Viewpoint-Based Legibility Optimization

Stefanos Nikolaidis*, Anca Dragan† and Siddhartha Srinivasa*
*The Robotics Institute, Carnegie Mellon University
†Department of Electrical Engineering and Computer Sciences, University of California, Berkeley
snikolai@andrew.cmu.edu, anca@berkeley.edu, siddh@cs.cmu.edu

***Abstract*—Much robotics research has focused on intent-expressive (legible) motion. However, algorithms that can autonomously generate legible motion have implicitly made the strong assumption of an *omniscient* observer, with access to the robot's configuration as it changes across time. In reality, human observers have a particular viewpoint, which biases the way they perceive the motion. In this work, we free robots from this assumption and introduce the notion of an observer with a specific point of view into legibility optimization. In doing so, we account for two factors: (1) depth uncertainty induced by a particular viewpoint, and (2) occlusions along the motion, during which (part of) the robot is hidden behind some object. We propose viewpoint and occlusion models that enable autonomous generation of viewpoint-based legible motions, and show through large-scale user studies that the produced motions are significantly more legible compared to those generated assuming an omniscient observer.**

I. Introduction

The development of new robotic systems that operate in the same physical space as people highlights the emerging need for robots that can integrate seamlessly into human group dynamics. An important factor of successful human-robot interactions is communication of intent. Previous studies have shown that humans have a universal tendency to interpret each other's actions as *intentional* and *goal-directed* [1], [2], [3], [4], [5], [6], [7]. Intent-expressive motion, sometimes referred to as legible [8], readable [9] or anticipatory [10], has repeatedly been cited as essential for robots that work around humans [11], [12], [13], [14].

Previous work [8] proposed a model to evaluate the legibility of a goal-directed motion based on the theory of action interpretation [15] in psychology, the result having strong motivations in the principle of rational action [16] and teleological reasoning [15]. This led to an algorithm that can autonomously generate legible motion, enabling a robot to successfully communicate its intent to the observer. Fig. 1 (top left) shows an example, in which the robot exaggerates the motion to the right to make to better convey that its goal is the bottle on the right. This works well when the observer looks at the scene from above, or sits across from the robot. Imagine, however, being side by side with the robot and observing the scene from the second viewpoint in the figure (top right). The exaggeration in the motion is barely perceptible in this case.

When generating this motion, the robot assumes an *omniscient* observer, with direct access to the robot's actual configuration as the robot is moving. This assumption is rarely accurate, however: a person on a wheelchair with an installed robotic arm, an EOD operator teleoperating a robot through an on-board camera and a mechanic working with a robotic assistant in the assembly line are examples of observers that have limited perception of the robot's actual configuration. The

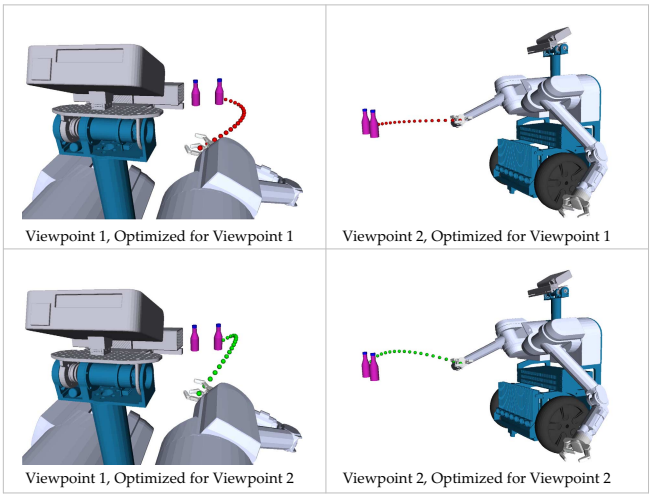

Fig. 1. Legible trajectories for two viewpoints. The red trajectory (top) is generated by an optimizer that assumes an omniscient observer. It exaggerates the motion to the right to convey that the robot is about to grasp the object on the right. This works well in the first viewpoint (left), but remains ambiguous in the second viewpoint (right). The green trajectory accounts for the observer's viewpoint, finding a way to exaggerate the motion that makes the robot's goal clear in the second viewpoint.

importance of viewpoint and its relation to the expressiveness of a motion has also deep roots in character animation [17], [18], [19], [20]. *Staging*, one of the twelve Disney principles of animation [21], emphasizes the importance of each action being presented "in the strongest and the simplest way ... communicating to the fullest extent with the viewers."

Whereas in animation the focus has mostly been on selecting a viewpoint that maximizes clarity, in robotics applications we typically cannot control for the viewpoint of the observer. Instead, we can enable robots to generate motion informed by and tailored to a given viewpoint. In this work, we introduce a framework for the mathematical underpinnings of this aspect of staging. Specifically, we introduce the notion of an observer with a specific point of view into legibility optimization (Fig. 1). This raises two challenges: accounting for that particular viewpoint and the induced depth uncertainty, and accounting for occlusions. We make the following three contributions: a viewpoint model, an occlusion model, and a user study evaluation testing the two.

Viewpoint Model. Our first contribution is legible motion optimization in the observer's viewpoint. We produce the motion via functional gradient optimization in the space of trajectories, echoing earlier works in motion planning [22], [23], [24], [25], [26], [27], [28], [29], with legibility as an

978-1-4673-8370-7/16/$31.00 © 2016 IEEE

optimization criterion. A key difference is that, rather than applying the principle of rational action in the world space, we model the viewplane of the observer and project the robot trajectory and goals to the viewplane. We then apply the principle of rational action in that space to generate *viewpoint-based* legible motion (Sec. III). Fig. 1 (bottom right) shows an example outcome of this optimization: the robot finds a different way of exaggerating, given the same constraints on efficiency, which is more legible in the second viewpoint.

Occlusion Model. Our second contribution is incorporating occlusions into legibility optimization. How does legibility change when the observer can only observe parts of it? We model the inference process that the human is making while the robot is about to enter, is inside, and has exited an area that occludes the trajectory. We then introduce generation of *occlusion-based* legible motion, taking into account occlusion in different parts of the robot trajectory (Sec. V).

