From Legible to Inscrutable Trajectories: (Il)legible Motion Planning Accounting for Multiple Observers

Ananya Yammanuru*, Maria Lusardi*, Nancy M. Amato, and Katherine Driggs-Campbell

* Equal Contribution; Authors listed randomly with seed '2026'
University of Illinois at Urbana-Champaign, Urbana IL 61801, USA

Abstract. In cooperative environments, such as in factories or assistive scenarios, it is important for a robot to communicate its intentions to observers, who could be either other humans or robots. A legible trajectory allows an observer to quickly and accurately predict an agent's intention. In adversarial environments, such as in military operations or games, it is important for a robot to not communicate its intentions to observers. An illegible trajectory leads an observer to incorrectly predict the agent's intention or delays when an observer is able to make a correct prediction about the agent's intention. However, in some environments there are multiple observers, each of whom may be able to see only part of the environment, and each of whom may have different motives. In this work, we introduce the **Mixed-Motive Limited-Observability Legible Motion Planning** (MMLO-LMP) problem, which requires a motion planner to generate a trajectory that is legible to observers with positive motives and illegible to observers with negative motives while also considering the visibility limitations of each observer. We highlight multiple strategies an agent can take while still achieving the problem objective. We also present DUBIOUS, a trajectory optimizer that solves MMLO-LMP. Our results show that DUBIOUS can generate trajectories that balance legibility with the motives and limited visibility regions of the observers. Future work includes many variations of MMLO-LMP, including moving observers and observer teaming.

Keywords: Path Planning · Trajectory Optimization · Privacy · Legibility

1 Introduction

Modern robotic systems, including warehouse fleets, autonomous vehicles, and patrol agents, often operate alongside other decision-making agents. Generally, individual agents are expected to account for the future trajectory of other nearby agents when planning their paths. Thus, when an agent moves, it is not only moving closer to its goal, but it is also signaling its intent to observers. The task of generating a trajectory that allows an observer to quickly and confidently infer an agent's goal is called *legible motion planning*.

In this work, we study two special cases of the legible motion planning problem: limited observability and mixed motive. Current work in legible motion planning generally optimizes a legibility score over an entire trajectory, thus assuming that an observer can reliably see an entire robot trajectory from start to finish. In most situations, an observer would not be able to see the agent's full trajectory, either due to obstacles or limited sensor range. Limited observability has been briefly addressed in prior literature [21], but does not account for multiple simultaneous observers with different but overlapping visibility regions. In contrast, we consider settings with multiple observers whose visibility regions may differ and overlap.

Aside from differences in visibility regions, observers may also have differing motives. Most previous work assumes that the observers are cooperative, meaning that the observers' objectives are in-line with the agent's goals. This implies that it is in the agent's best interest to quickly reveal its intent as early as possible in a trajectory. However, an agent may want to reveal its intention to trusted agents in the scene while concealing it from untrustworthy agents. Adversarial motion planning has been explored for a single, full-visibility observer. Main strategies have included trajectory optimization towards a decoy goal [6], or applying RRT* against a ML-based intent predictor [20]. These methods do not account for either multiple observers with different motives or limited visibility.

These challenges motivate us to introduce the **Mixed-Motive Limited-Observability Legible Motion Planning** (MMLO-LMP) problem. Our goal is to generate trajectories that clearly convey our intent to cooperative agents while concealing it from adversarial agents. We also present the **DU**al **BI**ased limited **O**bservability **U**nified trajectory **S**olver (DUBIOUS), which generates trajectories which satisfy the MMLO-LMP requirements.

While the legible motion planning problem can be solved by optimizing a legibility function, limited observability makes the problem more challenging. An observer with a very limited visibility range may see significantly less of a robot trajectory. Thus, the agent needs to optimize its time within an observer's visibility region in order to maximize the amount of information conveyed. The mixed motive observers add further complexities to this problem. A robot's strategy now not only depends on conveying information to its observers, but may also include either avoiding or intentionally misleading a malicious observer. If a good and a malicious agents' visibility regions overlap, the robot must appropriately strategize the information that it is conveying. This combination of limited observability and mixed motives changes how legibility can be optimized. Additionally, the location of a malicious observer's visibility region (near the robot's start position versus near the goals) can change the way an agent should behave. In this work, we:

- Formally define the Mixed-Motive Limited-Observability Legible Motion Planning problem (MMLO-LMP; Section 3.2). We also present interesting and important applications for MMLO-LMP. (Section 5)
- Present an optimization-based trajectory solver, the **DU**al **BI**ased limited **O**bservability **U**nified trajectory **S**olver (DUBIOUS) for addressing the MMLO-LMP problem. This includes 2 different strategies for addressing a malicious observer. (Section 3.3)
- Show some examples of interesting MMLO-LMP problem cases, and quantitatively and qualitatively analyze DUBIOUS's solutions. (Section 4).

2 Related Work

In this section we present work relevant to legible path planning, adversarial path planning, and observers with limited visibility. Then, we briefly discuss how our work addresses all three situations.

2.1 Legible Path Planning

There is much preexisting work on legible and illegible planning in both motion and task planning domains. Chakraborti et al. [4], highlight and differentiate between the terms used to describe legibility and similar problems, including "explicability", "predictability", "transparency", or conversely, "security" and "privacy". They define "plan legibility" as a plan which "reduces ambiguity over possible goals that might be achieved" and note that it is equivalent to plan "transparency".

Extensive work on legibility for robot motion planning has been done by Anca D. Dragan [7, 8, 6]. In [8], Dragan et al. define legible motion in the context of motion planning as "motion that enables an observer to quickly and confidently infer the correct goal" from a set of candidate goals. This is different from predictable motion, which is generated assuming that the observer already knows the goal. Legible and predictable motion can be the same if there is only one possible goal, but in many cases they are different, or even contradictory, since they are derived from different assumptions about the observer's knowledge. Drawing on psychology research in action interpretation, they model the legible inference as the goal which would be most efficiently reached by the observed trajectory. Therefore, we can calculate legibility score based on the probability of the true goal given the observed trajectory. Using this model, they evaluated human perception of legible motion and found that legible trajectories did enable more subjects to make correct predictions faster as opposed to predictable, straight-line trajectories from start to goal.

