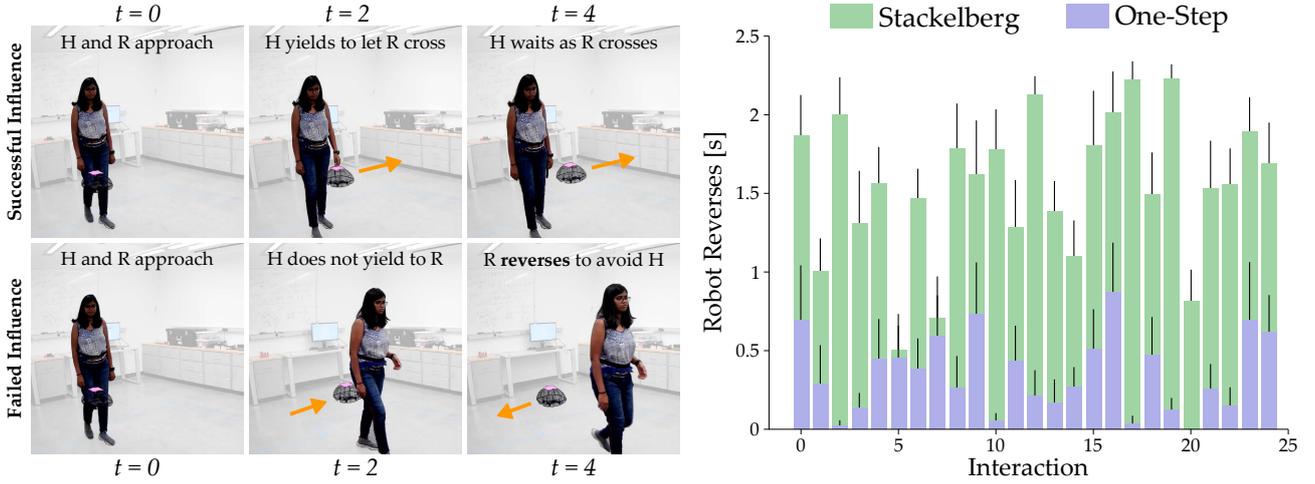


Figure 7. User study from Section 7.1 where in-person participants repeatedly cross paths with a drone. (Left) When the drone and human intersect, the drone tries to influence humans to yield so that it can cross first. If this influence fails, the robot temporarily reverses direction to avoid a collision. (Right) We plot the amount of time the robot reverses direction across 25 repeated interactions. The **Stackelberg** drone solves Equation (3) and Equation (4) to try and influence the human; this robot always yielded once the human was within a specific radius. By contrast, the **One-Step** drone used a simplification of the unified framework presented in Equation (15) to make its actions less predictable. This drone would occasionally yield at a larger radius or a smaller radius as compared to the **Stackelberg** drone. We found that this unpredictability reduced the amount of time the robot had to reverse and increased the number of successful influences. Error bars show standard error.

uncover x at the *next* timestep). As a result, our **One-Step** look-ahead greedily reasons over how its belief b will evolve, and incorporates that belief within its objective. Using this one-step simplification we reach the modified reward function (Brooks and Szafrir 2019):

$$R_{\mathcal{R}}(s^t, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{H}}) = \lambda \mathcal{H}(\phi^{t+1}) + \sum_{s \in \xi(s^t, \mathbf{a}_{\mathcal{R}}, \mathbf{a}_{\mathcal{H}})} r_{\mathcal{R}}(s, \theta_{\mathcal{R}}) \quad (15)$$

where \mathcal{H} is the Shannon entropy over the distribution ϕ at one timestep in the future, and $\lambda \geq 0$ is hyperparameter set by the designer. Comparing Equation (15) to Equation (2), we see that the key difference is the inclusion of the entropy term. Maximizing this modified reward causes the robot to take actions *now* that will increase the human’s uncertainty over the robot’s reward at the *next* timestep. Put another way, the robot proactively attempts to modify ϕ so that the human cannot uncover what the robot optimizing for or predict how the robot will behave. We expect that this **One-Step** simplification will outperform existing methods such as **Stackelberg** because i) the robot is reasoning over how ϕ evolves within long-term human dynamics and ii) the robot attempts to control that evolution so that the robot can maintain influence.

Experimental Setup. To put this simplification generated by our unified approach to the test, we created a task where participants interacted with an aerial drone (see Figure 7). The drone and human had similar tasks at opposite sides of the room. Participants walked back and forth across the room to pick up blocks and build a tower; each time the human started to cross, the drone moved orthogonally to monitor the blocks at two other locations. This resulted in repeated intersections between the drone and the human. Both agents had to determine how to navigate these intersections: e.g., the drone might move backwards to yield to the human, or the human might wait for the drone to cross. We tracked the

position of the drone using ceiling-mounted cameras, and humans wore an HTC Vive Tracker around their waist for real-time position measurements.

