

Integrating Human Observer Inferences into Robot Motion Planning

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

Abstract—Our goal is to enable robots to produce motion that is suitable for human-robot collaboration and co-existence. Most motion in robotics is purely *functional*, ideal when the robot is performing a task in isolation. In collaboration, however, the robot’s motion has an *observer*, watching and interpreting the motion.

In this work, we move beyond functional motion, and introduce the notion of an *observer* into motion planning, so that robots can generate motion that is mindful of how it will be interpreted by a human collaborator.

We formalize *predictability* and *legibility* as properties of motion that naturally arise from the inferences in opposing directions that the observer makes, drawing on action interpretation theory in psychology. We propose models for these inferences based on the principle of rational action, and derive constrained functional trajectory optimization techniques for planning motion that is predictable or legible.

Finally, we present experiments that test our work on novice users, and discuss the remaining challenges in enabling robots to generate such motion online in complex situations.¹

I. INTRODUCTION

In this paper, we explore the problem where a robot and a human are collaborating side by side to perform a tightly coupled physical task together, like clearing a table (a running example in our paper).

The task amplifies the burden on the robot’s motion. Most motion in robotics is purely *functional*: industrial robots move to package parts, vacuuming robots move to suck dust, and personal robots move to clean up a dirty table. This type of motion is ideal when the robot is performing a task in isolation.

Collaboration, however, does not happen in isolation. In collaboration, the robot’s motion has an *observer*, watching and interpreting the motion.

In this work, we move beyond functional motion, and introduce the notion of an *observer* and their inferences into motion planning, so that robots can generate motion that is mindful of how it will be interpreted by a human collaborator.

When we collaborate, we make two inferences (Fig.1, lower center) about our collaborator [12, 54, 56]: 1) we infer their goal based on their ongoing action (action-to-goal), and 2) if we know their goal, we infer their future action from it (goal-to-action). Our work




Fig. 1. Above: Predictable, day-to-day, expected handwriting vs. legible handwriting. Upper Center: A predictable and a legible trajectory of a robot’s hand for the same task of grasping the green object. Lower Center: Predictability and legibility stem from inferences in opposing directions. Below: The legibility optimization process for a reaching task. By exaggerating to the right, the robot is more clear about its intent to grasp the object on the right .

formalizes the two properties of motion that *enable* these two inferences: *predictability* and *legibility*.

Legibility stems from the first inference, and is about conveying intent – moving in a manner that makes the robot’s goal clear to observer. Predictability stems from the second inference, and is about matching the observer’s expectation – matching the motion they predict when they know the robot’s goal.

Predictable and legible motion can be correlated. For example, in an unambiguous situation, where an actor’s observed motion matches what is expected for a given intent (i.e. is predictable), then this intent can be used to explain the motion. If this is the only intent which explains the motion, the observer can immediately infer the actor’s intent, meaning that the motion is also legible.

¹This paper combines work from [15] and [18].

This is why we tend to assume that predictability implies legibility – that if the robot moves in an expected way, then its intentions will automatically be clear [2, 6, 35].

The writing domain, however, clearly distinguishes the two. The word *legibility*, traditionally an attribute of written text [52], refers to the quality of being easy to read. When we write legibly, we try consciously, and with some effort, to make our writing clear and readable to someone else. The word *predictability*, on the other hand, refers to the quality of matching expectation. When we write predictably, we fall back to old habits, and write with minimal effort.

As a consequence, our legible and predictable writings are *different*: our friends do not expect to open our diary and see our legible writing style. They rightfully assume the diary will be written for us, and expect our usual, day-to-day style. In this paper, by formalizing predictability and legibility as directly stemming from the two inferences in *opposing* directions, goal-to-action and action-to-goal, we show that the two are different in motion as well:

Predictability and legibility are fundamentally different properties of motion, stemming from observer inferences in opposing directions.

Ambiguous situations, occurring often in daily tasks, make this opposition clear: more than one possible intent can be used to explain the motion observed so far, rendering the predictable motion illegible. Fig.1(center) exemplifies the effect of this contradiction. The robot hand’s motion on the left is predictable in that it matches expected behavior. The hand reaches out directly towards the target. But, it is not legible, failing to make the intent of grasping the green object clear. In contrast, the trajectory on the right is more legible, making it clear that the target is the green object by deliberately bending away from the red object. But it is less predictable, as it does not match the expected behavior of reaching directly. We will show in Sections IV and V how we can quantify this effect with Bayesian inference, which enables us to derive the online probabilities of the motion reaching for either object, illustrated as bar graphs in Fig.1.

