

Generating Legible Motion

Anca Dragan and Siddhartha Srinivasa
The Robotics Institute, Carnegie Mellon University
{adragan, siddh}@cs.cmu.edu

Abstract—Legible motion — motion that communicates its intent to a human observer — is crucial for enabling seamless human-robot collaboration. In this paper, we propose a functional gradient optimization technique for autonomously generating legible motion. Our algorithm optimizes a legibility metric inspired by the psychology of action interpretation in humans, resulting in motion trajectories that purposefully deviate from what an observer would expect in order to better convey intent. A trust region constraint on the optimization ensures that the motion does not become too surprising or unpredictable to the observer.

Our studies with novice users that evaluate the resulting trajectories support the applicability of our method and of such a trust region. They show that within the region, legibility as measured in practice does significantly increase. Outside of it, however, the trajectory becomes confusing and the users’ confidence in knowing the robot’s intent significantly decreases.

I. INTRODUCTION

Robots perform remarkable superhuman acts of manipulation in our factories. Industrial manipulators are more precise than humans. But, how many of us would want to share a workspace with a robot? In contrast, we routinely share workspaces with less precise humans.

A key reason for this is *communication*. In addition to performing our tasks, we continuously communicate with each other via numerous channels, understanding each other’s intentions and responding appropriately.

We have a universal tendency to interpret each other’s actions as *intentional* and *goal-directed* [4, 6, 10, 11, 25, 31, 37], and our ability to communicate our intentions plays a crucial role in our collaborations [34].

The focus of our paper is to provide robots with this very ability: the ability to communicate their intent. Among the different channels, we focus on motion — a natural channel for communication in physical collaboration:

Our goal is to enable robots to *generate* intent-expressive motion — motion that is *legible*.

Legible motion, sometimes referred to as readable [32] or anticipatory [18], has repeatedly been cited as essential for robots that work around humans [2, 5, 14, 23]. Imagine, for example, the robot from Fig.1 cleaning up a dining room table together with a human collaborator. As it is reaching for one of the two remaining objects on the table, the human infers its goal and reaches for the other. Moving legibly means enabling the human to quickly and confidently make these type of predictions.

In order to achieve this, the robot needs a model of what the human will infer as he is observing the motion. In our prior work ([13], summarized in Sec.

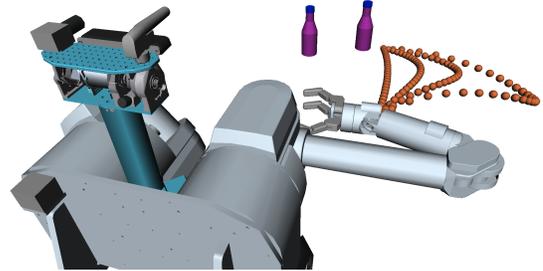
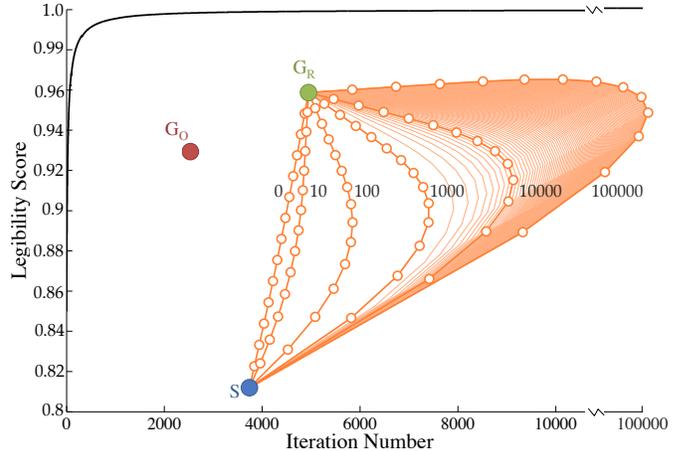


Fig. 1. The legibility optimization process for a reaching task. By moving the trajectory to the right, the robot is more clear about its intent to grasp the object on the right.

III), we proposed such a model based on the theory of action interpretation [12] in psychology, the result having strong motivations in the principle of rational action [17]. However, although this model enables us to *evaluate* how legible a motion trajectory is, and has been shown to correlate with legibility in practice, it does not enable us to *generate* trajectories that are legible.

Generation. Going from evaluation to generation means going beyond modeling the observer’s goal inference, to creating motion that results in the correct goal being inferred, i.e. going from "I can tell that you believe I am grasping this.", to "I know how to *make* you believe I am grasping this".

Our first contribution is to generate legible motion via functional gradient optimization in the space of trajectories (Sec. IV), echoing earlier works in motion planning [9, 21, 22, 27, 29, 33, 35], now with legibility as an optimization criterion. Fig.1 depicts this optimization process: by exaggerating the motion to the right, the robot makes the other goal option, G_O , far less likely to be inferred by the observer that the correct goal G_R .

Trust Region. The ability to optimize the legibility cri-

terion led us to a surprising observation: that there are cases in which the trajectory becomes too *unpredictable*. As our user studies show (Sec. VII, as well as our previous work [13]), *some unpredictability is often necessary to convey intent* — it is unpredictability beyond a threshold (like the outermost trajectory in Fig.1) that confuses users and lowers their confidence in what the robot is doing.

This phenomenon stems from the difficulty in capturing how humans make inferences when faced with high levels of unpredictability [30]. We address this fundamental limitation by prohibiting the optimizer to “travel to uncharted territory”, i.e. go outside of the region in which its assumptions have support — we call this a “trust region” of predictability. Our user studies indicate that indeed, there exists a size for this region in which legibility improves in practice, but outside of which the users’ confidence in knowing the robot’s goal drops. This is our second contribution.

New Research Threads. Finally, we use our optimization procedure to provide more insight into legibility, and discuss possible approaches for addressing the remaining challenges of producing legible motion in high dimensional spaces. Key among them is that the robot must learn what makes its motion predictable to a particular user — or perhaps do the opposite, and train the user’s very definition of predictability.