Independent Variables. We compared two robot controllers: the **Stackelberg** baseline introduced in Section 4 and the **One-Step** simplification of our unified framework. The robot was rewarded for crossing the room as quickly as possible while avoiding collisions with the human: to maximize its speed, the robot tried to influence humans to yield. The drone selected actions in real-time using the following reward functions:

$$r_{\mathcal{R}}(s, \theta_{\mathcal{R}}) = \dot{s}_{\mathcal{R}} - 10 \cdot \mathbb{1}\{\text{collision in } s\} \quad (16)$$

$$r_{\mathcal{H}}(s, \theta_{\mathcal{H}}) = \dot{s}_{\mathcal{H}} - 100 \cdot \mathbb{1}\{\text{collision in } s\} \quad (17)$$

where $\dot{s}_{\mathcal{R}} \subset s$ is the robot’s forward velocity and $\dot{s}_{\mathcal{H}} \subset s$ is the human’s velocity. We recognize that actual participants may not have followed this human reward function: $r_{\mathcal{H}}$ is used purely by the robot to calculate the **Stackelberg** human model in Equation (4).

Negative values for $\dot{s}_{\mathcal{R}}$ indicate that the drone is reversing direction and yielding to the human. Recall that the **One-Step** robot takes actions that make the human uncertain about what the robot is optimizing for (i.e., whether the robot is trying to avoid collisions). In practice, this caused the **One-Step** drone to randomly switch between crossing aggressively (only yielding if the human was within a small radius) and defensively (yielding if the human was anywhere within a larger radius).

Dependent Measures. A robot that maintains the right-of-way will always have a positive $\dot{s}_{\mathcal{R}}$. However, if the human insists on going first, then the robot must back off and give the participant space. To measure influence, we therefore recorded the amount of time the robot *reversed* during each interaction. Lower values correspond to higher influence (i.e., a robot that reverses direction the fewest number of times is effectively influencing the human to yield).

Participants. We recruited 11 participants from the Virginia Tech community (10 male, ages 22.1 ± 3.1 years). These participants provided informed written consent under IRB #20-755. We recognized that people may hesitate to walk close to a flying drone; before starting the experiment, we accordingly demonstrated the task and drone behaviors to help users become familiar and comfortable walking near the drone. We leveraged a within-subjects design: all participants interacted with a **One-Step** robot 25 times and a **Stackelberg** robot 25 times. The order of presentation was balanced across users, so that half started with **One-Step** and the other half started with **Stackelberg**.

Hypothesis. We hypothesized that:

H2. *A drone that approximates our unified framework while reasoning about long-term influence will cause the human to yield more consistently than the game-theoretic baseline.*

Results. Our results are summarized in Figure 7. Paired t -tests reveal that the **Stackelberg** robot spent significantly more time backing-up and yielding to the human as compared to **One-Step** ($t(327) = 13.02, p < .001$). We also noticed that — as participants became more familiar with the **Stackelberg** robot — they insisted on going first more frequently (perhaps because they were able to predict when this robot would yield). As a result, the final 8 interactions with **Stackelberg** had a higher average reverse time than the first 8 interactions, although this difference was not statistically significant ($t(10) = -1.45, p = .09$). Overall, the findings from this study supported **H2**: the **One-Step** robot consistently influenced the 11 participants to yield the right-of-way over 25 repeated interactions. This user study demonstrates that we can leverage our unified framework to generate novel simplifications (such as **One-Step**), and by carefully designing these simplifications we reach robots that effectively influence humans over long-term interaction.

7.2 Driving Alongside an Autonomous car

In our first user study (Section 7.1) we demonstrated how our unified framework can be leveraged to generate influential algorithms. This makes sense for cases where the problem setting is high dimensional, and we cannot practically solve the MOMDP. At the same time, we recognize that recent advances in optimization enable robots to reach near-exact solutions to the MOMDP in increasingly complex settings (Kurniawati 2022). In this experiment we therefore put our unified formalism to the test, and compare **Unified** to two baselines. The user study is performed in the same driving environment as in Sections 4.2 and 6.1. Participants drive a simulated vehicle alongside an autonomous car in both a highway and intersection environment. The autonomous car attempts to regulate the speed of the human on the highway, and cross in front of the human at the intersection. Importantly, this study focuses on long-term influence: participants interact with the autonomous car 30 times in each environment. Our results explore whether the unified formulation — in its complete form — regulates influence more effectively than state-of-the-art approximations.

Experiment Setup. We created driving environments consistent with both prior works on influence (Tian et al.

2023; Sadigh et al. 2016; Fisac et al. 2019; Schwarting et al. 2019) and our preliminary experiments from Sections 4.2 and 6.1. Participants shared a road with an autonomous car in two settings: Highway and Intersection (see Figure 8). The vehicles had point-mass dynamics from the CARLO environment (Cao et al. 2020). To control their car in real-time, users interacted with a Logitech G29 steering wheel and pedal. The length of the interactions was fixed; the next interaction started immediately after the previous one ended. As described in Equation (7), participants earned a score based on staying on the road, making lane progress, and avoiding collisions with the autonomous car. This score was displayed to participants throughout the experiment to encourage safe and efficient driving.

Independent Variables. We varied the autonomous car’s control strategy with three levels. To provide a baseline consistent with our previous studies, we included the game-theoretic **Stackelberg** algorithm. Next, we tested a modified version of this Stackelberg approach that added zero-mean Gaussian noise to the robot’s actions. Under **Noise** the robot solves Equation (3) and Equation (4) to find its action trajectory $\mathbf{a}_{\mathcal{R}}^* = (a_{\mathcal{R}}^{*,0}, \dots, a_{\mathcal{R}}^{*,T})$, and then we inject noise at each timestep:

$$a_{\mathcal{R}}^t = a_{\mathcal{R}}^{*,t} + \epsilon^t, \quad \epsilon^t \sim \mathcal{N}(0, \Sigma) \quad (18)$$

Intuitively, we thought that **Noise** might improve **Stackelberg** because it makes the robot’s actions less predictable, preventing the human from precisely anticipating the robot’s influential behavior over repeated interactions. Against the these baselines we compare our **Unified** formalism as described in Section 5.2. To tractably obtain $\pi_{\mathcal{R}}$ we again leveraged the POMCPOW approximation in Julia (Sunberg and Kochenderfer 2018). For each of the implemented controllers the robot reward was the same as Equation (6).