In [7], Dragan and Srinivasa present a method for generating legible trajectories using gradient descent with the CHOMP trajectory optimizer ([25]). More iterations produce trajectories with increasing legibility scores and increasing deviation from the most predictable trajectory. They found that more legible trajectories increased human subjects' prediction accuracy, but only up to a point. When the optimized legible trajectory becomes too dissimilar from the predictable trajectory, prediction confidence goes down. They define this point as the "trust region", a maximum area where legible trajectories are still considered trustworthy.

DUBIOUS uses the formulation of legibility similar to [8], but we do not use the CHOMP optimizer from [7], because we adapt the legibility cost function such that it is no longer differentiable. Instead we use Stochastic Trajectory Optimization for Motion Planning (STOMP) [11]. STOMP does not require a differentiable cost function and is less likely to get stuck in local minima.

2.2 Adversarial Path Planning

The inverse of legibility is goal "obfuscation" or "privacy". It entails any method which attempts to keep multiple goals likely in the observer's model given its observations [4].

Illegible behavior generation have been explored from a game theory perspective [23] and more recently with machine learning approaches such as RNNs for estimating observer predictions [20].

Masters and Sardina note that deception involves two components, simulation ("showing the false"), and dissimulation ("hiding the true") [18]. Simulation occurs when the probability of a decoy goal is strictly greater than the true goal. Dissimulation occurs when the probability of the true goal is less than or equal to the probability of any other goal. Thus deceptive planning in their formulation always involves dissimulation, but not necessarily simulation.

Dragan et al. present three methods which extend their legibility formulation to deceptive motion in [6]. Exaggerating motion maximizes legibility to a decoy goal instead of the true goal. Switching motion maximizes legibility to a randomly chosen goal, which may include the true goal, at each timestep. Ambiguous motion keeps the probability of all goals equal. They found exaggerating motion was most effective at deceiving human subjects because subjects were more confident in their incorrect predictions.

DUBIOUS employs a decoy or an ambiguous strategy analogous to the simulation and dissimulation components identified by [18] respectively. For the decoy strategy, we adapt the model for exaggerating motion as described in [6] by selecting a decoy goal to maximize for instead.

2.3 Limited Visibility Observers

There is a large body of work on the Minimal Exposure Problem, that is, the problem of planning a path through a region of observers that is observed as little as possible [5]. Many variations on the type of occlusions and visibility regions, such as terrain [15, 16], limitations of sensor measurement [24, 9, 5, 22], and graph formulations of the problem [17, 2] have been studied. Our work treats observer visibility as polygonal, weighted regions, similar to [3, 10]. Calculating visibility regions in a space is a prerequisite for our method, so the contributions of [3] and [10] are parallel to ours. However, while we recognize that avoiding observer visibility regions can be an effective adversarial strategy, we propose that paths which are purposely seen so as to communicate, or obfuscate, intent can, in some cases, be more effective.

Kulkarni et al. [13] present a framework for offline task planning which can produce legible or obfuscated plans for observers with partial visibility. Observers know the agent's planning model and a set of candidate goals, but do not know the agent's true goal and cannot directly observe the agent's actions. Observers receive observations emitted by an agents actions, but different actions can emit the same observation, introducing ambiguity into the observer's belief about the agent's plans or goals. Agents use this ambiguity to order tasks such that observations are consistent with at least k goals (in an adversarial context) or at most j goals (in a cooperative context). They test with only one observer who is either beneficial or adversarial, though their framework could be extended to scenarios with multiple observers with different motives and visibility. [14] also considers the problem in the context of proactive assistance .

Taylor et al. also consider observers with limited visibility, but instead define visibility by what portion of the environment an observer can see [21]. They evaluate the legibility of a robot serving tables in a restaurant based on the ability of diners (who cannot see the entire restaurant) to predict what table the robot is going towards. They also analyze the problem of remaining legible to multiple observers with differing visibility regions, which their method does not address.

Mavrogiannis et al. address the problem of legibility with multiple observers also active in an environment, namely navigating through pedestrians [19]. Given the application, they define legibility as communicating a collision avoidance strategy rather than communicating intended end goal.

Our work also presents a unified method for legible or obfuscated planning, but in an online motion planning context. The notion of legibility, and how it can be maximized or minimized, differs in many important ways in the motion vs. task planning domains as elaborated by [4]. Namely, in the motion planning context we can simplify the observer model by assuming shorter paths are preferred. Since we consider online observers, DUBIOUS is concerned not only with whether an observer predicts the correct goal, but when they predict the correct goal. Since we work in a motion planning context, like [21], visibility is defined by geometric area the observer can see over, and an observer receives accurate information about an agent's actions in that area. Similar to [19], we extend this problem to multiple observers, but we consider the case of observers with mixed motive. Our agent should be legible to some observers and illegible to others.

3 Methodology

3.1 Preliminaries and Problem Definition

We define a environment $E \subset \mathbb{R}^2$ which contains an agent $\mathcal{A}$ and a set of candidate goals $\mathcal{G}$. Let Ξ denote the set of all possible trajectories. A trajectory $\xi : [0, T] \to E$ has start position $\xi(0) = S$ and goal position $\xi(T) = G_R$, where $G_R \in \mathcal{G}$. Thus, $\xi(t)$ is the agent's position at time t.

Legibility. For a trajectory ξ, we define the cost $C : \Xi \to \mathbb{R}^+$ as a measure of efficiency. A lower C indicates a more efficient, or 'better' trajectory. For clarity, we refer to $\xi(t)$ for some arbitrary t as Q.

In this work, we use $C(\xi) = \frac{1}{2}\int ||\dot{\xi}||^2 dt$ (same as [7]). The trajectory which minimizes this cost is a straight-line path from S to G_R with constant velocity. We also define $C^*_G(Q) = \min_{\xi \in \Xi_{Q \to G}} C(\xi)$ as the cost of the optimal trajectory from Q to G.