II. NOTATION: FUNCTIONALS ON TRAJECTORIES

In this paper, we focus on goal-directed motion. Here, a robot executes a trajectory $\zeta : \mathbb{R} \rightarrow \mathcal{Q}$, lying in a Hilbert space of trajectories Ξ . ζ starts at a configuration S and ends at a goal G_R from a set of possible goals \mathcal{G} .

Measuring how legible a trajectory is requires a *functional*, mapping trajectory functions in Ξ to scores in \mathbb{R}_+ .

III. DEFINING LEGIBLE MOTION

Legibility and predictability are fundamental concepts in this paper: legible motion conveys intent (Fig.1,100), while predictable motion matches expectation (Fig.1,0). Our previous work [13] formalized these notions and proposed mathematical models that measure how legible or predictable a motion is, which we summarize below. A main result is that the two properties are fundamentally different, and that *a departure from predictability is often necessary to increase the legibility of the motion*. We tested this theoretical finding in practice, in a user study on three characters (including the robot from Fig.1).

Definitions. As the observer is watching a trajectory, he continually makes an inference as to what the goal of the trajectory might be. In the psychology of action interpretation, this is referred to as an “action-to-goal” inference [12], which we denote here

$$\mathcal{I}_L : \Xi \rightarrow \mathcal{G}$$

Legible motion enables an observer to *confidently* infer the *correct* goal configuration G_R after observing only a snippet of the trajectory, $\xi_{S \rightarrow Q}$, from the start S to the

configuration at a time t , $Q = \zeta(t)$: $\mathcal{I}_L(\xi_{S \rightarrow Q}) = G_R$. The *quicker* this happens (i.e. the smaller t is), the more legible the trajectory is.

On the other hand, if the observer knows that the goal is G_R , they anticipate what trajectory this might result in — an opposite, “goal-to-action” inference [12], which we denote here

$$\mathcal{I}_P : \mathcal{G} \rightarrow \Xi$$

Predictable motion is motion for which the trajectory $\xi_{S \rightarrow G_R}$ matches this inference: $\mathcal{I}_P(G_R) = \xi_{S \rightarrow G_R}$.

Inferences based on cost. If the observer sees the actor as a rational agent, applying the principle of rational action [17], then 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. We discuss the challenges of finding C , which is an input to our method, in Sec. IX.

Given C and applying the principle of maximum entropy, we can model the user as expecting a trajectory ζ with probability $P(\zeta) \propto \exp(-C[\zeta])$ (lower cost is exponentially preferred), leading to a score for predictability:

$$\text{PREDICTABILITY}[\zeta] = \exp(-C[\zeta]) \quad (1)$$

Therefore, the observer infers the trajectory with highest probability, i.e. lowest cost, given a goal G — the most *predictable* trajectory:

$$\mathcal{I}_P(G) = \arg \min_{\zeta \in \Xi_{S \rightarrow G}} C[\zeta] \quad (2)$$

Given an ongoing trajectory $\xi_{S \rightarrow Q}$, the observer infers the most probable goal:

$$\mathcal{I}_L(\xi_{S \rightarrow Q}) = \arg \max_{G \in \mathcal{G}} P(G|\xi_{S \rightarrow Q}) \quad (3)$$

where $P(G|\xi_{S \rightarrow Q})$ can be approximated as

$$P(G_R|\xi_{S \rightarrow Q}) = \frac{1}{Z} \frac{\exp(-C[\xi_{S \rightarrow Q}] - V_{G_R}(Q))}{\exp(-V_{G_R}(S))} P(G_R) \quad (4)$$

with Z a normalizer across \mathcal{G} and $V_G(q) = \min_{\zeta \in \Xi_{S \rightarrow q}} C[\zeta]$ [15].

As action interpretation suggests it should [12], this evaluates how efficient going to a goal is through the observed trajectory snippet $\xi_{S \rightarrow Q}$ relative to the optimal trajectory.

The Legibility Functional. The score for legibility tracks the probability assigned to the actual goal G_R across the trajectory: 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):

$$\text{LEGIBILITY}[\zeta] = \frac{\int P(G_R|\xi_{S \rightarrow \zeta(t)})f(t)dt}{\int f(t)dt} \quad (5)$$

with the goal probability from (4). While predictability optimizes C , legibility optimizes this more complex score, intimately related to C but focused on conveying intent.

In [13], we tested that a motion with higher LEGIBILITY score is indeed more legible to users, for both a point robot, as well as the robot in Fig.1.

IV. GENERATING LEGIBLE MOTION

In this section, we show how to generate legible trajectories via trajectory optimization of LEGIBILITY.

A. Gradient Ascent

In order to maximize the LEGIBILITY functional, we start from an initial trajectory ζ_0 and iteratively improve its score via functional gradient ascent (Fig.1).

At every iteration i , we maximize the regularized first order Taylor series approximation of LEGIBILITY about the current trajectory ζ_i :

$$\begin{aligned} \zeta_{i+1} = \arg \max_{\zeta} & \text{LEGIBILITY}[\zeta_i] + \bar{\nabla} \text{LEGIBILITY}^T(\zeta - \zeta_i) \\ & - \frac{\eta}{2} \|\zeta - \zeta_i\|_M^2 \end{aligned} \quad (6)$$

with $\frac{\eta}{2} \|\zeta - \zeta_i\|_M^2$ a regularizer restricting the norm of the displacement $\zeta - \zeta_i$ w.r.t. an M , as in [29].