Dependent Measures. For both environments we measured the human’s lane progress, total number of collisions, and robot reward. A participant that speeds up and safely passes the autonomous car will increase their own score while lowering the robot’s reward. Here the autonomous car successfully influences the human towards defensive driving by (i) reducing the speed of the human on the *Highway* or (ii) crossing the *Intersection* in front of the human.

Participants and Procedure. We advertised our experiment to the Virginia Tech community and recruited 20 college-age participants (16 male, ages 23.75 ± 2.79 years). All the participants had experience with driving vehicles, and 16 of the users had previously played racing video games. Participants provided informed written consent following university guidelines (IRB #23-784).

At the start of this user study participants answered a demographics survey. Next — similar to the experimental procedure in Section 4.2 — participants practiced driving in each environment for roughly five minutes. After this familiarization phase we began the recorded trials by following a within subjects design. Every participant interacted with **Stackelberg**, **Noise**, and **Unified** 60 times (30 per environment) so that individual users completed a total of $60 \times 3 = 180$ trials. The order of the tasks and influential algorithms was randomized and balanced across participants. After interacting with each algorithm

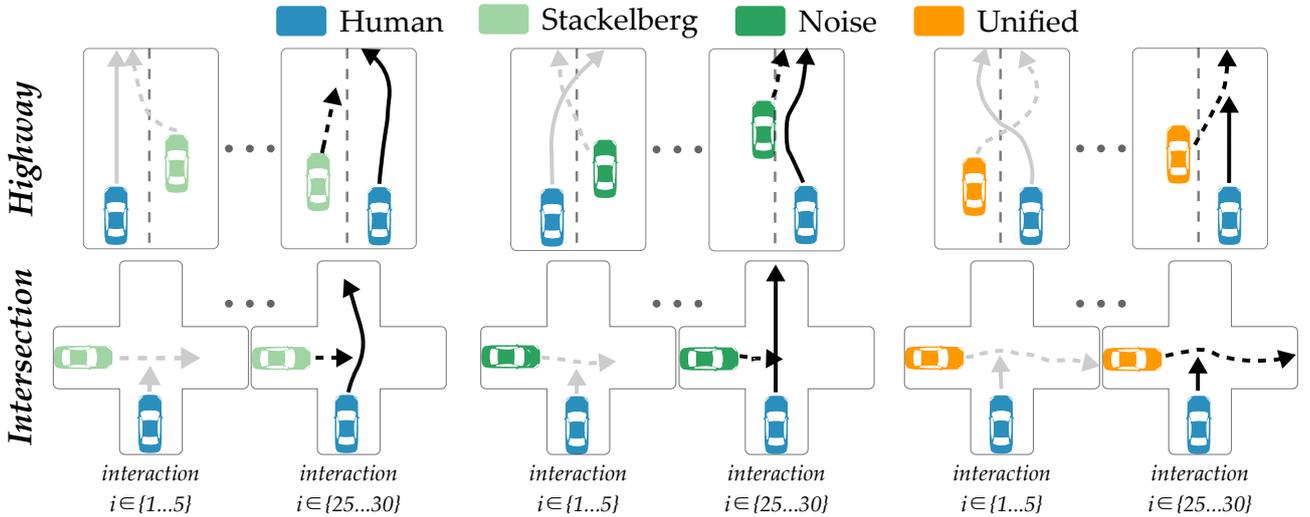


Figure 8. Example interactions from Section 7.2 between the robot and a user in Highway (Top) and Intersection (Bottom). The interactions are shown for robots using the three frameworks, **Stackelberg**, **Noise**, and **Unified**. We compare how influential the robot was in the beginning of the experiment versus at the end. All three frameworks enable the robot to successfully influence the user at the start. In Highway, the robot influences the user to slow down, reducing their lane progress. In Intersection, the robot influences the user to yield. Over long-term, repeated interactions with $N = 20$ users, only the **Unified** framework is able to maintain influence — slowing down the user in Highway or crossing in front of the user in Intersection.

the participants answered survey questions about their experience. On a 1-to-7 Likert scale, users indicated how much they agreed with the following prompts: the autonomous car’s response was *predictable*, and I *changed my inputs* based on what the autonomous car was doing.

Hypothesis. We hypothesized that:

H3. *Robots controlled with our unified framework will more effectively influence actual users over the long-term than current approaches.*

Results. Our average results across 20 participants are summarized in Figures 8 and 9.

In Figure 8 we show examples of the interaction trajectories for the two environments and three control algorithms. These trajectories are presented in pairs: on the left of each pair are the participant and robot behaviors at the start of the experiment, and on the right are the participant and robot behaviors after 30 repeated interactions. In general, we observe that **Stackelberg**, **Noise**, and **Unified** robots are all able to initially influence the human to slow down or yield. But as the human gains experience driving alongside the robot, with **Stackelberg** and **Noise** participants identify behaviors that avoid or mitigate the robot’s influential actions. By contrast, **Unified** is able to maintain influence — consider *Intersection*, where the **Unified** robot still crosses in front of the human driver at the end of 30 trials.