We define *legibility* similarly to [6–8]. A legible trajectory enables an observer to quickly infer the correct goal based on observing a portion $\xi_{S \to \xi(t)}$ of a trajectory. Thus, we can assign the legibility score of a trajectory as

$$\textsc{Legibility}(\xi) = \frac{\int P(G_R | \xi_{S \to \xi(t)})(T - t) dt}{\int (T - t) dt} \tag{1}$$

A high legibility score indicates that $P(G_R | \xi_{S \to \xi(t)})$ remains high from an early portion of the trajectory.

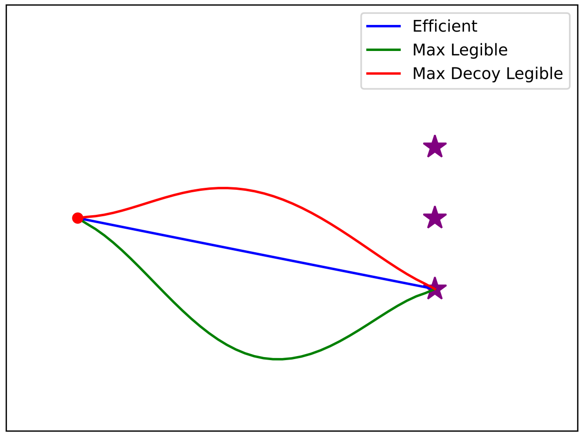

Fig. 1: An example of a legible and a decoy-legible trajectory. These baselines are generated with DUBIOUS in an environment that assumes full observability for a single +1 or -1 observer, for the max legible and max decoy illegible trajectories, respectively. The decoy trajectory is an example of the ILLEGIBLE-DECOY strategy. Note that the legible trajectory curves away from the other goal options, reducing the probability that the other goals are the true goal. Conversely, note that the illegible trajectory moves toward the other opponents before curving toward its correct goal.

It is shown in [8], that $P(\xi_{S\to Q}|G) \approx \frac{k\exp(-C(\xi_{S\to Q})-C^*_G(Q))}{\exp(-C^*_G(S))}$ for constant k. We calculate $P(G_R|\xi_{S\to Q})$ using Bayes rule and substitution:

$$P(G_R|\xi_{S\to Q}) = \frac{P(\xi_{S\to Q}|G_R)P(G_R)}{\sum_{G\in\mathcal{G}} P(\xi_{S\to Q}|G)P(G)} \tag{2}$$

There are two ways for a trajectory to increase $P(G_R|\xi_{S\to Q})$ and hence be more legible. First, by being predictable in regards to the true goal (thus increasing $P(\xi_{S\to Q}|G_R)$ Second by being unpredictable in regards to all other goals (thus decreasing $P(\xi_{S\to Q}|G)$ for all $G \neq G_R$. This can be understood when viewed through the lens of implicit communication [12]. In order for a trajectory to successfully communicate the intended goal, the action must be somewhat unpredictable, thus triggering the observer to seek an explanation, but more understandable once the true goal is known, thus providing a plausible explanation for the agent's behavior.

$$P(G_R|\xi_{S\to Q}) \propto \frac{\exp(-C(\xi_{S\to Q}) - C^*_{G_R}(Q) + C^*_{G_R}(S))P(G_R)}{\sum_{G\in\mathcal{G}} \exp(-C(\xi_{S\to Q}) - C^*_G(Q) + C^*_G(S))P(G)} \tag{3}$$

An example of a legible trajectory is shown in Fig. 1.

Illegibility. Since legibility is defined as quickly *and* confidently guessing the correct goal after seeing a portion of a trajectory, we can define illegibility as the opposite: either delaying a confidently correct guess for as long as possible, *or* quickly and confidently guessing the wrong goal.

We penalize early correctness and encourage incorrectness with the following equation:

$$\textsc{Illegibility}(\xi) = \frac{\int (1 - P(G_R|\xi_{S\to\xi(t)}))(T-t)dt}{\int (T-t)dt} \tag{4}$$

Coincidentally, note:

$$\textsc{Legibility}(\xi) + \textsc{Illegibility}(\xi) = 1 \tag{5}$$

An example of an illegible trajectory (using the decoy strategy) is shown in Fig. 1.

3.2 Mixed-Motive Limited-Observability Legible Motion Planning.

Now, consider that our environment also has observers O. Each observer $o \in O$ has a visibility region $V_o \subseteq E$ and motive $M_o \in [-1, 1]$. For simplicity, we refer to observers with $M_o < 0$ as "negative" observers and those with $M_o \geq 0$ as "positive" observers.

We define $v_o : [0, T] \to \Xi$, where $v_o(i)$ is the portion of $\xi_{S\to\xi(i)}$ that lies within V_o. Also, let T_o be the number of timesteps during which ξ lies in V_o.

The goal of the Mixed-Motive Limited-Observability Legible Motion Planning problem (MMLO-LMP) is to find a trajectory ξ from S to G_R such that ξ is legible to positive observers but is illegible to negative observers. In mathematical terms, we want to find ξ such that:

$$\xi = \arg\max_{\xi\in\Xi} \sum_{o\in O, M_o\geq 0} M_o \times \textsc{Legibility}(v_o(T)) + \sum_{o\in O, M_o<0} |M_o| \times \textsc{Illegibility}(v_o(T)) \tag{6}$$

If ξ is empty, then we assume $P(G_R|\xi) = 0$. We list a few variants of the MMLO-LMP problem in Section 5.

3.3 Method

We propose a trajectory optimization-based approach called **DU**al **BI**ased limited **O**bservability **U**nified trajectory **S**olver (DUBIOUS) solve MMLO-LMP. We begin with ξ_0 as the straight-line trajectory from S to G_R. To maximize legibility, we iteratively update ξ using STOMP [11] with the cost function $\mathcal{F}$ presented below. This function is not only designed to maximize legibility or illegibility as defined by the MMLO-LMP statement, but also to balance the trajectory construction in regions where observers with opposing motives overlap.

Since STOMP requires a cost for each point in the trajectory, $\mathcal{F}$ takes a timestep i as input and uses $\xi_{S\to\xi(i)}$ to calculate costs.