By taking the functional gradient of (6) and setting it to 0, we obtain the following update rule for ζ_{i+1} :

$$\zeta_{i+1} = \zeta_i + \frac{1}{\eta} M^{-1} \bar{\nabla} \text{LEGIBILITY} \quad (7)$$

To find $\bar{\nabla} \text{LEGIBILITY}$, let $\mathcal{P}(\zeta(t), t) = P(G_R | \zeta_{S \rightarrow \zeta(t)}) f(t)$ and $K = \frac{1}{\int f(t) dt}$. The legibility score is then

$$\text{LEGIBILITY}[\zeta] = K \int \mathcal{P}(\zeta(t), t) dt \quad (8)$$

and

$$\bar{\nabla} \text{LEGIBILITY} = K \left(\frac{\partial \mathcal{P}}{\partial \zeta} - \frac{d}{dt} \frac{\partial \mathcal{P}}{\partial \zeta'} \right) \quad (9)$$

\mathcal{P} is not a function of ζ' , thus $\frac{d}{dt} \frac{\partial \mathcal{P}}{\partial \zeta'} = 0$.

$$\frac{\partial \mathcal{P}}{\partial \zeta}(\zeta(t), t) = \frac{g'h - h'g}{h^2} P(G_R) f(t) \quad (10)$$

with $g = \exp(V_{G_R}(S) - V_{G_R}(Q))$ and $h = \sum_G \exp(V_G(S) - V_G(Q)) P(G)$, which after a few simplifications becomes

$$\begin{aligned} \frac{\partial \mathcal{P}}{\partial \zeta}(\zeta(t), t) &= \frac{\exp(V_{G_R}(S) - V_{G_R}(\zeta(t)))}{(\sum_G \exp(V_G(S) - V_G(\zeta(t)))) P(G)} P(G_R) \\ &\sum_G \left(\frac{\exp(-V_G(\zeta(t))) P(G)}{\exp(-V_G(S))} (V'_G(\zeta(t)) - V'_{G_R}(\zeta(t))) \right) f(t) \end{aligned} \quad (11)$$

Finally,

$$\bar{\nabla} \text{LEGIBILITY}(t) = K \frac{\partial \mathcal{P}}{\partial \zeta}(\zeta(t), t) \quad (12)$$

with $\frac{\partial \mathcal{P}}{\partial \zeta}(\zeta(t), t)$ from (11).

B. Parameters

Legibility depends on certain parameters: we list here what we use throughout the examples in the paper.

• **Trajectory parametrization.** We parametrized the trajectory as a vector of waypoint configurations.

• **Expected cost C .** We used sum squared velocities as the cost functional C capturing the user's expectation, $C[\zeta] = \frac{1}{2} \int \zeta'(t)^2 dt$. This cost, frequently used to encourage trajectory smoothness [29], produces trajectories that reach directly toward the goal, in line with users' expectation for a point robot (as evidenced in [13]). It also allows for an analytical V_G and its gradient, making the optimization process very fast.

• **Norms w.r.t. M .** We used the Hessian of C for M . As a result, the update rule in (7) propagates local gradient changes linearly to the rest of the trajectory.

• **Trajectory initialization.** We set $\zeta_0 = \arg \min_{\zeta} C[\zeta]$: we initialize with the most predictable trajectory, treating C (from (2)) as a prior.

V. THE UNPREDICTABILITY OF LEGIBILITY

Automating the generation of legible motion led us to a surprising observation: in some cases, by optimizing the legibility functional, one can become arbitrarily unpredictable.

Proof: Our gradient derivation in (11) enables us to construct cases in which this occurs. In a two-goal case like in Fig.1, with our example C (Sec. IV-B), the gradient for each trajectory configuration points in the direction $G_R - G_O$ and has positive magnitude everywhere but at ∞ , where $C[\zeta] = \infty$. Fig.2 (red) plots C across iterations.

The reason for this peculiarity is that the model for how observers make inferences in (3) and (4) fails to capture how humans make inferences in highly unpredictable situations. In reality, observers might get confused by the robot's behavior and stop reasoning about the robot's possible goals the way the model assumes they would — comparing the sub-optimality of its actions with respect to each of them. Instead, they might start believing that the robot is malfunctioning [30] or that it is not pursuing any of the goals — this is supported by our user studies in Sec. VII, which show that this belief significantly increases at higher C costs.

This complexity of action interpretation in humans, which is difficult to capture in a goal prediction model, can significantly affect the legibility of the generated trajectories in practice. Optimizing the legibility score outside of a certain threshold for predictability can actually lower the legibility of the motion as measured with real users (as it does in our study in Sec. VII-B). Unpredictability above a certain level can also be detrimental to the collaboration process in general [2, 20, 26].

We propose to address these issues by only allowing optimization of legibility where the model holds, i.e. where predictability is sufficiently high. We call this a "trust region" of predictability — a constraint that bounds the domain of trajectories, but that does so w.r.t. the cost functional C , resulting in $C[\zeta] \leq \beta$:

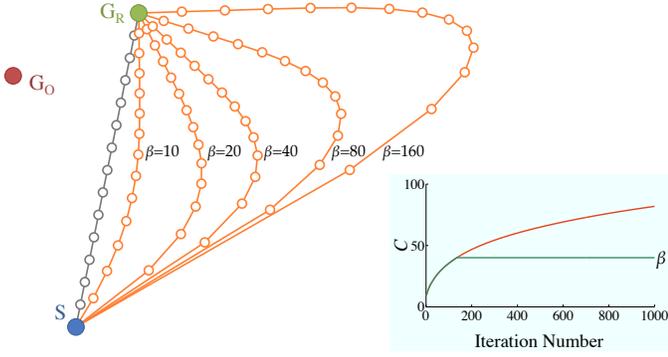


Fig. 2. The expected (or predictable) trajectory in gray, and the legible trajectories for different trust region sizes in orange. On the right, the cost C over the iterations in the unconstrained case (red) and constrained case (green).

The legibility model can only be trusted inside this trust region.

The parameter β , as our study will show, is identifiable by its effect on legibility as measured with users — the point at which further optimization of the legibility functional makes the trajectory less legible in practice.