The objective results from our final user study are presented in Figure 9 (Left). Here we plot the human’s lane progress, the robot’s reward, and the number of collisions. Robots that influence participants should *minimize* lane progress and collisions while *maximizing* robot reward. To analyze our results — and examine whether each algorithm was able to influence humans — we conducted repeated measures ANOVAs. We found that algorithm type had a significant effect on lane progress (Highway: $F(1, 599) = 97.7, p < .001$, Intersection: $F(1, 599) = 25.9, p < .001$,

robot reward (Highway: $F(1, 599) = 21.0, p < .001$, Intersection: $F(1, 599) = 19.8, p < .001$) and collisions (Highway: $F(1, 599) = 8.6, p = .003$, Intersection: $F(1, 599) = 75.2, p < .001$). Bonferroni post hoc tests revealed that **Unified** caused the human to have less lane progress than **Stackelberg** ($p < .001$) or **Noise** ($p < .001$) across Highway and Intersection. Similarly, **Unified** increased the robot’s reward over **Stackelberg** ($p < .001$) and **Noise** ($p < .001$). For the Highway environment, the number of collisions was roughly equal for **Unified** and **Noise** ($p = 0.87$). But in the Intersection environment **Unified** reduced the collisions as compared to both **Stackelberg** ($p < .001$) and **Noise** ($p < .001$). A possible explanation for this difference in the number of collisions is that human drivers perceived the Intersection to be a more structured environment (with clear expectations for going first or second), while the Highway was less structured — e.g., human drivers could weave between lanes repeatedly to try and pass the robot. This may have led to more collisions across the board.

Finally, in Figure 9 (Right) we plot the subjective results of our survey. This survey had two items: *Predictable* and *Change*. *Predictable* indicated whether participants perceived the robot’s behavior to be easy to anticipate, and *Change* indicated whether participants thought they changed their behavior in response to the robot’s actions. Repeated measures ANOVAs show that users gave significantly different scores for *Predictable* ($F(1, 19) = 11.0, p < .01$) and *Change* ($F(1, 19) = 8.0, p < .05$) based on the type of robot they interacted with. Post hoc tests suggest that with **Unified** participants found the robot to be *harder to predict* than with **Stackelberg** ($p < .05$) or **Noise** ($p < .001$). In addition, when interacting with **Unified** participants felt that they had to *change their behaviors more* than with both **Stackelberg** ($p < .05$) and **Noise** ($p < .05$). These subjective results match our previous findings for controlling influence. When humans are able to easily anticipate what the robot will do, they can plan around those behaviors and mitigate influential actions. By contrast, robots that