User Study Evaluation. We conduct large-scale user studies to evaluate the legibility of trajectories generated with the proposed algorithms. Our results suggest that trajectories generated when accounting for the observer viewpoint, which we call *viewpoint-based*, and the occlusion region, called *occlusion-based*, are significantly more legible than trajectories that assume an omniscient observer (Sec. VI, VII).

In summary, we introduce the notion of an observer's viewpoint into robot motion planning. We propose a viewpoint model and an occlusion model that enable autonomous generation of viewpoint-based legible motions, and test the models through large-scale user-studies. Results suggest that the produced motions are significantly more legible compared to those generated assuming an omniscient observer.

II. Legible motion

In this section, we discuss the notion of legibility in robot motion, and briefly present the algorithm introduced in [22] for generation of legible motion through iterative optimization.

Measuring Legibility. If the observer sees the actor as a rational agent, applying the principle of rational action [16], they expect the actor to be efficient. Efficiency can be modeled via a cost functional $C : \Xi \rightarrow \mathbb{R}_+$ with lower costs signifying more "efficient" (and thus more expected/predictable to the observer) trajectories ξ. Using the principle of maximum entropy, C induces a probability density over trajectories $\xi \in \Xi$ given a goal G as $P(\xi|G) \propto \exp\big(-C(\xi)\big)$. Using Bayes' rule, we can model the probability that a human observer would assign to a goal candidate based on an ongoing trajectory ξ: $P(G|\xi) \propto P(\xi|G)P(G)$. The legibility of the trajectory tracks the probability assigned to the actual goal G_R across time: trajectories are more legible if this probability is higher, with more weight being given to the earlier parts of the trajectory via a function $f(t)$ (e.g. $f(t) = T - t$, with T the total time):

$$\textsc{Legibility}(\xi) = \frac{\int P(G_R|\xi_{S\rightarrow\xi(t)})f(t)dt}{\int f(t)dt} \qquad (1)$$

In [8], the authors showed that a robot motion with a higher Legibility score is indeed more legible to users, i.e. users can more quickly infer the robot's goal.

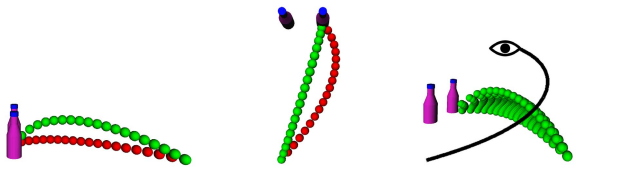

(a) Side view (b) Top view (c) Range of views from top to side

Fig. 2. (a-b) Incorporating observer viewpoint in trajectory generation. The green trajectory, optimized for the side view, appears more legible to a human viewing the motion from the side than the red trajectory, which is optimized for an omniscient observer. (c) Legible trajectories generated for 21 viewpoints starting from elevation 90°and azimuth 0°to elevation 10°and azimuth -90°, in decrements of 4°and 4.5°.

Generating Legible Motion. The Legibility score is a functional and following [22], we can generate a legible trajectory by starting from an initial trajectory ξ^0 and iteratively improving its score via functional gradient ascent.

$$\xi^{i+1} = \xi^i + \frac{1}{\eta}A^{-1}\bar{\nabla}\textsc{Legibility} \qquad (2)$$

A is used to measure the norm of a trajectory, $||\xi||_A^2 = \xi^T A\xi$.

Expected Cost C. Generating legible trajectories assumes access to the cost C that the human approximately expects the robot to optimize. Following previous work [22], [29], we use sum squared velocities as the cost functional $C(\xi) = \frac{1}{2}\int ||\xi'(t)||^2 dt$. This cost encourages smooth trajectories that go straight to the goal, matching user expectations [8].

III. Human Observer Viewpoint

We described in Sec. II that we model the observer as expecting the robot to move efficiently, where efficiency is modeled via a cost functional C. In this section, we derive a viewpoint-based legibility functional by first generalizing the notion of trajectory efficiency for the case of different viewpoints, and then using that to model the probability that a human observer would assign to the robot's goal.

A. Viewpoint-based Cost.

First, we apply a transformation on trajectories, $T : \Xi \rightarrow \bar{\Xi}$, so that $\bar{\xi} = T(\xi)$ is the projected trajectory onto the viewplane of the observer. We do this by transforming the waypoints of trajectory ξ from world to camera space, and then projecting them onto the 2D viewing plane [30].

We then define a cost function $\bar{C}(\xi)$ which computes the cost of the projected trajectories onto the observer viewing plane: $\bar{C}(\xi) = C(\bar{\xi}) = C(T(\xi))$. With the new definition, we posit that the observer's expectation is not on the actual trajectory, but on its projection on the observer's viewplane.

B. Viewpoint-based Legibility.

We compute the viewpoint-based Legibility from Eq. (1), where $P(G|\xi) \propto P(\xi|G)P(G)$. The probability distribution over trajectories given the goal G is now defined as: $P(\xi|G) \propto \exp\big(-\bar{C}(\xi)\big)$, i.e. it is defined in the observer's viewpoint. We can then use the algorithm of Sec. II to generate legible trajectories that account for the observer's viewpoint.

C. Implications.

In Fig. 2 we show an example of a generated trajectory considering the viewpoint of an observer, who views the scene from the side. The red trajectory assumes an omniscient observer and exaggerates the motion to the right. But because of the viewpoint, the red trajectory does not clearly convey the goal. Optimizing for viewpoint-based legibility yields the green trajectory instead. The bottle that is further away from the observer appears on the 2D viewplane to be higher than the other. Thus, by exaggerating upwards, instead of to the right, the green trajectory appears more legible to the observer from their viewpoint (but it is less legible to an omniscient observer). Both green and red trajectories have the same actual cost $C(\xi)$.

IV. Legibility Computation under Occlusion

We saw in Sec. II that the legibility of a trajectory ξ is computed by the integral of $P(G|\xi_{S\to\xi(t)})$ along time (Eq. (1)). In cases where the observer cannot see the ongoing trajectory, the probability distribution over goals is different from the case of full observability. In this section, we derive how $P(G|\xi_{S\to\xi(t)})$ changes in such a case.