VI. CONSTRAINED LEGIBILITY OPTIMIZATION

In order to prevent the legibility optimization from producing motion that is too unpredictable, we define a trust region of predictability, constraining the trajectory to stay below a maximum cost in C during the optimization in (6):

$$\begin{aligned} \zeta_{i+1} = \arg \max_{\zeta} & \text{LEGIBILITY}[\zeta_i] + \bar{\nabla} \text{LEGIBILITY}^T(\zeta - \zeta_i) \\ & - \frac{\eta}{2} \|\zeta - \zeta_i\|_M^2 \\ \text{s.t. } & C[\zeta] \leq \beta \end{aligned} \quad (13)$$

To solve this, we linearize the constraint, which now becomes $\bar{\nabla} C^T(\zeta - \zeta_i) + C[\zeta_i] \leq \beta$. The Lagrangian is

$$\begin{aligned} \mathcal{L}[\zeta, \lambda] = & \text{LEGIBILITY}[\zeta_i] + \bar{\nabla} \text{LEGIBILITY}^T(\zeta - \zeta_i) \\ & - \frac{\eta}{2} \|\zeta - \zeta_i\|_M^2 + \lambda(\beta - \bar{\nabla} C^T(\zeta - \zeta_i) - C[\zeta_i]) \end{aligned} \quad (14)$$

with the following KKT conditions:

$$\bar{\nabla} \text{LEGIBILITY} - \eta M(\zeta - \zeta_i) - \bar{\nabla} C \lambda = 0 \quad (15)$$

$$\lambda(\beta - \bar{\nabla} C^T(\zeta - \zeta_i) - C[\zeta_i]) = 0 \quad (16)$$

$$\lambda \geq 0 \quad (17)$$

$$C[\zeta] \leq \beta \quad (18)$$

Inactive constraint: $\lambda = 0$ and

$$\zeta_{i+1} = \zeta_i + \frac{1}{\eta} M^{-1} \bar{\nabla} \text{LEGIBILITY} \quad (19)$$

Active constraint: The constraint becomes an equality constraint on the trajectory. The derivation for ζ_{i+1} is analogous to [14], using the LEGIBILITY functional as

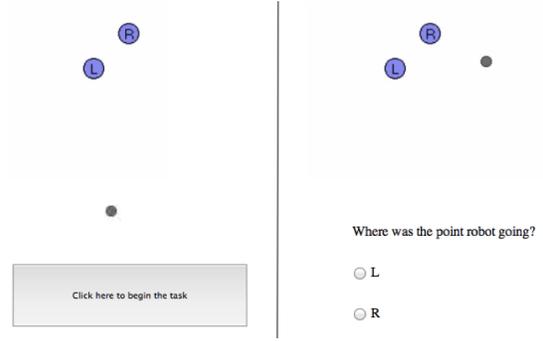


Fig. 3. We measure legibility by measuring at what time point along the trajectory users feel confident enough to provide a goal prediction, as well as whether the prediction is correct.

opposed to the classical cost used by the CHOMP motion planner[29]. From (15)

$$\zeta_{i+1} = \zeta_i + \frac{1}{\eta} M^{-1} \underbrace{(\bar{\nabla} \text{LEGIBILITY} - \lambda \bar{\nabla} C)}_{\bar{\nabla}(\text{LEGIBILITY} - \lambda C)} \quad (20)$$

Note that this is the functional gradient of LEGIBILITY with an additional (linear) *regularizer* λC penalizing *unpredictability*. Substituting in (16) to get the value for λ and using (15) again, we obtain a new update rule:

$$\begin{aligned} \zeta_{i+1} = \zeta_i + \frac{1}{\eta} M^{-1} \bar{\nabla} \text{LEGIBILITY} - \\ \underbrace{\frac{1}{\eta} M^{-1} \bar{\nabla} C (\bar{\nabla} C^T M^{-1} \bar{\nabla} C)^{-1} \bar{\nabla} C^T M^{-1} \bar{\nabla} \text{LEGIBILITY} -}_{\text{projection on } \bar{\nabla} C^T(\zeta - \zeta_i) = 0} \\ \underbrace{M^{-1} \bar{\nabla} C (\bar{\nabla} C^T M^{-1} \bar{\nabla} C)^{-1} (C[\zeta_i] - \beta)}_{\text{offset correction to } \bar{\nabla} C^T(\zeta - \zeta_i) + C[\zeta_i] = \beta} \end{aligned} \quad (21)$$

Fig.2 shows the outcome of the optimization for various β values. In what follows, we discuss what effect β has on the legibility of the trajectory in practice, as measured through users observing the robot's motion.

VII. FROM THEORY TO USERS

Legibility is intrinsically a property that depends on the observer: a real user. In this section, we test our legibility motion planner, as well as our theoretical notion of a trust region, on users observing motion. If our assumptions are true, then by varying $\beta \in [\beta_{min}, \beta_{max}]$, we expect to find that an intermediate value β^* produces the most legible result: much lower than β^* and the trajectory does not depart predictability enough to convey intent, much higher and the trajectory becomes too unpredictable, confusing the users and thus actually having a negative impact on legibility.

A. Main Experiment

Hypotheses.

H1 *The size of the trust region, β , has a significant effect on legibility.*

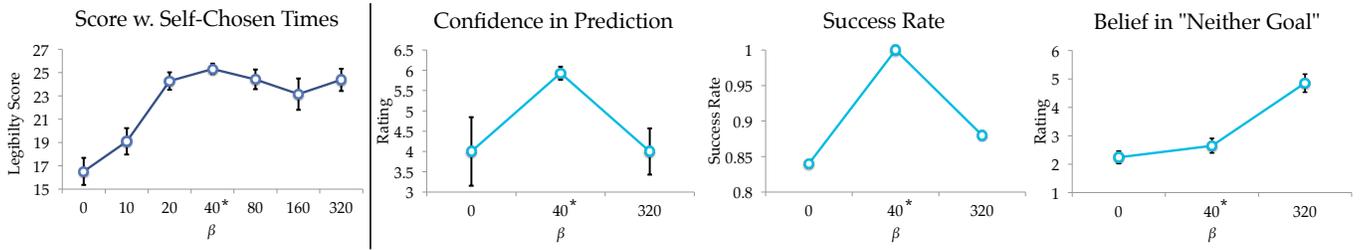


Fig. 4. Left: The legibility score for all 7 conditions in our main experiment: as the trust region grows, the trajectory becomes more legible. However, beyond a certain trust region size ($\beta = 40$), we see no added benefit of legibility. Right: In a follow-up study, we showed users the entire first half of the trajectories, and asked them to predict the goal, rate their confidence, as well as their belief that the robot is heading towards neither goal. The results reinforce the need for a trust region.