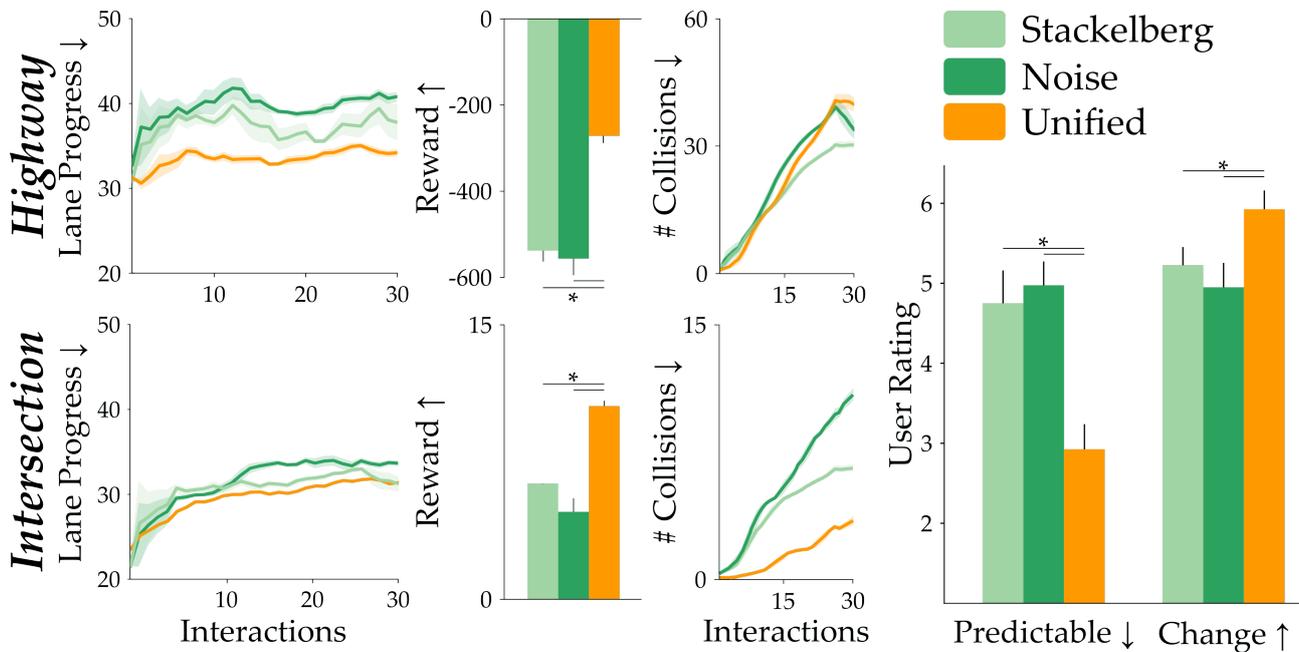


Figure 9. Objective and subjective results from our driving user study in Section 7.2. (Left) Objective results for the highway and intersection tasks. The leftmost plots show the lane progress made by the human; a lower lane progress indicates that the robot is successfully able to influence the users to slow down. The plots in the second column compare the robot’s reward across all the users. A higher reward suggests better robot performance. The plots in the third column display the number of collisions between the robot and the human in each interaction. The shaded region and the error bar in the plots show the standard error. (Right) Subjective results from Likert scale surveys. We plot the average user rating for *Predictable*, i.e., whether participants perceived the robot’s behavior to be easy to anticipate, and *Change*, i.e. whether participants thought they changed their behavior in response to the robot’s actions. A lower rating for *Predictable* suggests that the robot could maintain its influence over humans who could not easily anticipate the robot’s behavior. In contrast, for *Change* a higher rating suggests that the robot was able to influence the users.

effectively influence humans cause those humans to change their actions in response to the robot’s behavior.

Summary. Our final user study demonstrates that when robots solve our unified framework to select their policy, they are able to consistently influence humans across long-term interactions. Specifically, the **Unified** approach guided 20 participants to drive more defensively across 60 total interactions within two driving environments.

8 Conclusion

Just as AI agents have affected human decision making, embodied agents will influence the way people interact. In this paper we presented a unified framework to control how robots influence humans over long-term interaction. We first demonstrated that existing approaches are often unable to regulate influence during repeated interactions; this failure occurs because existing approaches rely on a static human model, and assume humans respond to robot behaviors in a fixed way. However, our experiments revealed that humans adapt to the robot — so that behaviors which were originally influential are later avoided or ignored.

To enable robots to reason over how humans adapt to their actions, we next developed an optimization framework for long-term influence. Our framework models humans as history-aware agents with short- and long-term dynamics, and then incorporates that human model into an augmented dynamical system. We formally write this as a single-agent system (i.e., the robot is the actor), where parameters of the human’s evolving model are unknown parts of the augmented state, and the robot gathers information about

this state by observing the human’s actions. Our overall formulation is an instance of a mixed-observability Markov decision process (MOMDP): existing optimization tools can be leveraged to obtain near-optimal robot policies from our framework. We refer to this formalism as *unified* because we can derive existing approaches from our method. For example, we show that state-of-the-art influential algorithms based on game theory and latent representations are actually simplifications of our MOMDP.

We tested our unified framework across simulations and user studies. Our simulations compared the unified method against state-of-the-art simplifications. These simplifications ignore the human’s long-term dynamics — so while they are able to influence the human at first, they struggle to maintain that influence as the human changes behavior. Accordingly, the simulation results supported our theoretical analysis and indicated that solving the unified framework outperforms state-of-the-art approximations. Next, we conducted two user studies where $N = 11$ participants interacted with an aerial drone and $N = 20$ participants drove a simulated vehicle alongside an autonomous car. These user studies focused on long-term interactions: the participants worked with the robot for 25+ trials. Our results suggest that i) designers can apply new simplifications to our method to reach tractable but influential policies, and ii) our unified approach is able to successfully regulate how it influences humans over repeated interactions.

Limitations. This paper is a step towards robots that consider how their behaviors will affect human decision making in the short-term and long-term. One limitation of this work is that our experiments focus on *one robot interacting*

with one human. Although the theoretical framework we developed in Sections 3–5 can extend to multiple users, it is not yet clear how this method will scale in practice. Theoretically, interacting with several humans increases the dimension of the latent parameters $z = (z_1, \dots, z_k)$ and $\phi = (\phi_1, \dots, \phi_k)$. Here k is the number of humans, and z_i and ϕ_i parameterize the robot’s model of the i -th human. Increasing the number of humans increases k , meaning that the latent vectors z and ϕ gain dimensions. This is a practical challenge for our approach: as discussed in Section 5.2, solving the MOMDP becomes intractable as the belief space grows. Designers can introduce simplifications to convert this multi-human MOMDP into a feasible optimization problem — but without additional experiments, it is not yet clear which simplifications are most appropriate. Overall, we see the problem of one robot influencing multiple humans as an open question where future work can build upon our formalism.

Acknowledgements

This research was supported in part by NSF Grants #2246446, #2246447, and #2246448.

Supplemental Material

Code for our simulations and experiments is available here:

<https://github.com/VT-Collab/influence>

Videos of our experiments are available here:

<https://youtu.be/nPekTUfUEbo>

References

- Adomavicius G, Bockstedt J, Curley SP, Zhang J and Ransbotham S (2019) The hidden side effects of recommendation systems. *MIT Sloan Management Review* 60(2): 1.
- Ayub A, De Francesco Z, Holthaus P, Nehaniv CL and Dautenhahn K (2025) Continual learning through human-robot interaction: Human perceptions of a continual learning robot in repeated interactions. *International Journal of Social Robotics* : 1–20.
- Bajcsy A, Loquercio A, Kumar A and Malik J (2024) Learning vision-based pursuit-evasion robot policies. In: *IEEE International Conference on Robotics and Automation*. pp. 9197–9204.
- Bestick A, Bajcsy R and Dragan AD (2017) Implicitly assisting humans to choose good grasps in robot to human handovers. In: *International Symposium on Experimental Robotics*. pp. 341–354.
- Bezou-Vrakatseli E, Brückner B and Thorburn L (2023) SHAPE: A framework for evaluating the ethicality of influence. In: *European Conference on Multi-Agent Systems*. pp. 167–185.
- Breazeal C and Scassellati B (1999) How to build robots that make friends and influence people. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 858–863.
- Brooks C and Szafr D (2019) Balanced information gathering and goal-oriented actions in shared autonomy. In: *ACM/IEEE International Conference on Human-Robot Interaction*. pp. 85–94.
- Bryant D, Borenstein J and Howard A (2020) Why should we gender? The effect of robot gendering and occupational stereotypes on human trust and perceived competency. In: *ACM/IEEE International Conference on Human-Robot Interaction*. pp. 13–21.
- Cao Z, Biyik E, Wang WZ, Raventos A, Gaidon A, Rosman G and Sadigh D (2020) Reinforcement learning based control of imitative policies for near-accident driving. In: *Robotics: Science and Systems*.
- Carroll M, Chan A, Ashton H and Krueger D (2023) Characterizing manipulation from AI systems. In: *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*.
- Carroll M, Foote D, Siththaranjan A, Russell S and Dragan A (2024) AI Alignment with changing and influenceable reward functions. In: *International Conference on Machine Learning*.
- Chari A, Chen R, Grover J and Liu C (2024) An optimal control framework for influencing human driving behavior in mixed-autonomy traffic. In: *American Control Conference*. pp. 292–298.
- Chen M, Nikolaidis S, Soh H, Hsu D and Srinivasa S (2020) Trust-aware decision making for human-robot collaboration: Model learning and planning. *ACM Transactions on Human-Robot Interaction* 9(2): 1–23.
- Cooper M, Lee JK, Beck J, Fishman JD, Gillett M, Papakipos Z, Zhang A, Ramos J, Shah A and Littman ML (2019) Stackelberg punishment and bully-proofing autonomous vehicles. In: *International Conference on Social Robotics*. pp. 368–377.
- Ding W, Veer S, Leung K, Cao Y and Pavone M (2025) Surprise potential as a measure of interactivity in driving scenarios. *arXiv preprint arXiv:2502.05677*.
- Donahue K, Immorlica N, Jagadeesan M, Lucier B and Slivkins A (2024) Impact of decentralized learning on player utilities in Stackelberg games. In: *International Conference on Machine Learning*. pp. 11253–11310.
- Egorov M, Sunberg ZN, Balaban E, Wheeler TA, Gupta JK and Kochenderfer MJ (2017) POMDPs.jl: A framework for sequential decision making under uncertainty. *Journal of Machine Learning Research* 18(26): 1–5.
- Eliot L (2020) *Zipper Merging Is A Frenzied Traffic Squeeze Play On Self-Driving Cars*. **Forbes**.
- Fisac JF, Bronstein E, Stefansson E, Sadigh D, Sastry SS and Dragan AD (2019) Hierarchical game-theoretic planning for autonomous vehicles. In: *International Conference on Robotics and Automation*. pp. 9590–9596.
- Foerster J, Chen RY, Al-Shedivat M, Whiteson S, Abbeel P and Mordatch I (2018) Learning with opponent-learning awareness. In: *International Conference on Autonomous Agents and MultiAgent Systems*. pp. 122–130.
- Franklin M, Ashton H, Gorman R and Armstrong S (2022) Recognising the importance of preference change: A call for a coordinated multidisciplinary research effort in the age of AI. *arXiv preprint arXiv:2203.10525*.
- Fridovich-Keil D, Ratner E, Peters L, Dragan AD and Tomlin CJ (2020) Efficient iterative linear-quadratic approximations for nonlinear multi-player general-sum differential games. In: *IEEE International Conference on Robotics and Automation*. pp. 1475–1481.
- Hadfield-Menell D, Russell SJ, Abbeel P and Dragan A (2016) Cooperative inverse reinforcement learning. In: *Advances in Neural Information Processing Systems*.
- Hagen L, Neely S, Keller TE, Scharf R and Vasquez FE (2022) Rise of the machines? Examining the influence of social bots on a political discussion network. *Social Science Computer Review* 40(2): 264–287.