An occlusion region induces three parts to the trajectory ξ: (i) the robot has not entered the occlusion area, (ii) the robot is inside the area, and (iii) the robot has exited the area (Fig 3). We use the derivation on these cases as building blocks to compute $P(G|\xi_{S\to\xi(t)})$ for the general case, where the robot enters and exits multiple occlusion regions.

A. Robot has not entered the occlusion region

When the robot trajectory thus far has been fully visible to the observer (Fig. 3(a)), we compute $P(\xi_{S\to Q}|G)$ as the ratio of all trajectories from S to G that pass through $\xi_{S\to Q}$ to all trajectories from S to G [8]:

$$P(\xi_{S\to Q}|G) = \frac{\exp\big(-C(\xi_{S\to Q})\big)\int_{\xi_{Q\to G}}\exp\big(-C(\xi_{Q\to G})\big)}{\int_{\xi_{S\to G}}\exp\big(-C(\xi_{S\to G})\big)} \tag{3}$$

where $\xi_{S\to Q}$ is the trajectory segment from the start configuration S to the current configuration Q.

We can simplify Eq. (3) by approximating $C(\xi_{X\to Y})$ by its second order Taylor series expansion around $\xi^*_{X\to Y} = \arg\min_{\xi_{X\to Y}} C(\xi_{X\to Y})$ [31]:

$$\int_{\xi_{X\to Y}}\exp\big(-C(\xi_{X\to Y})\big) \approx \exp\big(-V_Y(X)\big)\frac{\sqrt{2\pi^k}}{\sqrt{|H_{X\to Y}|}} \tag{4}$$

where $H_{X\to Y}$ the Hessian of the cost function around $\xi^*_{X\to Y}$, and $V_Y(X) = \min_{\xi\in\Xi_{X\to Y}} C(\xi)$. When the cost is quadratic, the Hessian is constant and Eq. (3) simplifies to

$$P(\xi_{S\to Q}|G) = \frac{\exp\big(-C(\xi_{S\to Q}) - V_G(Q)\big)}{\exp\big(-V_G(S)\big)} \tag{5}$$

The probability $P(\xi_{S\to Q}|G)$ is equal to the cost of going through the current trajectory and continuing optimally to the goal, over the optimal cost from the start to the goal.

Using Bayes' rule, we can compute the probability of the actual goal of the robot G_R given the observed trajectory:

$$P(G_R|\xi_{S\to Q}) = \frac{1}{Z}\frac{\exp\big(-C(\xi_{S\to Q}) - V_{G_R}(Q)\big)}{\exp\big(-V_{G_R}(S)\big)}P(G_R) \tag{6}$$

with Z a normalizer across all candidate goals G.

B. Robot is inside the occlusion region

Let us assume that at time t, the robot is at some configuration $\xi(t)$ inside the hidden area, as shown in Fig. 3(b). The observer has viewed a trajectory snippet $\xi_{S\to U}$, where U is the last visible point before the robot enters the area. We compute $P(\xi_{S\to\xi(t)}|G)$ as the ratio of all trajectories from S that include $\xi_{S\to U}$ and pass through the occlusion region. To do that, we integrate over all configurations x_t of the set X_t inside the occlusion region that the robot may be at time t:

$$\begin{aligned} P(\xi_{S\to\xi(t)}|G) =& \frac{P(\xi_{S\to U})}{\int_{\xi_{S\to G}} P(\xi_{S\to G})}\times \\ & \int_{X_t} P_{x_t}\int_{\xi_{x_t\to G}} P(\xi_{x_t\to G}) \end{aligned} \tag{7}$$

P_{x_t} is the probability of being in state x_t at time t when the last visible state was U. It is equal to the integral of the probabilities of all trajectories starting from U and ending at x_t, at timepoint t: $P_{x_t} = \int_{U\to x_t} P(\xi_{U\to x_t})$. Eq. (7) can be written as:

$$\begin{aligned} P(\xi_{S\to\xi(t)}|G) =& \frac{\exp\big(-C(\xi_{S\to U})\big)}{\int_{\xi_{S\to G}}\exp\big(-C(\xi_{S\to G})\big)}\times \\ & \int_{X_t} P_{x_t}\int_{\xi_{x_t\to G}}\exp\big(-C(\xi_{x_t\to G})\big) \end{aligned} \tag{8}$$

Whereas the observer does not directly observe the current robot configuration, they know how much time has passed since the robot became occluded. To account for that, we use a Hidden Markov Model and compute P_{x_t} via filtering. The HMM has initial state U and transition probabilities:

$$\begin{aligned} P(s_{t+1}|s_t, G) =& \frac{1}{Z}\exp\big(-C(\xi_{s_t\to s_{t+1}})\big) \\ & \int_{\xi_{s_{t+1}\to G}}\exp\big(-C(\xi_{s_{t+1}\to G})\big) \end{aligned} \tag{9}$$

Z is a normalizing constant. Eq. (9) stems from the principle of maximum entropy, applied for the case of a human predicting a trajectory segment from state s_t to s_{t+1}, given a goal G [32]. In this work, we assume a binary observation model, where no observation is given when the robot is inside the occluded area and a perfect observation is given when the robot is outside.

By approximating $C(\xi_{s_{t+1}\to G})$ as in Sec. IV-A and including the Hessian in Z, Eq. (9) simplifies to:

$$P(s_{t+1}|G, s_t) = \frac{1}{Z}\exp\big(-C(\xi_{s_t\to s_{t+1}})\big)\exp\big(-V_G(s_{t+1})\big) \tag{10}$$

Eq. (8) becomes:

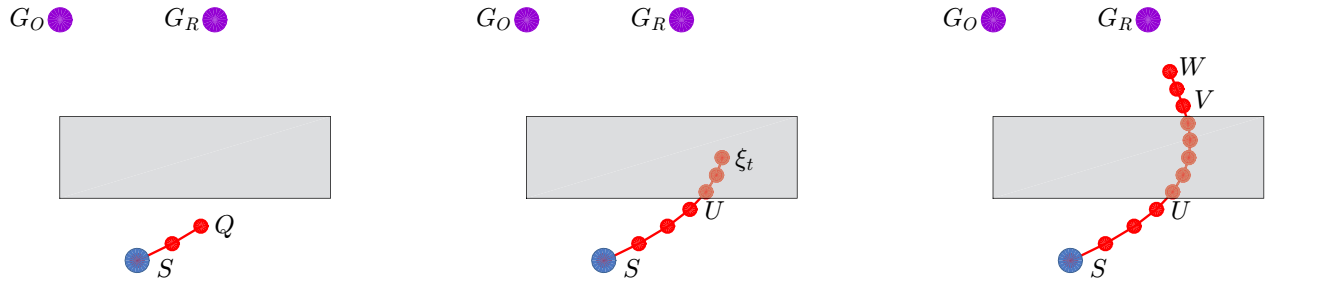

(a) Robot has not entered the occlusion region (b) Robot is inside the occlusion region (c) Robot has exited the occlusion region

Fig. 3. Different cases when robot path is occluded. The human does not see the robot, when it is inside the occlusion region.

$$P(\xi_{S\to\xi(t)}|G) = \frac{\exp\big(-C(\xi_{S\to U})\big)\int_{X_t} P_{x_t}\exp\big(-V_G(X_t)\big)}{\exp\big(-V_G(S)\big)} \tag{11}$$

Using Bayes Rule, we have for the robot goal G_R:

$$P(G_R|\xi_{S\to\xi(t)}) = \frac{1}{Z}\exp\big(V_{G_R}(S)\big) \int_{X_t} P_{x_t}\exp\big(-V_{G_R}(X_t)\big)P(G_R) \tag{12}$$

Therefore, at time t we can estimate the inference the human is making about a goal G, having observed a part of the robot trajectory from S to U.

C. Robot has exited the occlusion region

We assume that part of the robot trajectory has been occluded but the robot is now visible (Fig. 3(c)).

$$\begin{aligned} P(\xi_{S\to W}|G) =& P(\xi_{S\to V}\wedge\xi_{V\to W}|G) \\ =& P(\xi_{S\to V}|G)P(\xi_{V\to W}|G,\xi_{S\to V}) \end{aligned} \tag{13}$$

For additive cost functions:

$$P(\xi_{S\to W}|G) = P(\xi_{S\to V}|G)P(\xi_{V\to W}|G) \tag{14}$$

To compute the first term $P(\xi_{S\to V}|G)$ in Eq. (14), we note that the human has seen $\xi_{S\to U}$, but has no information on the actual trajectory followed inside the occlusion region from U to V. Therefore, we integrate over all possible trajectories from U to V inside the occlusion region:

$$P(\xi_{S\to V}|G) = \exp\big(-C(\xi_{S\to U})\big)\times \frac{\int_{\xi_{U\to V}}\exp\big(-C(\xi_{U\to V})\big)\int_{\xi_{V\to G}}\exp\big(-C(\xi_{V\to G})\big)}{\int_{\xi_{S\to G}}\exp\big(-C(\xi_{S\to G})\big)} \tag{15}$$

We have for $P(\xi_{V\to W}|G)$:

$$P(\xi_{V\to W}|G) = \int_{\xi_{W\to G}}\exp\big(-C(\xi_{W\to G})\big)\times \frac{\exp\big(-C(\xi_{V\to W})\big)}{\int_{\xi_{V\to G}}\exp\big(-C(\xi_{V\to G})\big)} \tag{16}$$

Using the the approximation described in Sec. IV-A on Eq. (15) and (16), Eq. (14) becomes:

$$P(\xi_{S\to W}|G) = \frac{\exp\big(-C(\xi_{S\to U})-V_V(U)\big)}{\exp\big(-V_G(S)\big)c}\times \exp\big(-C(\xi_{V\to W})-V_G(W)\big) \tag{17}$$

where c is a constant accounting for the Hessian term in Eq. (4). Using Bayes' rule, we have

$$P(G_R|\xi_{S\to W}) = \frac{1}{Z}\frac{\exp\big(-C(\xi_{S\to U})-V_V(U)\big)}{\exp\big(-V_{G_R}(S)\big)}\times \exp\big(-C(\xi_{V\to W})-V_{G_R}(W)\big)P(G_R) \tag{18}$$

Note that, once the robot exits the hidden area and becomes visible again at point V, the human inference of the goal $P(G_R|\xi_{S\to W})$ is independent of the actual trajectory $\xi_{U\to V}$ followed inside the invisible area. We show that this is actually the case in User Study I of Sec. VI.

D. General Case

In the general case, the robot can execute any arbitrary trajectory, entering and exiting different occlusion regions. In that case, we can combine the derivations from the previous sections: We split the trajectory into segments, using as dividing points the entry and exit points of each occlusion region previously entered, and use an HMM to do inference for the part of the trajectory that is currently occluded, if any.

For instance, assume that after point W, the robot becomes hidden again. As in Sec. IV-B and IV-C, we integrate over all possible trajectories from U to V inside the first occlusion region, as well as over all configuration points x_t of the set X_t inside the new occlusion region that the robot is at time t:

$$P(\xi_{S\to\xi(t)}|G) = \frac{\exp\big(-C(\xi_{S\to U})\big)}{\int_{\xi_{S\to G}}\exp\big(-C(\xi_{S\to G})\big)}\times \int_{\xi_{U\to V}}\exp\big(-C(\xi_{U\to V})\big)\exp\big(-C(\xi_{V\to W})\big)\times \int_{X_t} P_{x_t}\int_{\xi_{x_t\to G}}\exp\big(-C(\xi_{x_t\to G})\big) \tag{19}$$

$$P(G_R|\xi_{S\to\xi(t)}) = \frac{1}{Z}\frac{\exp\big(-C(\xi_{S\to U}) - C(\xi_{V\to W})\big)}{\exp\big(-V_{G_R}(S)\big)}\times$$
$$\exp\big(-V_V(U)\big)\int_{X_t} P_{x_t}\exp\big(-V_{G_R}(x_t)\big)P(G_R) \tag{20}$$

We compute P_{x_t} using filtering on an HMM with start state W. We see that the terms $\exp\big(-C(\xi_{S\to U})\big)$, $\exp\big(-V_V(U)\big)$, $\exp\big(-C(\xi_{V\to W})\big)$ that refer to the trajectory part $\xi_{S\to W}$, are goal-independent and cancel out when divided by the probability sum Z over all goals G. Therefore, $P(G_R|\xi_{S\to\xi(t)})$ can be described using Eq. (12) in Sec. IV-B. We can show likewise that when the robot has exited an occlusion region for the second time, the probability of the goal does not depend on the past trajectory, therefore is independent from the number of occlusion regions and the trajectory followed inside.