We first consider the positive observers. We need to maximize legibility for each observer in only what lies within their visibility region. Thus, we re-formulate the legibility function and use the observer's motive as a scaling factor:

$$\mathcal{F}_{pos}(i) = \sum_{o \in O, M_o \geq 0} M_o \frac{\int P(G_R|v_o(i))(T_o - t)dt}{\int (T_o - t)dt} \tag{7}$$

This is, for each observer, the legibility score only considers points within that observer's visibility region.

Then, we consider the negative observers. We propose two strategies that an agent can follow: the decoy strategy and the ambiguous strategy.

The decoy strategy misleads the observer to quickly and confidently predict the agent is moving towards a decoy goal instead of the true goal We can measure this strategy similarly to [6] for some $G_D \neq G_R$ as:

$$\textsc{Illegibility-decoy}(\xi) = \frac{\int P(G_D|\xi_{S \to \xi(t)})(T - t)dt}{\int (T - t)dt} \tag{8}$$

The ambiguous strategy maintains ambiguity among multiple candidate goals by minimizing the average difference between $P(G_R|\xi)$ and $P(G|\xi)$ for all other $G \in \mathcal{G} \setminus \{G_R\}$. This extends the method in [6] to an arbitrary number of goals. So:

$$\textsc{Illegibility-ambiguous}(\xi) = \frac{\int (T - t)(1 - \frac{1}{|\mathcal{G}|} \sum_{G \in \mathcal{G} \setminus \{G_R\}} |P(G_R|\xi) - P(G|\xi)|)dt}{\int (T - t)dt} \tag{9}$$

We construct a cost function for negative observers similar to the one for positive observers, however, using the decoy goal G_D:

$$\mathcal{F}_{neg}(i) = \sum_{o \in O, M_o < 0} \alpha_{neg}|M_o| \frac{\int P(G_D|v_o(i))(T_o - t)dt}{\int (T_o - t)dt} \tag{10}$$

Even though the illegibility score above uses the better of the decoy and ambiguous strategies, DUBIOUS does not do this calculations. Rather, we leave the strategy type as a hyperparameter α_{neg}. If $\alpha_{neg} = 1$, then the agent maximizes legibility of the decoy goal for the observer. If $\alpha_{neg} = -1$, then the negative opponent's visibility region is costly to navigate through, and thus the agent avoids the negative opponent. We note that an avoidance strategy is an extremely efficient interpretation of the illegibility ambiguous strategy in limited-observability environments. This strategy is also quite similar to the approach of the minimum exposure problem [5].

While we have integrated motives into the costs above, they must be normalized over the total number of observers who can see $\xi(i)$. And lastly, STOMP's goal is to minimize cost, so we flip the sign on our cost function. Thus, our final cost function $\mathcal{F}$ is:

$$\mathcal{F}(i) = -\frac{\mathcal{F}_{pos}(i) + \mathcal{F}_{neg}(i)}{\sum_{o \in O} \mathbb{I}_{q_i \in V_o} |M_o|} \tag{11}$$

We optimize ξ using $\mathcal{F}$ with STOMP, initialized with the most predictable trajectory as defined in [7]. With our cost function, this is the straight-line trajectory from S to G_R. For this work, we assume that $P(G)$ is uniformly distributed across all goals.

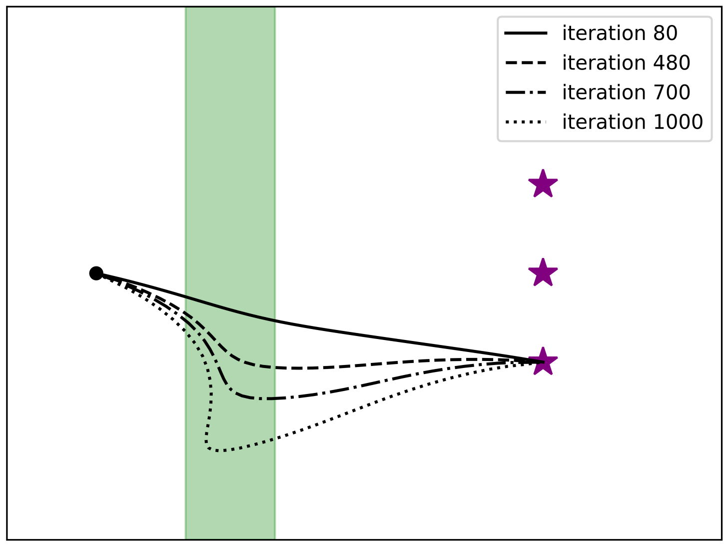

Fig. 2: STOMP Over-optimization. As the number of iterations increases, STOMP creates increasingly extreme trajectories. Here we show legible trajectories generated through the region in view of an observer with motive (green) +1 at iterations 80, 480, 700, and 1000.

4 Experiments & Results

In this section, we present examples of the MMLO-LMP problem, and we demonstrate that DUBIOUS can produce trajectories that achieve the problem objective. The environments we use allow for multiple different strategies for achieving our objective. We compare DUBIOUS against the full-environment legible and illegible trajectories (from Fig. 1). We find that DUBIOUS can produce multiple successful strategies and adapts its trajectories based on observer motive and goal location.

It is important to note that while our goal is to return a trajectory with maximum legibility or illegibility, there is a point after which a trajectory becomes *too* legible or illegible. The more iterations we run with any optimizer, the more extreme our output trajectories become. Thus, we do not report or compare to maximum legibility or illegibility; instead, for each experiment, unless otherwise noted, we present the trajectory generated after 1000 iterations of STOMP. Figure 2 provides some examples of overly-optimized trajectories.

Evaluation Metrics. For all cases, we qualitatively analyze our trajectories. Additionally, we define an observer's *guess* at timestep i of a trajectory's goal by calculating $\arg\max_{\xi \in \Xi} P(G_R | \xi_{S \to \xi(i)})$ as in [8]. We call an observer's guess "correct" when $P(G_R|\xi_{S\to\xi(t)}) \geq P(G_{\text{other}}|\xi_{S\to\xi(t)}) + 5\%$ for any other goal G_{other}- in other words, we require a 5% margin of confidence in an observer's guess.