- Hellström T and Bensch S (2018) Understandable robots-what, why, and how. *Paladyn, Journal of Behavioral Robotics* 9(1): 110–123.
- Hong J, Levine S and Dragan A (2023) Learning to influence human behavior with offline reinforcement learning. In: *Advances in Neural Information Processing Systems*. pp. 36094–36105.
- Hsu KC, Leung K, Chen Y, Fisac JF and Pavone M (2023) Interpretable trajectory prediction for autonomous vehicles via counterfactual responsibility. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 5918–5925.
- Hu H, Zhang Z, Nakamura K, Bajcsy A and Fisac JF (2023) Deception game: Closing the safety-learning loop in interactive robot autonomy. In: *Conference on Robot Learning*.
- Huang Z, Liu H, Wu J and Lv C (2023) Conditional predictive behavior planning with inverse reinforcement learning for human-like autonomous driving. *IEEE Transactions on Intelligent Transportation Systems* 24(7): 7244–7258.
- Jaques N, Lazaridou A, Hughes E, Gulcehre C, Ortega P, Strouse D, Leibo JZ and De Freitas N (2019) Social influence as intrinsic motivation for multi-agent deep reinforcement learning. In: *International Conference on Machine Learning*. pp. 3040–3049.
- Kedia K, Bhardwaj A, Dan P and Choudhury S (2024) Interact: Transformer models for human intent prediction conditioned on robot actions. In: *IEEE International Conference on Robotics and Automation*. pp. 621–628.
- Khaund T, Kirdemir B, Agarwal N, Liu H and Morstatter F (2021) Social bots and their coordination during online campaigns: A survey. *IEEE Transactions on Computational Social Systems* 9(2): 530–545.
- Kim DK, Riemer M, Liu M, Foerster J, Everett M, Sun C, Tesauro G and How JP (2022) Influencing long-term behavior in multiagent reinforcement learning. *Advances in Neural Information Processing Systems* : 18808–18821.
- Kurniawati H (2022) Partially observable Markov decision processes and robotics. *Annual Review of Control, Robotics, and Autonomous Systems* 5(1): 253–277.
- Lazar DA, Pedarsani R, Chandrasekher K and Sadigh D (2018) Maximizing road capacity using cars that influence people. In: *IEEE Conference on Decision and Control*. pp. 1801–1808.
- Le Cleac’h S, Schwager M and Manchester Z (2021) LUCIDGames: Online unscented inverse dynamic games for adaptive trajectory prediction and planning. *IEEE Robotics and Automation Letters* 6(3): 5485–5492.
- Letcher A, Foerster J, Balduzzi D, Rocktaschel T and Whiteson S (2019) Stable opponent shaping in differentiable games. In: *International Conference on Learning Representations*.
- Li M, Kwon M and Sadigh D (2021) Influencing leading and following in human–robot teams. *Autonomous Robots* 45: 959–978.
- Littman ML, Cassandra AR and Kaelbling LP (1995) Learning policies for partially observable environments: Scaling up. In: *Machine Learning Proceedings*. Elsevier, pp. 362–370.
- Liu C, Tang T, Lin HC, Cheng Y and Tomizuka M (2018) SERoCS: Safe and efficient robot collaborative systems for next generation intelligent industrial co-robots. *arXiv preprint arXiv:1809.08215*.
- Losey DP and Sadigh D (2019) Robots that take advantage of human trust. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 7001–7008.
- Lu C, Willi T, De Witt CAS and Foerster J (2022) Model-free opponent shaping. In: *International Conference on Machine Learning*. pp. 14398–14411.
- Mainprice J and Berenson D (2013) Human-robot collaborative manipulation planning using early prediction of human motion. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 299–306.
- Marshall A (2018) *How Robocars Handle the Frustratingly Human Act of Merging*. **WIRED**.
- Mehr N, Wang M, Bhatt M and Schwager M (2023) Maximum-entropy multi-agent dynamic games: Forward and inverse solutions. *IEEE Transactions on Robotics* 39(3): 1801–1815.
- Natarajan M and Gombolay M (2020) Effects of anthropomorphism and accountability on trust in human robot interaction. In: *ACM/IEEE International Conference on Human-Robot Interaction*. pp. 33–42.
- Newman BA, Biswas A, Ahuja S, Girdhar S, Kitani KK and Admoni H (2020) Examining the effects of anticipatory robot assistance on human decision making. In: *International Conference on Social Robotics*. pp. 590–603.
- Nicas J (2018) *How YouTube Drives People to the Internet’s Darkest Corners*. **The Wall Street Journal**.
- Nikolaidis S, Nath S, Procaccia AD and Srinivasa S (2017) Game-theoretic modeling of human adaptation in human-robot collaboration. In: *ACM/IEEE International Conference on Human-Robot Interaction*. pp. 323–331.
- Obaid M, Sandoval EB, Zlotowski J, Moltchanova E, Basedow CA and Bartneck C (2016) Stop! That is close enough. how body postures influence human-robot proximity. In: *IEEE International Symposium on Robot and Human Interactive Communication*. pp. 354–361.
- Ong SC, Png SW, Hsu D and Lee WS (2010) Planning under uncertainty for robotic tasks with mixed observability. *The International Journal of Robotics Research* 29(8): 1053–1068.
- Pandya R, Liu C and Bajcsy A (2025) Robots that learn to safely influence via prediction-informed reach-avoid dynamic games. In: *IEEE International Conference on Robotics and Automation*.
- Pandya R, Wang Z, Nakahira Y and Liu C (2024) Towards proactive safe human-robot collaborations via data-efficient conditional behavior prediction. In: *IEEE International Conference on Robotics and Automation*. pp. 12956–12963.
- Parekh S, Habibian S and Losey DP (2022) RILI: Robustly influencing latent intent. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Parekh S and Losey DP (2023) Learning latent representations to co-adapt to humans. *Autonomous Robots* 47(6): 771–796.
- Pearce M, Mutlu B, Shah J and Radwin R (2018) Optimizing makespan and ergonomics in integrating collaborative robots into manufacturing processes. *IEEE Transactions on Automation Science and Engineering* 15(4): 1772–1784.
- Peters L, Fridovich-Keil D, Royo VR, Tomlin CJ and Stachniss C (2021) Inferring objectives in continuous dynamic games from noise-corrupted partial state observations. In: *Robotics: Science and Systems*.
- Peters L, Fridovich-Keil D, Tomlin CJ and Sunberg ZN (2020) Inference-based strategy alignment for general-sum differential games. In: *International Conference on Autonomous Agents and MultiAgent Systems*. pp. 1037–1045.