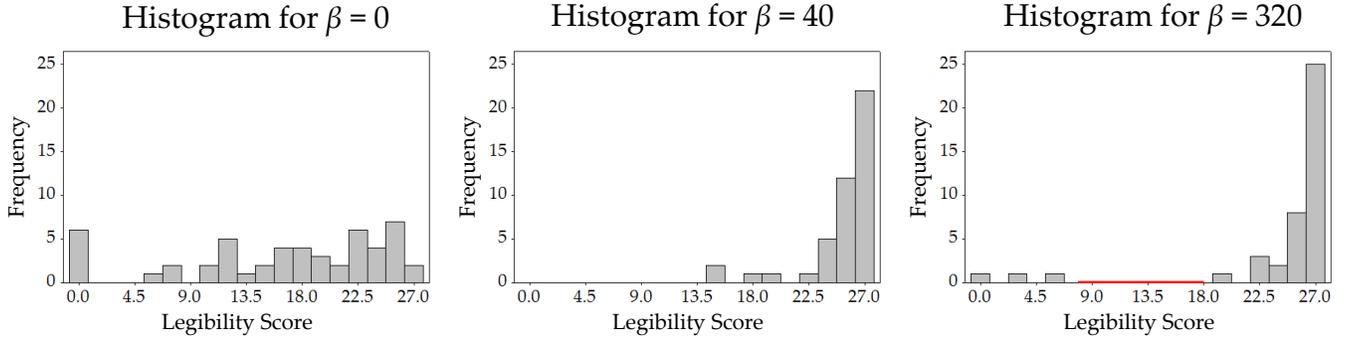


Fig. 5. The distribution of scores for three of the conditions. With a very large trust region, even though the legibility score does not significantly decrease, the users either infer the goal very quickly, or they wait until the end of the trajectory, suggesting a legibility issue with the middle portion of the trajectory.

H2 Legibility will significantly increase with β at first, but start decreasing at some large enough β .

Manipulated Variables. We manipulated β , selecting values that grow geometrically (with scalar 2) starting at 10 and ending at 320, a value we considered high enough to either support or contradict the expected effect. We also tested $\beta = \min_{\xi} C[\xi]$, which allows for no additional legibility and thus produces the predictable trajectory (we denote this as $\beta = 0$ for simplicity). We created optimal trajectories for each β in the scene from Fig.3: a point robot reaching for one of two goals.

Dependent Measures. We measured the legibility of the seven trajectories. Our measurement method follows [13, 18]: we showed the users a video of the trajectory, and asked them to stop the video as soon as they felt confident in their prediction of which goal the robot is headed toward (Fig.3). We recorded their goal prediction and the time from the start of the video to the point where they stopped it, and combined the two into a single metric based on the Guttman score [7]. Incorrect predictions received a score of 0, and correct ones received a linearly higher score when the response time was lower, i.e. when they became confident in the correct prediction earlier. We used slow videos (28s) to control for response time effects.

Subject Allocation. We chose a between-subjects design in order to not bias the users with trajectories from previous conditions. We recruited 320 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 all participants 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 297 samples.

Analysis. An ANOVA using β as a factor supported H1, showing that the factor had a significant effect on legibility ($F(6,290) = 12.57, p < 0.001$). Fig.4(left) shows the means and standard errors for each condition.

An all-pairs post-hoc analysis with Tukey corrections for multiple comparisons revealed that all trajectories with $\beta \geq 20$ were significantly more legible than the predictable trajectory ($\beta = 0$), all with $p \leq 0.001$, the maximum being reached at $\beta = 40$. This supports the first part of H2, that legibility significantly increases with β at first: *there is no practical need to become more unpredictable beyond this point.*

The maximum mean legibility was the trajectory with $\beta = 40$. Beyond this value, the mean legibility stopped increasing. Contrary to our expectation, it did not significantly decrease. In fact, the difference in score between $\beta = 40$ and $\beta = 320$ is in fact significantly less than 2.81 ($t(84) = 1.67, p = 0.05$). At a first glance, the robot’s overly unpredictable behavior seems to not have caused any confusion as to what its intent was.

Analyzing the score histograms (Fig.5) for different β values, we observed that for the high β s, users did not stop the trajectory in the middle: the guessed the goal in the beginning, or waited until the end. The consequence

is that our legibility measure failed to capture whether the mid-part of the trajectory becomes illegible. Thus, we ran a follow-up study to verify that legibility in this region does decrease at $\beta = 320$ as compared to our $\beta^* = 40$.

B. Follow-Up Study

Our follow-up study was designed to investigate legibility during the middle of the trajectories. The setup was the same, but rather than allowing the users to set the time at which they provide an answer, we fixed the time and instead asked them for a prediction and a rating of their confidence on a Likert scale from 1 to 7. We hypothesize that in this case, the users' confidence (aggregated with success rate such that a wrong prediction with high confidence is treated negatively) will align with our H2: it will be higher for $\beta = 40$ than for $\beta = 320$.

We conducted this study with 90 users. Fig.4 plots the confidences and success rates, showing that they are higher for $\beta = 40$ than they are for both of the extremes, 0 and 320. An ANOVA confirmed that the confidence effect was significant ($F(2,84) = 3.64, p = 0.03$). The post-hoc analysis confirmed that $\beta = 40$ had significantly higher confidence ($t(57) = 2.43, p = 0.45$).

We also asked the users to what extent they believed that the robot was going for neither of the goals depicted in the scene (also Fig.4). In an analogous analysis, we found that users in the $\beta = 40$ condition believed this significantly less than users in the $\beta = 320$ condition ($t(57) = 5.7, p < 0.001$).