Using this, we calculate and report 3 quantities for each trajectory:

- Percent correct guesses for each observer. This number represents for how many timesteps the observer confidently achieves the correct goal guess. This metric indicates whether the path is successful (or unsuccessful) at conveying its intended goal to the observers.
- Score functions. We report Legibility score (equation 1). Because of equation 5, we do not need to report Illegibility. For comparisons in negative observer regions, we report Illegibility-ambiguous or Illegibility-decoy, depending on the trajectory being analyzed (Equations 8 and 9, respectively). All of these scores are calculated with the probability function from Equation 3.
  It is important to note that these numbers are intended to be compared against each other. For example, the highest possible score for Legibility is 1, however, this requires $P(G_R|\xi_{S\to\xi(t)}) = 1$ for every possible value t. This is only true when there is 1 possible goal. The same applies for Illegibility-decoy. The maximum value of Illegibility-ambiguous is $\frac{1}{|\mathcal{G}|}$, as demonstrated in Table 1.

Legibility Strategy. Fig. 3 shows DUBIOUS's performance in environments with a single observer, with either +1 or -1 motive, and partial observability. Fig. 3a shows a +1 motive observer that can see the green portion of the environment. We show 3 candidate trajectories: the "Efficient" trajectory ("Baseline"), which minimizes $C(\xi)$, the maximally legible trajectory as calculated for Fig. 1, and a trajectory generated by DUBIOUS.

The maximally legible trajectory only optimizes for legibility and will not consider the observer, thus inadvertently avoiding the observer. As a result, the observer would not have enough information for goal prediction. In this case, the efficient straight line trajectory provides the observer with more information since it goes through its visibility region. DUBIOUS, however, employs legible motion while maximizing its time in the observer's view, enabling more correct guesses (see Table 1). The bottom portion of Figure 3a shows the probability distribution for $P(G_i|\xi_{S\to\xi(t)})$ for each goal as time progresses across the trajectory generated by DUBIOUS. We see that our trajectory distinguishes Goal 3 as the most likely goal early in the trajectory. However, given the observer's viewpoint, any path that maximizes time in that region will be more in the direction of the incorrect goals. For this reason, despite producing a higher percentage of correct guesses, our method has a lower legibility score than the efficient method. Our method does not determine how much time spent in the observer's view is optimal for legibility or illetibility. We leave determining the ideal trade-off for future work.

Illegible Decoy Strategy Fig. 3b shows 4 trajectories: the efficient trajectory, the maximum decoy legibility as calculated for Fig 1, a DUBIOUS trajectory with $\alpha_{neg} = 1$

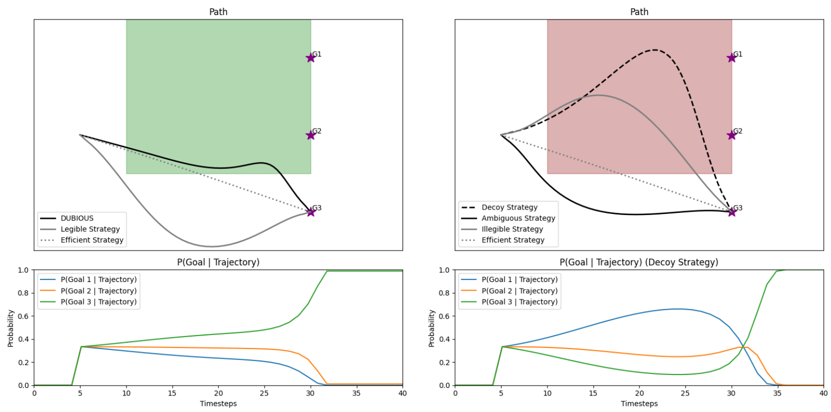

(a) +1 motive observer in green region and observer probabilities

(b) -1 motive observer in red region and probabilities

Fig. 3: (Top) Environments with one observer with partial legibility. Fig. 3a has one observer with +1 motive that can see the agent when it is in the green box. Fig. 3b has one observer with -1 motive that can see the agent's location when it is in the red box. (Bottom) Probabilities $P(G_i|\xi_{S\rightarrow\xi(t)})$ for each goal, over time, as seen by the observer. The ambiguous trajectory does not enter the observer's visibility region, so its probabilities are not reported.

('Decoy'), and a DUBIOUS trajectory with $\alpha_{neg} = -1$ ('Ambiguous'). We also show an interesting example of MMLO-LMP: the negative observer is blocking some of the goals. At some point along the trajectory, the agent's true motives must be revealed.

DUBIOUS can employ two different strategies for handling negative observers. The Decoy strategy initially moves toward a decoy goal, and changes its trajectory as late as possible. The ambiguous strategy attempts to minimize time spent in the observer's region. Thus, in this example, it moves down to go below the observer's visibility region. The bottom of Figure 3b shows the observer's goal guesses over the course of both of DUBIOUS's trajectories. The decoy strategy spends longer in the observer's region, but does not provide definitive information to the observer until it begins to turn. The ambiguous strategy enters the observer's region as late as possible, but reveals its goal instantly. Table 1 validates the efficacy of the decoy strategy in this environment, and further reaffirms that the ambiguous strategy is not ideal for this environment.

Illegible Ambiguous Strategy There are some scenarios where using an ambiguous or avoidance strategy is more effective than a decoy strategy. Figure 4 shows an environment with two overlapping observers in front of the goal, one positive and one negative. It is possible for the agent to bypass one or both regions in order to reach its goal. Figure 4a shows the agent traveling through the positive observer's region, providing as much information as possible about its goal. Then, it briefly travels through the negative observer's visibility region, leading the negative observer to infer that the agent

Table 1: Metrics for single observer environments illustrated in Fig. 3. Baseline for Fig. 3a and Fig. 3b are the max legible and max decoy legible trajectories from Fig. 1 respectively.