Our work makes the following contributions:

1. We formalize legibility and predictability in the context of goal-directed motion in Sec. III as stemming from inferences in *opposing* directions. The formalism emphasizes their difference, and directly relates to the theory of action interpretation [12] and the concepts of “action-to-goal” and “goal-to-action” inference.
2. Armed with mathematical definitions of legibility and predictability, we propose a way in which a robot could model these inferences in order to evaluate how legible or predictable a motion is (Sections IV and V). The models are based on cost optimization, resonate with

the principle of rational action [11, 23], and echo earlier works on action understanding via inverse planning [5].

3. We derive methods for generating predictable and legible motion. Although our model enables us to *evaluate* how predictable or legible a motion trajectory is, it does not enable us to *generate* trajectories that are predictable or legible. 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”. We do so via functional gradient optimization in the space of motion trajectories, echoing earlier works in motion planning [9, 28, 31, 43, 45, 53, 55, 61], now with optimization criteria based on the observer’s inferences.

4. We propose a method for optimizing for legibility while still maintaining predictability. The ability to optimize the legibility criterion leads us to a surprising observation: that there are cases in which the trajectory becomes too *unpredictable*. As our user studies show (Sec. X-B), *some unpredictability is often necessary to convey intent* — it is unpredictability beyond a threshold (like the outermost trajectory in Fig.4) that confuses users and lowers their confidence in what the robot is doing. 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.

5. We test our work in two experiments with novice users, one focused on the models, and the other on motion generation. We demonstrate legibility and predictability are contradictory not just in theory, but also in practice, and follow-up with an analysis of the trade-off between the two. We experiment with three characters that differ in their complexity and anthropomorphism: a simulated point robot, the bi-manual mobile manipulator HERB[48], and a human (Sec. X).

Understanding and planning motion that is predictable or legible can greatly enhance human-robot collaboration, but many challenges remain in generating such motion online in complex situations, and in deciding how to trade off between the two properties in each situation. We discuss these in Sec. XI.

II. RELATION TO PRIOR WORK

“Predictable” means “expected”, and is usually used to refer to a desirable property of motion without needing additional clarification [2, 6]. “Legible”, on the other hand, is typically reserved for writing, and appears in the robot motion literature accompanied by an explanation: being legible means that observers are able to recognize [6], infer [39], or understand [3, 35] the intentions of the robot, or that robot indicates the goal it will

reach [35] – overall, legibility is about expressing intent. Building on this, we define predictability and legibility as enabling inferences that the collaborator needs to make, and propose trajectory optimization criteria for the two.

One exception to using legibility to mean expressing intent is a definition of legibility as both being intent-expressive *and* matching expectation [40]. Our work, however, shows that matching expectation and expressing intent are fundamentally different properties that can at times contradict.

Existing methods for generating legible motion fall in three categories. (1) increasing legibility indirectly by increasing predictability (e.g. via learning from demonstration) [6]; (2) increasing legibility indirectly by increasing visibility to the observer [3, 29]; and (3) increasing legibility directly by encoding animation principles like anticipation in the motion [24, 49].

In contrast, our work explicitly formalizes intent-expressiveness for goal-directed motion in the form of a trajectory optimization criterion: rather than targeting a related property, the robot directly maximizes the probability of the correct intent being inferred. The resulting motion can be interpreted using animation principles, but we do not encode these explicitly – instead, they emerge out of the mathematics of legible motion.

III. FORMALIZING LEGIBILITY AND PREDICTABILITY

In common use, legible motion is intent-expressive, and predictable motion matches what is expected. Here, we formalize these definitions for the context of goal-directed motion, where a human or robot is executing a trajectory $\zeta : [0, 1] \rightarrow \mathcal{Q}$, lying in a Hilbert space of trajectories Ξ . ζ starts at a configuration S and ends at a goal G from a set of possible goals \mathcal{G} , like in Fig.1. In this context, G is central to both properties:

Definition 1: Legible motion is motion that enables an observer to quickly and confidently infer the correct goal G .

Definition 2: Predictable motion is motion that matches what an observer would expect, given the goal G .