C. Interpretation

Overall, the results support the existence of a *trust region of expectation within which legibility optimization can make trajectories significantly more legible* to novice users. Outside of this trust region, being more legible w.r.t. LEGIBILITY an impractical quest, because it no longer improves legibility in practice. Furthermore, the unpredictability of the trajectory can actually confuse the observer enough that they can no longer accurately and confidently predict the goal, and perhaps even doubt that they have the right understanding of how the robot behaves. They start believing in a "neither goal" option that is not present in the scene. *Indeed, the legibility formalism can only be trusted within this trust region.*

VIII. UNDERSTANDING LEGIBLE TRAJECTORIES

Armed with a legible motion generator, we investigate legibility further, looking at factors that affect the final trajectories.

Ambiguity. Certain scenes are more ambiguous than others, in that the legibility of the predictable trajectory is lower. The more ambiguous a scene is, the greater the need to depart from predictability and exaggerate the motion. Fig.6(a) compares two scenes, the one on the right being more ambiguous by having the candidate goals closer and thus making it more difficult to distinguish between them. This ambiguity is reflected

in its equivalent legible trajectory (both trajectories are obtained after 1000 iterations). The figure uses the same cost C from Sec. IV-B.

Scale. The scale does affect legibility when the value functions V_G are affected by scale, as in our running example. Here, reaching somewhere closer raises the demand on legibility (Fig.6(b)). Intuitively, the robot could still reach for G_O and suffer little penalty compared to a larger scale, which puts an extra burden on its motion if it wants to institute the same confidence in its intent.

Weighting in Time. The weighting function f (5) qualitatively affects the shape of the trajectory by placing the emphasis (or exaggeration) earlier or later (Fig.6(c)).

Multiple Goals. Although for simplicity, our examples so far were focused on discriminating between two goals, legibility does apply in the context of multiple goals (Fig.8(a)). Notice that for the goal in the middle, the most legible trajectory coincides with the predictable one: any exaggeration would lead an observer to predict a different goal — *legibility is limited by the complexity in the scene.*

Obstacle Avoidance. In the presence of obstacles in the scene, a user would expect the robot to stay clear of these obstacles, which makes C more complex. We plot in Fig.7 an example using the cost functional from the CHOMP motion planner[29], which trades off between the sum-squared velocity cost we have been using thus far, and a cost penalizing the robot from coming too close to obstacles. Legibility in this case will move the predictable trajectory much closer to the obstacle in order to disambiguate between the two goals.

Local optima. There is no guarantee that LEGIBILITY is concave. This is clear for the case of a non-convex C , where we often see different initializations lead to different local maxima, as in Fig.8(b).

In fact, even for quadratic V_G s, $P(G_R | \xi_S \rightarrow Q)$ is – aside from scalar variations – a ratio of sums of Gaussian functions of the form $\exp(-V_G(\xi(t)))$. Convergence to local optima is thus possible even in this simple case.

As a side-effect, it is also possible that initializing the optimizer with the most predictable trajectory leads to convergence to a local maxima.

IX. LEGIBILITY IN HIGH-DIMENSIONAL SPACES

So far, our studies and examples focused on a two-dimensional space. Our optimization method for legibility does apply to high-dimensional spaces, but comes with two big challenges that are much easier addressed in low dimensions: 1) finding the cost functional C describing user expectation; and 2) computing its value function V_G for every candidate goal $G \in \mathcal{G}$.

In the case of mobile manipulator like in Fig.9, legibility implies going beyond end effector position, to orientation, elbow location, etc. If we assume the same C as in examples so far (sum squared velocities in configuration space), then V has an analytical form, and local legibility optimization happens in real-time despite

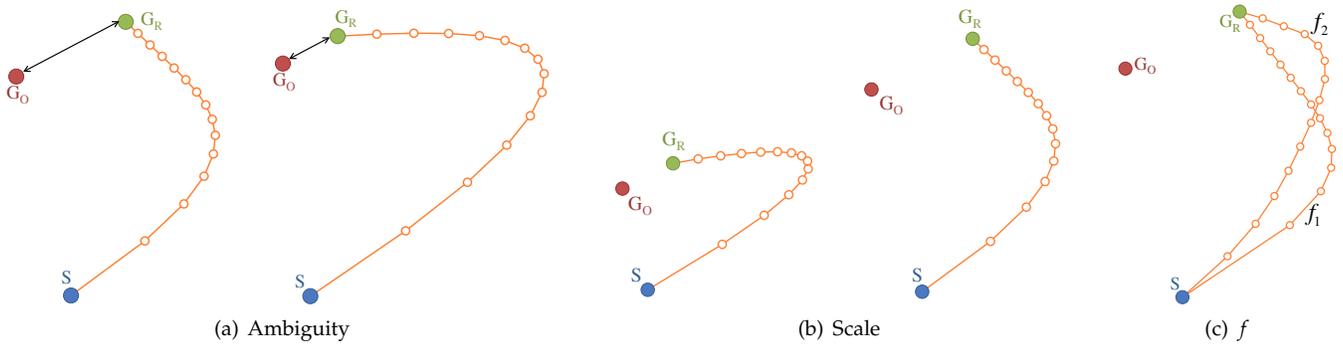


Fig. 6. The effects of ambiguity, scale, and the weighting function f on legibility.

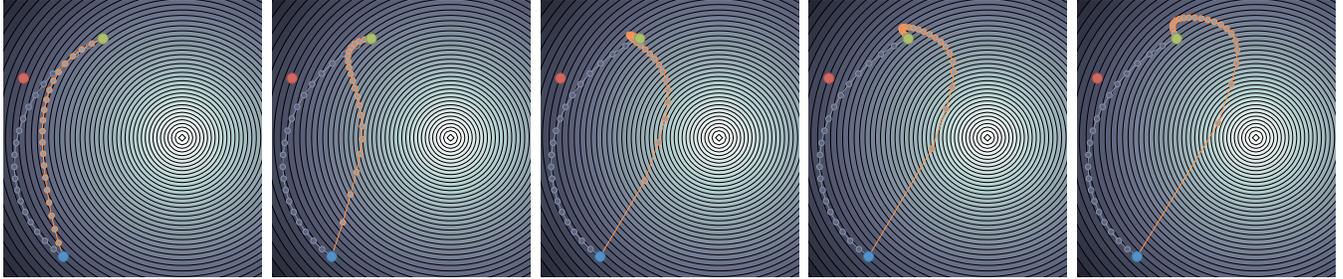


Fig. 7. Legibility given a C that accounts for obstacle avoidance. The gray trajectory is the predictable trajectory (minimizing C), and the orange trajectories are obtained via legibility optimization for 10, 10^2 , 10^3 , 10^4 , and 10^5 iterations. Legibility purposefully pushes the trajectory closer to the obstacle than expected in order to express the intent of reaching the goal on the right.

the high-dimensionality of the space. We show the result in Fig.9, in which a 7DOF arm is reaching for one of two objects. In this case, the end effector traces for the predictable and the resulting legible trajectories are similar to our 2D examples, as well as to the trajectories we used in [13] (which has shown the legible trajectory to be significantly more legible to users than its predictable counterpart).