		% Correct	LEGIBILITY	ILLEG-DECOY	ILLEG-AMBIGUOUS
Figure 3a	Efficient	65%	**0.421**	—	—
	Max Legible	0%	0	—	—
	DUBIOUS	**72.5%**	0.407	—	—
Figure 3b	Efficient	65%	0.421	0.085	0.273
	Max Decoy Legible	34.5%	0.335	0.188	0.208
	DUBIOUS Decoy	**22.5%**	**0.219**	**0.360**	**0.293**
	DUBIOUS Ambiguous	**0%**	**0**	**0**	**0.33**

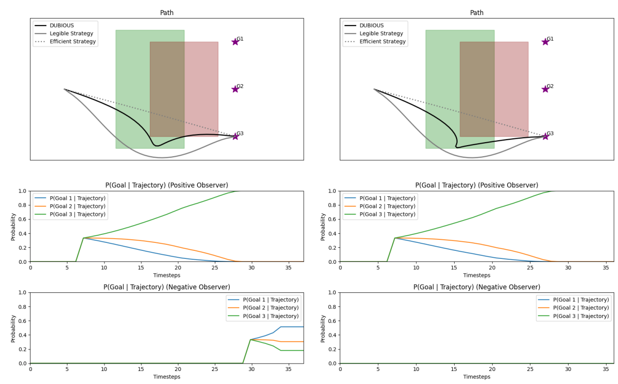

(a) DUBIOUS-decoy path and observer probabilities.

(b) DUBIOUS-ambiguous path and observer probabilities

Fig. 4: (Top) Environments with two observers with partial and overlapping legibility. (Bottom) Probabilities $P(G_i|\xi_{S\to\xi(t)})$ for each goal and each observer over time. The ambiguous trajectory does not enter the observer's visibility region, so its probabilities are not reported.

is moving to Goal 1. The agent leaves the visibility region quickly enough that the observer's inference does not change. Figure 4b travels through the positive observer's region in a similar manner, but skips the negative observer's region completely. This trajectory provides absolutely no information to the negative observer. While the decoy agent has successfully misled the negative observer with its trajectory, the ambiguous agent's strategy is more successful by avoiding the negative region completely.

Goals in Opposite Directions Figure 5 illustrates a challenging case where all potential goals are in opposing directions. The most legible trajectory and the most ambiguous trajectory would both be the straight-line shortest path in this scenario. The straight-line path is completely dissimilar from the most efficient path to any other goal, making it legible. Since there is no way to avoid the negative observer's visibility region, the

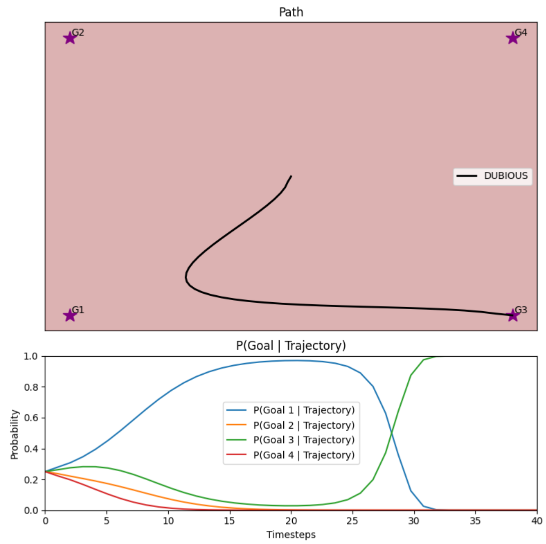

Fig. 5: Decoy strategy on environment with potential goals in opposite directions. The agent goes toward decoy goal 1 before pivoting to the true goal later in the trajectory.

ambiguous strategy will try to minimize time spent there, and thus time spend on the path in total. Therefore it will also choose the straight-line path. The decoy strategy, however, elicits some interesting behavior. The opposing directions of the goals make it easier to maximize the probability of the decoy goal. Indeed, for several timesteps at the middle of the trajectory, the observer is near 100% confident that the decoy is the true goal. It's only until almost timestep 30 that the observer shifts to predicting the correct goal. This makes the decoy method ideal for adversarial planning in such scenarios.

Overlapping Visibility Regions. Figures 4 and 6 show environments with two observers with opposing motives and overlapping visibility regions. We show the most legible and most decoy illegible paths from Fig. 1. For DUBIOUS trajectories, we show the decoy and ambiguous strategy for both experiments. Fig. 6a shows a visibility region of an observer with motive 1 within the visibility region of an observer with motive -0.25. Fig. 6b shows a visibility region of an observer with motive -1 within the visibility region of an observer with motive 0.25. In overlapping regions, DUBIOUS weights the legibility score of each observer according to their motive. In other words, in Fig. 6a, while it is not good for a trajectory to be legible in the view of the -0.25 observer, it is more important to be legible in view of the +1 observer. Hence, the DUBIOUS trajectory focuses on legibility in the overlapping region. On the other hand, in Fig. 6b, while it is good for a trajectory to be legible in the view of the 0.25 observer, it is more important to be illegible in view of the -1 observer. Hence, the DUBIOUS trajectory focuses on moving toward the decoy goal in the overlapping region.

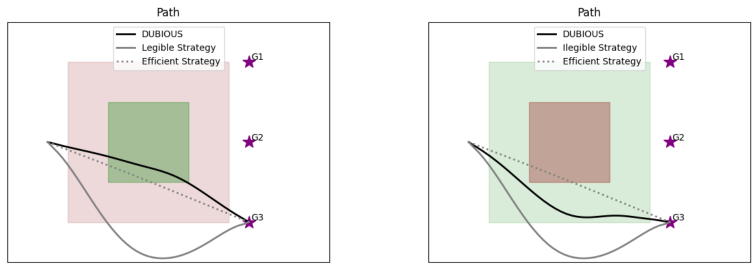

(a) +1 observer within -0.25 observer (b) -1 observer within +0.25 observer

Fig. 6: Both figures illustrate multi-observer environments with overlapping visibility regions. In Fig. 6a, the visibility region of a +1 observer is completely covered by the visibility region of a -0.25 observer. In Fig. 6b, the visibility region of a -1 observer is completely covered by the visibility region of a +0.25 observer.

5 Potential Variants of MMLO-LMP

Here, we presented a simplified version of the MMLO-LMP problem. Since this is a novel problem, there are many interesting variations. In this section, we detail the directions and potential applications of MMLO-LMP.