- Ratliff LJ, Dong R, Sekar S and Fiez T (2019) A perspective on incentive design: Challenges and opportunities. *Annual Review of Control, Robotics, and Autonomous Systems* 2(1): 305–338.
- Reily B, Reardon C and Zhang H (2020) Leading multi-agent teams to multiple goals while maintaining communication. In: *Robotics: Science and Systems*.
- Sadigh D, Sastry S, Seshia SA and Dragan AD (2016) Planning for autonomous cars that leverage effects on human actions. In: *Robotics: Science and Systems*.
- Sagheb S, Mun YJ, Ahmadian N, Christie BA, Bajcsy A, Driggs-Campbell K and Losey DP (2023) Towards robots that influence humans over long-term interaction. In: *IEEE International Conference on Robotics and Automation*. pp. 7490–7496.
- Sanneman L, Fourie C and Shah JA (2021) The state of industrial robotics: Emerging technologies, challenges, and key research directions. *Foundations and Trends in Robotics* 8(3): 225–306.
- Schaefer S, Leung K, Ivanovic B and Pavone M (2021) Leveraging neural network gradients within trajectory optimization for proactive human-robot interactions. In: *IEEE International Conference on Robotics and Automation*. pp. 9673–9679.
- Schneble CO and Shaw DM (2021) Driver’s views on driverless vehicles: Public perspectives on defining and using autonomous cars. *Transportation Research Interdisciplinary Perspectives* 11: 100446.
- Schwartz W, Pierson A, Alonso-Mora J, Karaman S and Rus D (2019) Social behavior for autonomous vehicles. *Proceedings of the National Academy of Sciences* 116(50): 24972–24978.
- Siegel M, Breazeal C and Norton MI (2009) Persuasive robotics: The influence of robot gender on human behavior. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 2563–2568.
- Silver D and Veness J (2010) Monte-carlo planning in large POMDPs. In: *Advances in Neural Information Processing Systems*.
- Somani A, Ye N, Hsu D and Lee WS (2013) DESPOT: Online POMDP planning with regularization. In: *Advances in Neural Information Processing Systems*.
- Stella M, Ferrara E and De Domenico M (2018) Bots increase exposure to negative and inflammatory content in online social systems. *Proceedings of the National Academy of Sciences* 115(49): 12435–12440.
- Sunberg Z and Kochenderfer M (2018) Online algorithms for POMDPs with continuous state, action, and observation spaces. In: *International Conference on Automated Planning and Scheduling*. pp. 259–263.
- Tang C, Zhan W and Tomizuka M (2022) Interventional behavior prediction: Avoiding overly confident anticipation in interactive prediction. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems*. pp. 11409–11415.
- Tian R, Sun L, Bajcsy A, Tomizuka M and Dragan AD (2022) Safety assurances for human-robot interaction via confidence-aware game-theoretic human models. In: *International Conference on Robotics and Automation*. pp. 11229–11235.
- Tian R, Tomizuka M, Dragan AD and Bajcsy A (2023) Towards modeling and influencing the dynamics of human learning. In: *ACM/IEEE International Conference on Human-Robot Interaction*. pp. 350–358.
- Tolstaya E, Mahjourian R, Downey C, Vadarajan B, Sapp B and Anguelov D (2021) Identifying driver interactions via conditional behavior prediction. In: *IEEE International Conference on Robotics and Automation*. pp. 3473–3479.
- Tran AD, Pallant JI and Johnson LW (2021) Exploring the impact of chatbots on consumer sentiment and expectations in retail. *Journal of Retailing and Consumer Services* 63: 102718.
- Unhelkar VV, Lasota PA, Tyroller Q, Buhai RD, Marceau L, Deml B and Shah JA (2018) Human-aware robotic assistant for collaborative assembly: Integrating human motion prediction with planning in time. *IEEE Robotics and Automation Letters* 3(3): 2394–2401.
- Vidal R, Shakernia O, Kim HJ, Shim DH and Sastry S (2002) Probabilistic pursuit-evasion games: Theory, implementation, and experimental evaluation. *IEEE Transactions on Robotics and Automation* 18(5): 662–669.
- Wang C, Pérez-D’Arpino C, Xu D, Fei-Fei L, Liu K and Savarese S (2022a) Co-GAIL: Learning diverse strategies for human-robot collaboration. In: *Conference on Robot Learning*. pp. 1279–1290.
- Wang WZ, Shih A, Xie A and Sadigh D (2022b) Influencing towards stable multi-agent interactions. In: *Conference on Robot Learning*. pp. 1132–1143.
- Waymo (2024) *That’s a Wrap! Waymo’s 2024 Year in Review*. [Waymo Blog](#).
- Wright JL, Chen JY and Lakhmani SG (2019) Agent transparency and reliability in human–robot interaction: The influence on user confidence and perceived reliability. *IEEE Transactions on Human-Machine Systems* 50(3): 254–263.
- Xie A, Losey DP, Tolsma R, Finn C and Sadigh D (2021) Learning latent representations to influence multi-agent interaction. In: *Conference on Robot Learning*. pp. 575–588.
- Yang J, Li A, Farajtabar M, Sunehag P, Hughes E and Zha H (2020) Learning to incentivize other learning agents. *Advances in Neural Information Processing Systems* : 15208–15219.
- Zhao G, Zhu B, Jiao J and Jordan M (2023) Online learning in Stackelberg games with an omniscient follower. In: *International Conference on Machine Learning*. pp. 42304–42316.