A. Formalism

1) *Legibility*: Imagine an observer watching the orange trajectory from Fig.1. As the robot’s hand departs the starting configuration and moves along the trajectory, the observer is running an inference, predicting which of the two goals it is reaching for. We denote this inference function that maps (snippets of) trajectories from all trajectories Ξ to goals as

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

The bar graphs next to the hands in Fig.1 signify the observer’s predictions of the two likely goals. At the very beginning, the trajectory is confusing and the observer

has little *confidence* in the inference — in what follows, we model this confidence based on the probability that the observer assigns to the inferred goal. However, the observer becomes confident very *quickly* – even from the second configuration of the hand along the trajectory, it becomes clear that the green object is the target. This quick and confident inference is the hallmark of legibility.

We thus formalize *legible* motion as motion that enables an observer to *confidently* infer the *correct* goal configuration G after observing only a snippet of the trajectory, $\zeta_{S \rightarrow Q}$, from the start S to the configuration at a time t , $Q = \zeta(t)$:

$$\mathcal{I}_L(\zeta_{S \rightarrow Q}) = G$$

The *quicker* this happens (i.e. the smaller t is), the more legible the trajectory is.

The definition in terms of \mathcal{I}_L is an interpretation in the motion domain for terms like “readable” [49], or “understandable” [3], and encourages “anticipatory” motion [24] because it brings the relevant information for goal prediction towards the beginning of the trajectory, thus lowering t . The formalism can also generalize to outcome-directed motion (e.g. gestures such as pointing at, waving at, etc.) by replacing the notion of goal with that of an outcome – here, legible motion becomes motion that enables quick and confident inference of the desired outcome. Our recent work illustrated this for pointing gestures [20].

2) *Predictability*: Now imagine someone knowing that the hand is reaching towards the green goal. Even before the robot has moved, the observer creates an expectation, making an inference on how the hand will move – for example, that the hand will start turning towards the green object as it is moving directly towards it. We denote this inference function mapping goals to trajectories as

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

We formalize *predictable* motion as motion for which the trajectory $\zeta_{S \rightarrow G}$ matches this inference:

$$\mathcal{I}_P(G) = \zeta_{S \rightarrow G}$$

The better the actual trajectory matches the inference, measurable for example using a distance metric between $\mathcal{I}_P(G)$ and $\zeta_{S \rightarrow G}$, the more predictable the trajectory is.

B. Connection to Action Interpretation in Psychology

A growing amount of research in psychology suggests that humans interpret observed behaviors as goal-directed actions [10, 12, 26, 42, 47, 58, 59], a result supported by studies observing infants and how they show surprise when exposed to inexplicable action-goal pairings. Csibra and Gergely [12] summarize two types

Human Inference Type	Example	Analogy in Motion	Property of Motion
action \mapsto goal	... pour beans in grinder \mapsto coffee	$\xi_{S \rightarrow Q} \mapsto G$	<i>legibility</i>
goal \mapsto action	coffee \mapsto ... pour beans in grinder ...	$G \mapsto \xi_{S \rightarrow G}$	<i>predictability</i>

TABLE I

LEGIBILITY AND PREDICTABILITY AS ENABLING INFERENCES FROM ACTION INTERPRETATION IN OPPOSING DIRECTION.




Fig. 2. Our models for \mathcal{I}_P and \mathcal{I}_L : the observer expects the robot’s motion to optimize a cost function C (\mathcal{I}_P , left), and identifies based on C which goal is most probable given the robot’s motion so far (\mathcal{I}_L , right).

of inference stemming from the interpretation of actions as goal directed: “action-to-goal” and “goal-to-action”.

“Action-to-goal” refers to an observer’s ability to infer someone’s goal state from their ongoing actions (e.g. because they are pouring coffee beans into the grinder, they will eventually hold a cup of coffee). “Action-to-goal” inference answers the question “What is the function of this action?”.

“Goal-to-action” refers to an observer’s ability to predict the actions that someone will take based on their goal (e.g. because they want to make coffee, they will pour coffee beans into the grinder). “Goal-to-action” inference answers the question “What action would achieve this goal?”.