5.1 Scenario Modifications

Our current formulation is a bounded 2-D environment with no obstacles. In order to bring DUBIOUS solutions closer to the real-world situations, the environment must include obstacles. Aside from collision avoidance challenges, obstacles provide visibility region complexities. Obstacles also provide new potential strategies, such as hiding behind obstacles to avoid negative observers.

The problem formulation can also include a task planning component. Previous work [13] addresses legibility challenges in abstract task ordering problems. An example of a task planning extension to MMLO-LMP includes providing the agent with a list of goals that it has to reach. The agent would need to plan both the goal ordering and the trajectory such that an observer would need to be able to infer both.

Application of the MMLO-LMP problem to different domains highlights different challenges. For example, in a home robotics application, the robot will need to share space with humans and possibly animals. The robot will need to plan legible paths around the other agents in the home, though it will be more important to be legible for some agents (human beings) than for other agents (pets or other robots). Visibility will be limited by the human user's range of vision and their attention as well. By accounting for legibility only when being observed, the robot could dynamically balance efficiency and legibility in path planning.

Another application is a competitive game such as capture-the-flag. The robot will need to be legible about its intent (defending home base or grabbing the opponent flag) to teammates while hiding its intent from opponents. The environment may contain occlusions limiting observability. By accounting for where its actions are observable and by whom, an agent can effectively strategize with teammates and surprise the opponent.

5.2 Observer Modifications

We do not consider dynamic observers in our current formulation of MMLO-LMP. In fact, we do not even consider whether observers are physically present in the agent's environment. However, in many scenarios, particularly multi-robot or crowd navigation, the observers are also agents acting in the environment. This adds a concrete meaning to 'motive', where the predictions of observers enable them to take actions that help or hinder the agent in reaching its goal. What those actions are can vary widely by application. In an autonomous driving scenario, motive could encode how likely other drivers are to allow the agent to merge into their lane. In a competitive game, motive could encode which other agents are teammates and which are opponents.

We assume visibility regions of each observer are known and static, and certain within range. In many scenarios, however, the agent may be uncertain about what observers can see. The boundaries of an observer's visibility may be uncertain due to sensor uncertainty, resulting in more certain observations in the middle of the region and less certain observations at the boundaries. Observer certainty is important for an agent to compensate for (in legible scenarios) or take advantage of (in illegible scenarios). The observer may be concealing where their regions of visibility are (similar to hiding linear constraints in constrained policy optimization [1]), requiring an agent to infer their locations. Visibility regions may also evolve online. Active observers could take actions to increase their information gain, requiring the agent to adapt or even predict where the observers will move. This makes the agent an observer as well.

In this work we assumed that each observer made its inference of the agent's goal based only on the information from its own visibility region. In a modern world, observers in adversarial settings often have secure ways to communicate with each other. Thus, future modifications could include positive and negative observers working in teams. An observer's goal inference could be based on information from other agents and combined with its own observations. Some observers could be more trustworthy than others, meaning that the observer team would weight trustworthy observations higher. An agent's strategy could include how to create maximal certainty or uncertainty among an observer team.

In a real-world obstacle avoidance application as presented at the beginning of this work, inferring legibility could be used for collision avoidance. Collision avoidance in dynamic environments could be greatly improved by including trajectory prediction to the existing framework for goal prediction. Observers would need to accurately estimate the agent's trajectory to the goal. Thus, the agent would need to strike a balance between two opposing ideas: predictability and legibility.

5.3 Agent Modifications

Many improvements could be made to the agent's strategy. Paths which maximize legibility, or illegibility, are often not the most efficient. Realistically, an agent will have to balance legibility with efficiency if, for example, there are constraints on the amount of fuel, or the amount of time, the agent can spend getting to the goal. Our current solution plans trajectories offline since all aspects of the environment are static. However, as described above, in many real-world scenarios the goals, observers, and visibility regions will all change over time, requiring an online method. Additionally, we assume a static observer model based on a simple cost function. An observer's model may evolve over time if, for example, a negative observer learns that an agent will attempt a deceptive motion. By maintaining an estimate of the observer's model, the agent can evolve it's strategy to remain unpredictable, such as by selecting unexpected decoy goals, or switching from a decoy to an ambiguous strategy. One method for doing this is described using an RNN as the estimate of the observer model in [20]. As stated in the section above, in a typical multi-robot or crowd navigation scenario there will be multiple agent-observers, each with their own goals that may align or misalign with the agent's own goals. The motives of other agents may not be known, requiring the agent to learn over time which observers to be legible to and which to be illegible to.

We encourage readers to brainstorm and explore more MMLO-LMP modifications.

6 Conclusion

In this work, we presented and mathematically defined the Mixed-Motive Limited-Observability Legible Motion Planning problem. We introduced a trajectory optimization-based method, DUBIOUS, for finding trajectories that satisfy the requirements of being legible to some observers and illegible to others, while considering their visibility regions. We demonstrate the validity of DUBIOUS on multiple environments with varying numbers of observers and types of visibility regions. We provided multiple ways of producing illegible behavior (decoy or delay) and demonstrate the benefits and drawbacks of each method. Our experiments show that DUBIOUS can generate trajectories are equivalent to or outperform trajectories which maximize only for efficiency, legibility, and illegibility when observers have limited visibility or mixed motives. Lastly, we present many interesting directions for future work on the MMLO-LMP problem.