E. Implications

Fig. 4 shows the probability $P(G_R|\xi_{S\to\xi(t)})$, as it changes with time for the robot trajectory of Fig. 3. We compare against $P_{omni}(G_R|\xi_{S\to\xi(t)})$ computed for the same trajectory assuming that the robot is always visible [8]. We see that until the robot enters the occluded region the two probabilities are identical. Then, the probability of the goal computed with the proposed model increases at a lower rate, since the human observer has no reason to believe that the robot will keep exaggerating to the right, but instead makes an estimate on the current state of the robot, as described in Sec. IV-B. When the robot exits the occluded region, the probability of the correct goal becomes the same as in the omniscient case, because now the observer has all the information relevant to the inference.

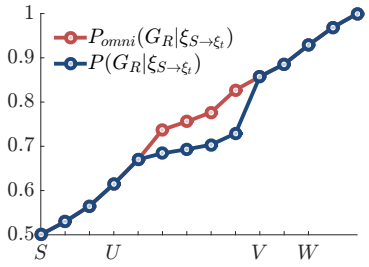

Fig. 4. Computation of the probability of the goal, as it is changes with time, for the robot trajectory in Fig. 3. The probability of G_R computed using the proposed model is in blue color and the one assuming an omniscient observer is in red color. $P(G_R|\xi_{S\to\xi_t})$ keeps increasing while the robot is occluded since, based on the entry point of the robot and the transition model of the HMM, the most probable states are closer to the G_R.

V. Generating Viewpoint-Based Legible Motion

As described in Sec. II, Eq. (2), the update rule for ξ^{i+1} is:

$$\xi^{i+1} = \xi^i + \frac{1}{\eta}A^{-1}\bar{\nabla}\text{LEGIBILITY} \tag{21}$$

The LEGIBILITY score is given from Eq. (1). We describe below the generation of legible motion for each of the three cases of Sec. IV: the robot has not entered the occlusion area, the robot is inside the area, and the robot has exited the area.

A. Robot has not entered the occlusion region

The gradient of legibility functional was shown to be [22]:

$$\bar{\nabla}\text{LEGIBILITY}(t) = K * \frac{\exp\big(V_{G_R}(S) - V_{G_R}(\xi(t))\big)}{\big(\sum_G \exp\big(V_G(S) - V_G(\xi(t))\big)\big)^2}$$
$$\sum_G\left(\frac{\exp\big(-V_G(\xi(t))\big)}{\exp\big(-V_G(S)\big)}(V'_G(\xi(t)) - V'_{G_R}(\xi(t)))\right)P(G_R)f(t) \tag{22}$$

where $K = \frac{1}{\int f(t)dt}$, with $f(t)$ a function that assigns more weight to the earlier parts of the trajectory.

B. Robot is inside the occlusion region

$\bar{\nabla}$LEGIBILITY can be computed as follows:

$$\bar{\nabla}\text{LEGIBILITY} = K\left(\frac{\partial\mathcal{P}}{\partial\xi} - \frac{d}{dt}\frac{\partial\mathcal{P}}{\partial\xi'}\right) \tag{23}$$

with $\mathcal{P}(\xi(t),t) = P(G_R|\xi_{S\to\xi(t)})f(t)$, where $P(G_R|\xi_{S\to\xi(t)})$ is given by Eq. (12), and $K = \frac{1}{\int f(t)dt}$.

$$\frac{\delta\mathcal{P}}{\delta\xi}(\xi(t),t) = \frac{g'h - h'g}{h^2}P(G_R)f(t) \tag{24}$$

with g, h as follows:
$g = \exp\big(V_{G_R}(S)\big)\int P_{x_t}\exp\big(-V_{G_R}(x_t)\big)$
and
$h = \sum_G \exp\big(V_G(S)\big)\int P_{x_t}\exp\big(-V_G(x_t)\big)$.
When the robot is inside the occlusion region, the values of g and h do not depend on ξ, thus $g' = 0$, $h' = 0$. This leads to $\frac{\partial\mathcal{P}}{\partial\xi}(\xi(t),t) = 0$. Additionally, $\frac{d}{dt}\frac{\delta\mathcal{P}}{\delta\xi'} = 0$, since $\mathcal{P}$ is not a function of ξ'. Therefore:

$$\bar{\nabla}\text{LEGIBILITY}(t) = 0 \tag{25}$$

C. Robot has exited the occlusion region

We saw in Sec. IV-C, that

$$P(G_R|\xi_{S\to W}) = \frac{1}{Z}\frac{\exp\big(-C(\xi_{S\to U}) - V_V(U)\big)}{\exp\big(-V_{G_R}(S)\big)}\times$$
$$\exp\big(-C(\xi_{V\to W}) - V_{G_R}(W)\big)P(G_R)$$

Including in the normalization term Z the goal-independent terms $V_V(U)$ and $\exp\big(-C(\xi_{S\to U})\big)$ we end up with:

$$P(G_R|\xi_{S\to W}) = \frac{1}{Z}\frac{\exp\big(-C(\xi_{V\to W}) - V_{G_R}(W)\big)P(G_R)}{\exp\big(-V_{G_R}(S)\big)}$$

This equation is similar to Eq. (6), therefore the derivation of the gradient is identical to the one followed in Sec. V-A and described in [22].

D. Implications

We observe that the LEGIBILITY gradient of the trajectory is 0 when the robot is inside the occlusion region. This result reflects our intuition, that when the robot is not visible, changes in its trajectory do not result in a change in its legibility, as perceived by the human observer. Therefore, the robot will exaggerate its motion when outside the occlusion region, but follow a straight line inside. Fig. 5(a) shows the trajectory generated with the proposed algorithm (green) and the one ignoring occlusion (red). Both trajectories have the same cost. The lack of exaggeration in the occlusion region allows the green trajectory to exaggerate more in the visible regions. Although the gradient for the trajectory points inside the region is 0, the points are not stationary, as the gradient of the visible points outside the region propagates to the points inside through the metric tensor A in the update rule of Eq. (2).