This has a natural connection to our formalism. In goal-directed motion, actions are trajectories and goals are goal configurations. Thus the inference occurring in legibility, from trajectory to goal, $\xi_{S \rightarrow Q} \mapsto G$, relates naturally to “action-to-goal” inference. Likewise, the inference occurring in predictability, from goal to trajectory, $G \mapsto \xi_{S \rightarrow G}$, relates naturally to “goal-to-action”.

C. Summary

Our formalism emphasizes the difference between legibility and predictability in theory: they stem from inferences in *opposing directions* (from trajectories to goals vs. from goals to trajectories), with strong parallels in the theory of action interpretation. Table I shows a summary.

In what follows, we introduce one way for a robot to model these two inferences (summarized in Fig.2), derive motion generating algorithms based on this model, and present experiments that emphasize the difference between the two properties in practice.

IV. PREDICTABLE MOTION

A. Modeling the Trajectory Inference \mathcal{I}_P

To model \mathcal{I}_P is to model the observer’s expectation. One way the robot could do so is by assuming that the human observer expects it to be a rational agent acting efficiently[12] or justifiably[11] to achieve a goal. This is known as the principle of rational action[11, 23], and it has been shown to apply to non-human agents, including robots[33]. The robot can model this notion of “efficiency” via a cost function defining what it means to be efficient. For example, if the observer expected the robot’s hand to move directly towards the object it wants to grasp (as opposed to taking an unnecessarily long path to it), then “efficiency” would be defined by the cost function penalizing the trajectory’s length.

Throughout this paper, we will refer to the cost function modeling the observer’s expectation as C :

$$C : \Xi \rightarrow \mathbb{R}^+$$

with lower costs signifying more “efficient” trajectories.

The principle of rational action suggests that the most predictable trajectory is the most “efficient”, for some definition of efficiency C :

$$\mathcal{I}_P(G) = \arg \min_{\xi \in \Xi_{S \rightarrow G}} C[\xi] \quad (1)$$

C represents what the observer expects the robot to optimize, and therefore encompasses every aspect of the observer’s expectation, including (when available) body motion, hand motion, arm motion, and gaze.

B. Predictability Score

Predictability can be evaluated based on C : the lower the cost, the more predictable (expected) the trajectory. We propose a predictability *score* normalized from 0 to 1, where trajectories with lower cost are exponentially

more predictable (following the principle of maximum entropy, which we detail in Sec. V-A):

$$\text{PREDICTABILITY}(\xi) = \exp(-C[\xi]) \quad (2)$$

C. Generating Predictable Motion

Generating predictable motion means maximizing the predictability score, or equivalently minimizing the cost function C – as in (1).

One way to do so is via functional gradient descent. We start from an initial trajectory ξ_0 and iteratively improve its score. At every iteration i , we maximize the regularized first order Taylor series approximation of C about the current trajectory ξ_i :

$$\xi_{i+1} = \arg \min_{\xi \in \Xi} C[\xi_i] + \nabla C_{\xi_i}^T (\xi - \xi_i) + \frac{\eta}{2} \|\xi - \xi_i\|_M^2 \quad (3)$$

with $\frac{\eta}{2} \|\xi - \xi_i\|_M^2$ a regularizer restricting the norm of the displacement $\xi - \xi_i$ w.r.t. an M , as in [45].

By taking the functional gradient of (12) and setting it to 0, we obtain the following update rule for ξ_{i+1} :

$$\xi_{i+1} = \xi_i - \frac{1}{\eta} M^{-1} \bar{\nabla} C \quad (4)$$

Recent work has successfully applied this method to motion planning, using a C that trades off between an efficiency and an obstacle avoidance component [61]. The non-convexity of this type of C makes trajectory optimization prone to convergence to high-cost local minima, which can be mediated by learning from experience [13, 14].

A key challenge, which we discuss in Sec. XI, is finding the C that each observer expects the robot to minimize.

V. LEGIBLE MOTION

A. Modeling the Goal Inference \mathcal{I}_L

To model \mathcal{I}_L is to model how the observer infers the goal from a snippet of the trajectory $\xi_{S \rightarrow Q}$. One way to do so is by assuming that the observer compares the possible goals in the scene in terms of how probable each is given $\xi_{S \rightarrow Q}$. This is supported by action interpretation: Csibra and Gergeley [12] argue, based on the principle of rational action, that humans assess which end state would be most efficiently brought about by the observed ongoing action. Taking trajectory length again as an example for the observer's expectation, this translates to predicting a goal because $\xi_{S \rightarrow Q}$ moves directly toward it and away from the other goals, making them less probable.