However, this positive result should be taken with a grain of salt. Unlike for the point robot case, we do not actually know what makes a trajectory predictable in this higher-dimensional space. The fact that our C had a reasonable effect here does not mean that this is the C that users would expect, or that the result would generalize to other situations — especially for less anthropomorphic robots, for which straight lines in configuration space could be far from predictable. This leads us to the first challenge of high dimensional spaces:

Finding C . If the human observer expects *human-like* motion, cues from animation or biomechanics [16, 19, 24, 36] can help provide good approximations for C . However, our previous studies suggest that efficiency of *robot* motion has different meanings for different observers [13]. A possibility is to learn from demonstrations provided by the observer. Here, the robot can learn a C that explains the demonstrations[3], using tools like Inverse Optimal Control (IOC) [1, 28, 38]. However, extending these tools to higher dimensions is an open problem [28].

Aside from investigating the extension of IOC to high-dimensional spaces, we also propose a second thread of research: the idea of habituating users to robot behavior.

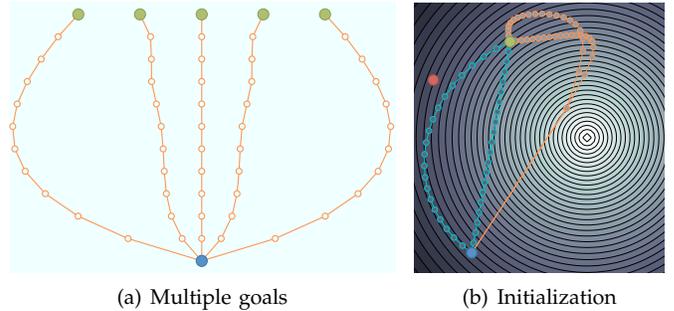


Fig. 8. (a) Legible trajectories for multiple goals. (b) Legibility is dependent on initialization.

Can users *be taught* a particular C over time?

Computing V . Given a C , legibility optimization requires access to its value function for every goal. In simple cases, like the one we focused on in this paper, V has an analytical form. But this is not the case, for instance, for non-convex functions that require obstacle avoidance. In such cases, finding good approximations for V becomes crucial, many techniques value function approximation techniques can be applied toward this goal [8].

What makes our problem special, however, is that the quality of the approximation is *defined in terms of its impact on legibility*, and not on the original value function itself. There could be approximations, such as ignoring entire components of C , or only focusing on some lower-dimensional aspects, which are very poor approximations of V itself, but might have little effect on legibility in practice.

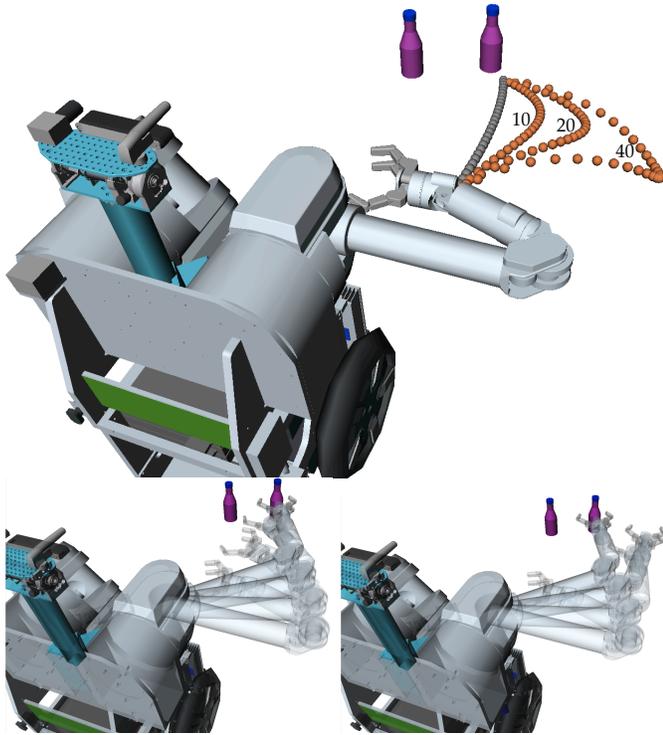


Fig. 9. Legible trajectories on a robot manipulator assuming C , computed by optimizing LEGIBILITY in the full dimensional space. The figure shows trajectories after 0 (gray), 10, 20, and 40 iterations. Below, a full-arm depiction of the trajectories at 0 and 20 iterations.

X. DISCUSSION

Limitations. Our work is limited in many ways. As the previous section discussed, in optimizing legibility, we inherit the challenges of learning and optimizing non-convex functions in high-dimensional spaces. Furthermore, adding a trust region to the optimization is a way to prevent the algorithm for traveling on “uncharted territory” — from reaching trajectories where the model’s axioms stop holding. It does not, however, fix the model itself, as it does not capture the inferences that observers would make in those regions.

Implications. Legibility will play a crucial role in enabling robots to seamlessly collaborate with humans. In this paper, we proposed a method that can generate legible motion, and illustrated a path of future work for addressing the remaining challenges. In addition, we are excited to explore applications of legibility beyond robotics, for example in animation, as well as applications of our method beyond legible motion, to purposefully ambiguous or deceptive motion.