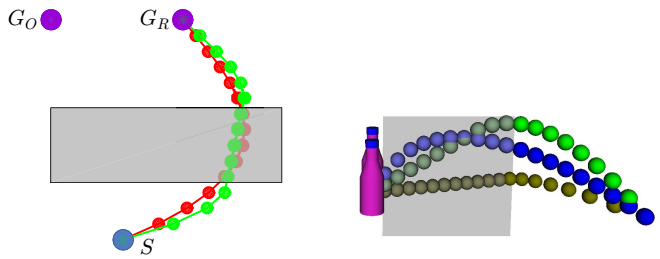

Fig. 5. (a) Generated occlusion-based trajectory in green color and baseline trajectory in red color. Both trajectories have the same cost C. (b) A trajectory optimized for both viewpoint and occlusion (green), compared to viewpoint only (blue) and occlusion only (olive). The goal of the robot is the bottle that is farther away and appears to be higher than the other.

Integrating Occlusion and Human Observer Viewpoint. The results of Sec. IV and V hold for the general case where the observer does not have information on parts of the robot trajectory. We can combine the occlusion and viewpoint models by using in the equations of this section the cost $\bar{C}(\xi)$ of the trajectory projected to the observer's viewplane, instead of $C(\xi)$, as described in Sec. III. Fig. 5(b) shows an example of a viewpoint-based only, an occlusion-based only and a combined trajectory, all of the same cost C. The occlusion-based only trajectory ignores the viewpoint by exaggerating in the horizontal plane, and the viewpoint-based only keeps exaggerating inside the occlusion region. The combined one does all the exaggeration before entering the occlusion region and goes straight to the goal when inside the region.

VI. FROM THEORY TO USERS

We conduct a large-scale user study to evaluate the legibility of trajectories generated with the proposed algorithms. We test that trajectories generated when accounting for the observer viewpoint (Sec. III), called *viewpoint-based*, and the occlusion region (Sec. V), called *occlusion-based*, are significantly more legible compared to trajectories assuming an omniscient observer. To control for confounds arising from timing and efficiency, all trajectories have the same duration and cost.

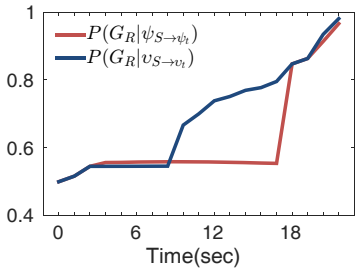

Fig. 6. Computation of the probability of the goal with the model of Sec. IV, for the occlusion-based trajectory υ (blue color) and baseline trajectory ψ (red color) of User Study I. We note that the confidence score plot for User Study I in Fig. 8 (left) is in line with the proposed model.

A. Hypotheses

H1. *Participants will find the occlusion-based trajectory significantly more legible than the baseline.*
H2. *Participants will find the viewpoint-based trajectory significantly more legible than the baseline.*

B. User Study I: Occlusion Test

We present users with a point robot entering an occlusion region with a small "window," as shown in Fig. 7. The actual goal G_R is on the right. We run the algorithm of Sec. V, and produce a trajectory that goes through the window, since passing through the window generates a steep rise in the perceived probability of the right goal, as shown in Fig. 6. On the other hand, the baseline algorithm that assumes an omniscient observer, ignoring occlusion, exaggerates to the right missing the window.

Dependent Measures. We measured the legibility of the trajectory that takes into account occlusion, which we call "occlusion-based," and the one that assumes an omniscient observer, called "baseline." We showed users three videos of the trajectory, stopping the video after 6, 12 and 18 sec (out of 22 sec). We asked them to predict the robot goal (Fig. 7), briefly explain their prediction and rate their confidence on their prediction on a Likert scale from 1 to 7. We then combined the three answers into a single legibility score metric, by computing a weighted sum of the ratings. The weights decreased linearly with the time of the rating, thus assigning more weight to earlier responses [22]. We assigned a score of 0 to incorrect predictions.

Subject Allocation. We chose a between-subjects design in order to not bias the users with trajectories from previous conditions. We recruited 100 participants through Amazon's Mechanical Turk service, and took several measures to ensure reliability of the results. All participants were located in the USA to avoid language barriers, and they all had an approval rate of over 95%. We asked users a control question that tested their attention to the task, and eliminated data associated with wrong answers to this question, as well as incomplete data, resulting in a total of 94 samples.

Analysis. An unpaired two-tailed t-test supported H1, showing that the robot trajectory had a significant effect on the legibility score of the users ($t(92) = 2.648, p = 0.010$). As shown in Fig. 8, the confidence score and prediction success rate in the