References

1. Benvenuti, A., Bialy, B., Dennis, M., Hale, M.: Guaranteed Feasibility in Differentially Private Linearly Constrained Convex Optimization. IEEE Control Systems Letters **8**, 2745–2750 (2024). https://doi.org/10.1109/LCSYS.2024.3513232
2. van Bevern, R., Fluschnik, T., Tsidulko, O.Y.: Parameterized algorithms and data reduction for the short secluded \$s\$-\$t\$-path problem. Networks **75**(1), 34–63 (Jan 2020). https://doi.org/10.1002/net.21904
3. Buchin, K., Polishchuk, V., Sedov, L., Voronov, R.: Geometric secluded paths and planar satisfiability (Mar 2020). https://doi.org/10.48550/arXiv.1902.06471
4. Chakraborti, T., Kulkarni, A., Sreedharan, S., Smith, D.E., Kambhampati, S.: Explicability? Legibility? Predictability? Transparency? Privacy? Security? The Emerging Landscape of Interpretable Agent Behavior. Proceedings of the International Conference on Automated Planning and Scheduling **29**, 86–96 (Jul 2019). https://doi.org/10.1609/icaps.v29i1.3463
5. Djidjev, H.N.: Efficient Computation of Minimum Exposure Paths in a Sensor Network Field. In: Aspnes, J., Scheideler, C., Arora, A., Madden, S. (eds.) Distributed Computing in Sensor Systems, vol. 4549, pp. 295–308. Springer Berlin Heidelberg, Berlin, Heidelberg (2007). https://doi.org/10.1007/978-3-540-73090-3_20
6. Dragan, A., Holladay, R., Srinivasa, S.: An Analysis of Deceptive Robot Motion. In: Robotics: Science and Systems X. Robotics: Science and Systems Foundation (Jul 2014). https://doi.org/10.15607/RSS.2014.X.010
7. Dragan, A., Srinivasa, S.: Generating Legible Motion. In: Robotics: Science and Systems IX. Robotics: Science and Systems Foundation (Jun 2013). https://doi.org/10.15607/RSS.2013.IX.024
8. Dragan, A.D., Lee, K.C., Srinivasa, S.S.: Legibility and predictability of robot motion. In: 2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI). pp. 301–308. IEEE, Tokyo, Japan (Mar 2013). https://doi.org/10.1109/HRI.2013.6483603
9. Feng, H., Luo, L., Wang, Y., Ye, M., Dong, R.: A novel minimal exposure path problem in wireless sensor networks and its solution algorithm. International Journal of Distributed Sensor Networks **12**(8), 155014771666424 (Aug 2016). https://doi.org/10.1177/1550147716664245
10. Gewali, L., Meng, A., Mitchell, J.S., Ntafos, S.: Path planning in 0/1/weighted regions with applications. In: Proceedings of the fourth annual symposium on Computational geometry. pp. 266–278 (1988)
11. Kalakrishnan, M., Chitta, S., Theodorou, E., Pastor, P., Schaal, S.: STOMP: Stochastic trajectory optimization for motion planning. In: 2011 IEEE International Conference on Robotics and Automation. pp. 4569–4574. IEEE, Shanghai, China (May 2011). https://doi.org/10.1109/ICRA.2011.5980280
12. Knepper, R.A., Mavrogiannis, C.I., Proft, J., Liang, C.: Implicit Communication in a Joint Action. In: Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction. pp. 283–292. ACM, Vienna Austria (Mar 2017). https://doi.org/10.1145/2909824.3020226
13. Kulkarni, A., Srivastava, S., Kambhampati, S.: A Unified Framework for Planning in Adversarial and Cooperative Environments. Proceedings of the AAAI Conference on Artificial Intelligence **33**(01), 2479–2487 (Jul 2019). https://doi.org/10.1609/aaai.v33i01.33012479
14. Kulkarni, A., Srivastava, S., Kambhampati, S.: Planning for Proactive Assistance in Environments with Partial Observability (Sep 2021). https://doi.org/10.48550/arXiv.2105.00525
15. Lebeck, N., Mølhave, T., Agarwal, P.K.: Computing highly occluded paths on a terrain. In: Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems. pp. 14–23. ACM, Orlando Florida (Nov 2013). https://doi.org/10.1145/2525314.2525363

16. Lebeck, N., Mølhave, T., Agarwal, P.K.: Computing highly occluded paths using a sparse network. In: Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems. pp. 3–12. ACM, Dallas Texas (Nov 2014). https://doi.org/10.1145/2666310.2666394
17. Luckow, M.J., Fluschnik, T.: On the computational complexity of length- and neighborhood-constrained path problems. Information Processing Letters **156**, 105913 (Apr 2020). https://doi.org/10.1016/j.ipl.2019.105913
18. Masters, P., Sardina, S.: Deceptive path-planning. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence. pp. 4368–4375. IJCAI'17, AAAI Press, Melbourne, Australia (Aug 2017)
19. Mavrogiannis, C.I., Thomason, W.B., Knepper, R.A.: Social Momentum: A Framework for Legible Navigation in Dynamic Multi-Agent Environments. In: Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction. pp. 361–369. ACM, Chicago IL USA (Feb 2018). https://doi.org/10.1145/3171221.3171255
20. Nichols, H., Jimenez, M., Goddard, Z., Sparapany, M., Boots, B., Mazumdar, A.: Adversarial Sampling-Based Motion Planning. IEEE Robotics and Automation Letters **7**(2), 4267–4274 (Apr 2022). https://doi.org/10.1109/LRA.2022.3148464
21. Taylor, A.V., Mamantov, E., Admoni, H.: Observer-Aware Legibility for Social Navigation. In: 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN). pp. 1115–1122 (Aug 2022). https://doi.org/10.1109/RO-MAN53752.2022.9900676
22. Veltri, G., Huang, Q., Qu, G., Potkonjak, M.: Minimal and maximal exposure path algorithms for wireless embedded sensor networks. In: Proceedings of the 1st International Conference on Embedded Networked Sensor Systems. pp. 40–50. ACM, Los Angeles California USA (Nov 2003). https://doi.org/10.1145/958491.958497
23. Wagner, A.R., Arkin, R.C.: Acting Deceptively: Providing Robots with the Capacity for Deception. International Journal of Social Robotics **3**(1), 5–26 (Jan 2011). https://doi.org/10.1007/s12369-010-0073-8
24. Yuning Song, Liang Liu, Huadong Ma, Vasilakos, A.V.: A Biology-Based Algorithm to Minimal Exposure Problem of Wireless Sensor Networks. IEEE Transactions on Network and Service Management **11**(3), 417–430 (Sep 2014). https://doi.org/10.1109/TNSM.2014.2346080
25. Zucker, M., Ratliff, N., Dragan, A.D., Pivtoraiko, M., Klingensmith, M., Dellin, C.M., Bagnell, J.A., Srinivasa, S.S.: CHOMP: Covariant Hamiltonian optimization for motion planning. Int. J. Rob. Res. **32**(9-10), 1164–1193 (Aug 2013). https://doi.org/10.1177/0278364913488805