ACKNOWLEDGMENTS

This material is based upon work supported by NSF-IIS-0916557, NSF-EEC-0540865, ONR-YIP 2012, and the Intel Embedded Computing ISTC. We thank the members of the Personal Robotics Lab for very helpful discussion and advice.

REFERENCES

[1] P. Abbeel and A. Y. Ng. Apprenticeship learning via inverse reinforcement learning. In *ICML*, 2004.
 [2] R. Alami, A. Albu-Schaeffer, A. Bicchi, R. Bischoff, R. Chatila, A. D. Luca, A. D. Santis, G. Giralt, J. Guiochet, G. Hirzinger, F. Ingrand, V. Lippiello,

R. Mattone, D. Powell, S. Sen, B. Siciliano, G. Tonietti, and L. Villani. Safe and Dependable Physical Human-Robot Interaction in Anthropic Domains: State of the Art and Challenges. In *IROS Workshop on pHRI*, 2006.
 [3] B. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *RAS*, 57(5):469 – 483, 2009.
 [4] D. A. Baldwin, J. A. Baird, M. M. Saylor, and M. A. Clark. Infants parse dynamic action. *Child Development*, 72(3):708–717, 2001.
 [5] M. Beetz, F. Stulp, P. Esden-Tempski, A. Fedrizzi, U. Klank, I. Kresse, A. Maldonado, and F. Ruiz. Generality and legibility in mobile manipulation. *Autonomous Robots*, 28:21–44, 2010.
 [6] T. Behne, M. Carpenter, J. Call, and M. Tomasello. Unwilling Versus Unable: Infants’ Understanding of Intentional Action. *Developmental Psychology*, 41:328–337, 2005.
 [7] G. Bergersen, J. Hannay, D. Sjoberg, T. Dyba, and A. Karahasanovic. Inferring skill from tests of programming performance: Combining time and quality. In *ESEM*, 2011.
 [8] J. Boyan and A. Moore. Generalization in reinforcement learning: Safely approximating the value function. *NIPS*, 1995.
 [9] O. Brock and O. Khatib. Elastic strips: A framework for motion generation in human environments. *IJRR*, 21(12):1031, 2002.
 [10] M. Carpenter, N. K., Tomasello, G. M., Butterworth, and C. Moore. Social cognition, joint attention, and communicative competence from 9 to 15 months of age. *Monographs of the Society for Research in Child Development*, 63(4):1–174.
 [11] E. J. Carter, J. K. Hodgins, and D. H. Rakison. Exploring the neural correlates of goal-directed action and intention understanding. *NeuroImage*, 54(2):1634–1642, 2011.
 [12] G. Csibra and G. Gergely. Obsessed with goals: Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, 124(1):60 – 78, 2007.
 [13] A. Dragan, K. Lee, and S. Srinivasa. Legibility and predictability of robot motion. In *ACM/IEEE HRI*, 2013.
 [14] A. Dragan, N. Ratliff, and S. Srinivasa. Manipulation planning with goal sets using constrained trajectory optimization. In *ICRA*, May 2011.
 [15] A. Dragan and S. Srinivasa. Formalizing assistive teleoperation. In *RSS*, July 2012.
 [16] T. Flash and N. Hogan. The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci*, 5:1688–1703, July 1985.
 [17] G. Gergely, Z. Nadasdy, G. Csibra, and S. Biro. Taking the intentional stance at 12 months of age. *Cognition*, 56(2):165 – 193, 1995.
 [18] M. Gielniak and A. Thomaz. Generating anticipation in robot motion. In *RO-MAN*, 2011.
 [19] M. Gielniak and A. L. Thomaz. Spatiotemporal correspondence as a metric for human-like robot motion. In *ACM/IEEE HRI*, 2011.
 [20] J. Heinzmann and A. Zelinsky. The safe control of human-friendly robots. In *IEEE/RSJ IROS*, 1999.
 [21] C. Igel, M. Toussaint, and W. Weishui. Rprop using the natural gradient. *Trends and Applications in Constructive Approximation*, pages 259–272, 2005.
 [22] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal. STOMP: Stochastic trajectory optimization for motion planning. In *IEEE ICRA*, 2011.
 [23] 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*, nov.-dec. 2004.
 [24] J. Lasseter. Principles of traditional animation applied to 3d computer animation. In *SIGGRAPH*, 1987.
 [25] A. N. Meltzoff. Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31(5):838–850, 1995.
 [26] S. Nikolaidis and J. Shah. Human-robot teaming using shared mental models. In *ACM/IEEE HRI*, 2012.
 [27] S. Quinlan. *The Real-Time Modification of Collision-Free Paths*. PhD thesis, Stanford University, 1994.
 [28] N. Ratliff, J. A. Bagnell, and M. Zinkevich. Maximum margin planning. In *ICML*, 2006.
 [29] N. Ratliff, M. Zucker, J. A. D. Bagnell, and S. Srinivasa. Chomp: Gradient optimization techniques for efficient motion planning. In *ICRA*, May 2009.
 [30] E. Short, J. Hart, M. Vu, and B. Scassellati. No fair!! an interaction with a cheating robot. In *ACM/IEEE HRI*, 2010.
 [31] B. Sodani and C. Thoenner. Infants’ understanding of looking, pointing, and reaching as cues to goal-directed action. *Journal of Cognition and Development*, 5(3):289–316, 2004.
 [32] L. Takayama, D. Dooley, and W. Ju. Expressing thought: improving robot readability with animation principles. In *HRI*, 2011.
 [33] E. Todorov and W. Li. A generalized iterative lqg method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *ACC*, 2005.
 [34] M. Tomasello, M. Carpenter, J. Call, T. Behne, and H. Moll. Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and Brain Sciences*, 2004.
 [35] M. Toussaint. Robot trajectory optimization using approximate inference. In *International Conference on Machine Learning*, 2009.
 [36] A. Witkin and M. Kass. Spacetime constraints. In *SIGGRAPH*, 1988.
 [37] A. L. Woodward. Infants selectively encode the goal object of an actor’s reach. *Cognition*, 69(1):1 – 34, 1998.
 [38] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, 2008.