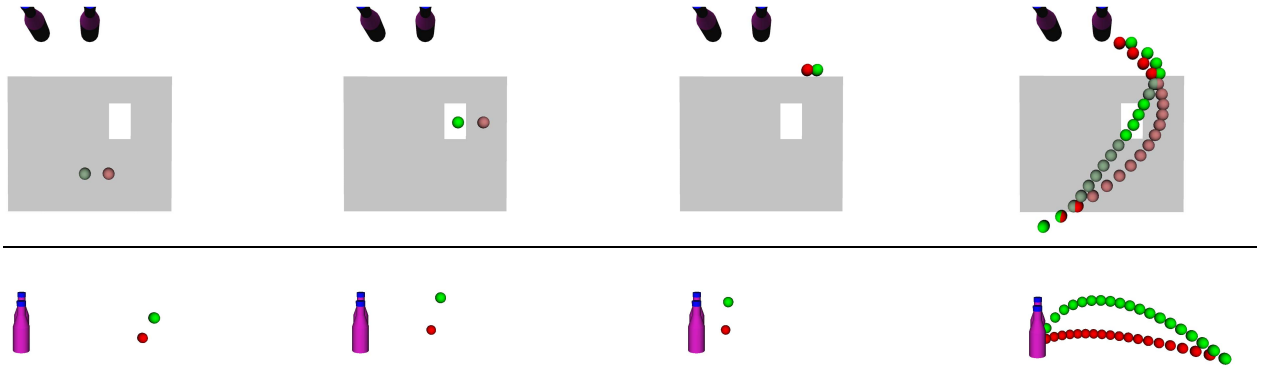

(a) T = 6 sec (b) T = 12 sec (c) T = 18 sec (d) Complete Trajectories

Fig. 7. Environment and trajectory waypoints of User Study I(top) and II(bottom). Left/Center: The human subject is asked to predict the actual goal of the robot in each of the three timepoints. Right: Complete trajectories. In User Study I/II, the occlusion-based/viewpoint-based trajectory is in green color and the baseline in red color. Both trajectories are overlayed for clarity, but each user watched only one of the trajectories. In User study I, the user views the trajectories from the top and they do not see the trajectory waypoints that are inside the grey occlusion region. In User Study II, the user views the trajectories from the side.

first and third timepoints of the trajectory are nearly identical. The gap between the scores of the two trajectories increases in the second timepoint, and becomes very small in the third timepoint when the robot exits the occlusion area. This result is in line with the computed inference of the goal (Fig. 6), based on the model described in Sec. IV. This supports the prediction of our model that, once the robot is visible again, the path inside the hidden area does not affect human inference.

C. User Study II: Viewpoint Test

We used the same environment of the previous study, but this time we removed the occlusion area, and changed the viewpoint to 9° elevation and -90° azimuth (side view, Fig. 7).

We then compared the viewpoint-based trajectory to the baseline.

Dependent Measures. Users were asked to predict the goal of the robot, briefly explain their prediction, and rate their confidence, identically to Case Study I. A legibility score was computed in the same way as in Sec. VI-B.

Subject Allocation. The subject allocation and data selection were the same as in Case Study I, resulting in a total of 196 samples from Amazon Mechanical Turk.

Analysis. An unpaired two-tailed t-test supported H2. The legibility score of the viewpoint-based trajectory was significantly higher $(t(194) = 2.06, p = 0.040)$. We noticed that the score decreased for both trajectories at the third timepoint (Fig. 8), when the point-robot moved downwards as shown in Fig. 7(c). Looking at the open-ended responses, several subjects noted that in the third timepoint the robot "appeared to be changing goals." Similarly to how going up allowed subjects to infer that the robot was moving farther away, moving down seemed to have the opposite effect. We hypothesize that this effect could be mitigated by having the robot going on top of the goal and then moving vertically downwards, rather than following a curved path to the end, and leave this for future work.

VII. Generalization to Arm Motion

The previous studies revealed that the mathematical model for generation of viewpoint-based legible motion performs well in practice. In this section, we show that the proposed mathematical model generalizes beyond a simple 2D robot character, by applying the model to the 7DOF right arm of a simulated mobile bi-manual manipulator (Fig. 9). We generated trajectories for a side (elevation, azimuth: 22°, -90°) and a $3/4$ viewpoint (29°, -118°).

Hypothesis. *Participants will find the viewpoint-based trajectories significantly more legible than the baseline.*

Dependent Measures. Users were asked to predict the goal of the robot, briefly explain their prediction, and rate their confidence, identically to Case Study II. A legibility score was computed in the same way as in Sec. VI-B.

Subject Allocation. We chose a between-subjects design, with different groups for each trajectory. The subject and data selection were the same as in Case Study II, resulting in a total of 189 samples for the side viewpoint and 89 samples for the $3/4$ viewpoint.

Analysis. A two-way ANOVA showed no significant interaction effects between the viewpoint factor and viewpoint optimization factor. Additionally, the test showed a statistically significant main effect of both viewpoint $(F(1, 274) = 14.40, p < 0.001)$ and viewpoint optimization $(F(1, 274) = 43.17, p < 0.001)$ on trajectory legibility, supporting our hypothesis. Interestingly, the average success rate of the goal prediction in the omniscient trajectory was 20%, as opposed to 44% for the viewpoint trajectory. Whereas these results show a significant improvement in legibility, they also indicate that for some viewpoints simply changing the motion of the arm might not be sufficient to guarantee successful inference. In these cases, adding *secondary actions*, a technique widely used in animation [21], by jointly optimizing the trajectories of the robot arm, head, and torso, could significanly improve the communication of intent, but we leave this for future work.

VIII. Discussion

Limitations. In real-world environments, humans use environmental cues to assess the 3D position of objects. Additionally, even if the robot end-effector is occluded, the observer may make an estimate of its position based on the angle of the

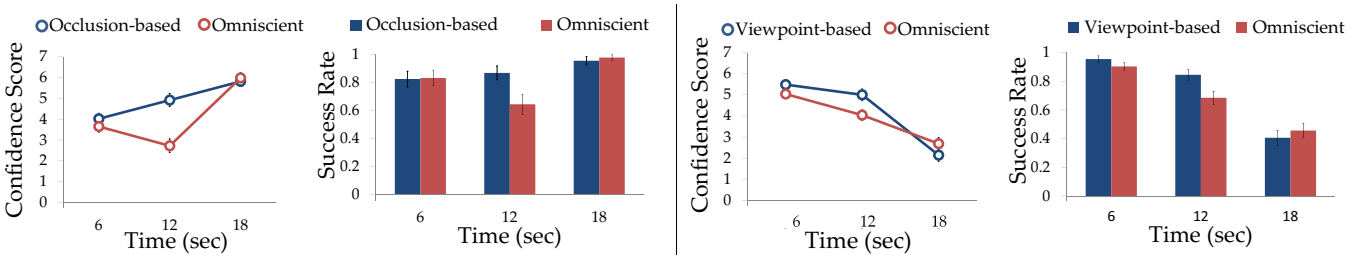

Fig. 8. Confidence scores and success rate of goal prediction for the three timepoints of User Study I (left) and II (right), averaged over all subjects of each group. We assigned a confidence score of 0 to incorrect goal predictions.

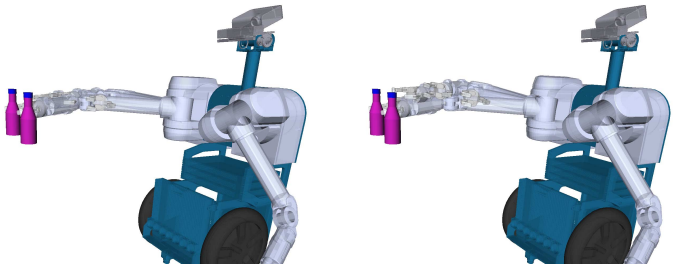

Fig. 9. A full-arm depiction of omniscient (left) and viewpoint-based (right) legible trajectories for the $3/4$ viewpoint.

base and the noise of the motors. We plan to address these issues by extending the proposed model to partial observability domains, where the spectator is making inferences based on noisy estimates of the robot trajectory.

Implications. Motion clarity has been a key element of Disney's animation principle *staging* [21]. While at Disney this was initially achieved by staging actions in silhouette, 3D animation evolved the principle, by placing the camera so that the entire motion would be clear to the spectator [17]. We showed how to automatically generate viewpoint-based legible motions, but the proposed models could also be used to solve the inverse problem: given a motion, find the legibility-optimal viewpoint. We are additionally excited to explore applications of this work in theater and puppeteering. Finally, exploiting occlusion opens an exciting range of possibilities for generation of purposefully ambiguous or deceptive motion [33].

Acknowledgments

This work was funded by the DARPA SIMPLEX program through ARO contract number 67904LSDRP, National Institute of Health R01 (#R01EB019335), National Science Foundation CPS (#1544797), and the Office of Naval Research. We also acknowledge the Onassis Foundation as a sponsor.

References

[1] D. A. Baldwin, J. A. Baird, M. M. Saylor, and M. A. Clark, "Infants parse dynamic action," *Child Development*, 2001.
[2] A. L. Woodward, "Infants selectively encode the goal object of an actor's reach," *Cognition*, 1998.
[3] B. Sodian and C. Thoermer, "Infants' understanding of looking, pointing, and reaching as cues to goal-directed action," *J. Cog. Dev.*, 2004.
[4] E. J. Carter, J. K. Hodgins, and D. H. Rakison, "Exploring the neural correlates of goal-directed action and intention understanding." *NeuroImage*, 2011.
[5] M. Carpenter, N. K., Tomasello, G. M., Butterworth, and C. Moore, "Social cognition, joint attention, and communcative competence from 9 to 15 months of age," *Monogr. Soc. Res. Child Dev.*
[6] T. Behne, M. Carpenter, J. Call, and M. Tomasello, "Unwilling Versus Unable: Infants' Understanding of Intentional Action," *Dev. Psychol.*, 2005.
[7] A. N. Meltzoff, "Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children," *Dev. Psychol.*, 1995.
[8] A. Dragan, K. Lee, and S. Srinivasa, "Legibility and predictability of robot motion," in *HRI*, 2013.
[9] L. Takayama, D. Dooley, and W. Ju, "Expressing thought: improving robot readability with animation principles," in *HRI*, 2011.
[10] M. Gielniak and A. Thomaz, "Generating anticipation in robot motion," in *RO-MAN*, 2011.
[11] M. Beetz *et al.*, "Generality and legibility in mobile manipulation," *Auton. Robots*, 2010.
[12] R. Alami *et al.*, "Safe and dependable physical human-robot interaction in anthropic domains: State of the art and challenges," in *IROS*, 2006.
[13] G. Klien, D. Woods, J. Bradshaw, R. Hoffman, and P. Feltovich, "Ten challenges for making automation a "team player" in joint human-agent activity," *Intelligent Systems*, 2004.
[14] A. Dragan, N. Ratliff, and S. Srinivasa, "Manipulation planning with goal sets using constrained trajectory optimization," in *ICRA*, May 2011.
[15] G. Csibra and G. Gergely, "Obsessed with goals: Functions and mechanisms of teleological interpretation of actions in humans," *Acta Psychologica*, vol. 124, no. 1, pp. 60 – 78, 2007.
[16] G. Gergely, Z. Nadasdy, G. Csibra, and S. Biro, "Taking the intentional stance at 12 months of age," *Cognition*, 1995.
[17] J. Lasseter, "Principles of traditional animation applied to 3d computer animation," in *ACM SIGGRAPH*, 1987.
[18] T. Shiratori, M. Mahler, W. Trezevant, and J. K. Hodgins, "Expressing animated performances through puppeteering," in *IEEE 3DUI*, 2013.
[19] M. Dontcheva, G. Yngve, and Z. Popović, "Layered acting for character animation," in *ACM Trans. Graph.*, 2003.
[20] T. Porter and G. Susman, "On site: Creating lifelike characters in pixar movies," *Commun. ACM*, 2000.
[21] F. Thomas and O. Johnston, *The illusion of life: Disney animation*. Hyperion New York, 1995.
[22] A. Dragan and S. Srinivasa, "Generating legible motion," in *RSS*, 2013.
[23] S. Quinlan, "The real-time modification of collision-free paths," Ph.D. dissertation, Stanford University, 1994.
[24] O. Brock and O. Khatib, "Elastic strips: A framework for motion generation in human environments," *IJRR*, 2002.
[25] M. Toussaint, "Robot trajectory optimization using approximate inference," in *ICML*, 2009.
[26] C. Igel, M. Toussaint, and W. Weishui, "Rprop using the natural gradient," *Trends and Applications in Constructive Approximation*, 2005.
[27] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "STOMP: Stochastic trajectory optimization for motion planning," in *ICRA*, 2011.
[28] E. Todorov and W. Li, "A generalized iterative lqg method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *ACC*, 2005.
[29] N. Ratliff, M. Zucker, J. A. D. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *ICRA*, 2009.
[30] B. Horn, *Robot vision*. MIT press, 1986.
[31] A. Dragan and S. Srinivasa, "Formalizing assistive teleoperation," in *RSS*, July 2012.
[32] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. Dey, "Maximum entropy inverse reinforcement learning," in *AAAI*, 2008.
[33] A. Dragan, R. Holladay, and S. Srinivasa, "An analysis of deceptive robot motion," in *RSS*, 2014.