One model for \mathcal{I}_L is to compute the probability for each goal candidate G and to choose the most likely:

$$\mathcal{I}_L(\xi_{S \rightarrow Q}) = \arg \max_{G \in \mathcal{G}} P(G | \xi_{S \rightarrow Q}) \quad (5)$$

To compute this probability, we start with Bayes' Rule:

$$P(G | \xi_{S \rightarrow Q}) \propto P(\xi_{S \rightarrow Q} | G) P(G) \quad (6)$$

where $P(G)$ is a prior on the goals which can be uniform in the absence of prior knowledge, and $P(\xi_{S \rightarrow Q} | G)$ is the probability of seeing $\xi_{S \rightarrow Q}$ when the robot targets goal G . The is in line with the notion of action understanding as inverse planning proposed by Baker et al.[5], here $P(\xi_{S \rightarrow Q} | G)$ relating to the forward planning problem of finding a trajectory given a goal.

We compute $P(\xi_{S \rightarrow Q} | G)$ as the ratio of all trajectories from S to G that pass through $\xi_{S \rightarrow Q}$ to all trajectories from S to G (Fig.3):

$$P(\xi_{S \rightarrow Q} | G) = \frac{\int_{\xi_{Q \rightarrow G}} P(\xi_{S \rightarrow Q \rightarrow G})}{\int_{\xi_{S \rightarrow G}} P(\xi_{S \rightarrow G})} \quad (7)$$

Following [60], we assume trajectories are separable, i.e. $P(\xi_{X \rightarrow Y \rightarrow Z}) = P(\xi_{X \rightarrow Y}) P(\xi_{Y \rightarrow Z})$, giving us:

$$P(\xi_{S \rightarrow Q} | G) = \frac{P(\xi_{S \rightarrow Q}) \int_{\xi_{Q \rightarrow G}} P(\xi_{Q \rightarrow G})}{\int_{\xi_{S \rightarrow G}} P(\xi_{S \rightarrow G})} \quad (8)$$

At this point, the robot needs a model of how probable a trajectory ξ is in the eye of an observer. The observer expects the trajectory of minimum cost under C . It is unlikely, however, that they would be completely surprised (i.e. assign 0 probability) by all other trajectories, especially by one ever so slightly different. One way to model this is to make suboptimality w.r.t. C still possible, but exponentially less probable, i.e. $P(\xi) \propto \exp(-C[\xi])$, adopting the principle of maximum entropy [60]. With this, (8) becomes:

$$P(\xi_{S \rightarrow Q} | G) \propto \frac{\exp(-C[\xi_{S \rightarrow Q}]) \int_{\xi_{Q \rightarrow G}} \exp(-C[\xi_{Q \rightarrow G}])}{\int_{\xi_{S \rightarrow G}} \exp(-C[\xi_{S \rightarrow G}])} \quad (9)$$

Computing the integrals is still challenging. In [19], we derived a solution by approximating the probabilities using Laplace's method (also proposed independently in [38]). If we approximate C as a quadratic, its Hessian is constant and according to Lapace's method, $\int_{\xi_{X \rightarrow Y}} \exp(-C[\xi_{X \rightarrow Y}]) \approx k \exp(-C[\xi_{X \rightarrow Y}^*])$ (with k a constant and $\xi_{X \rightarrow Y}^*$ the optimal trajectory from X to Y w.r.t. C). Plugging this into (9) and using (6) we get:

$$P(G | \xi_{S \rightarrow Q}) = \frac{1}{Z} \frac{\exp(-C[\xi_{S \rightarrow Q}] - V_G(Q))}{\exp(-V_G(S))} P(G_R) \quad (10)$$

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

Much like teleological reasoning suggests [12], this evaluates how efficient (w.r.t. C) going to a goal is through the observed trajectory snippet $\xi_{S \rightarrow Q}$ relative to the most efficient (optimal) trajectory, $\xi_{S \rightarrow G}^*$. In ambiguous situations like the one in Fig.1, a large portion




Fig. 3. $\zeta_{S \rightarrow Q}$ in black, examples of $\zeta_{Q \rightarrow G}$ in blue, and further examples of $\zeta_{S \rightarrow G}$ in purple. Trajectories more costly w.r.t. C are less probable.

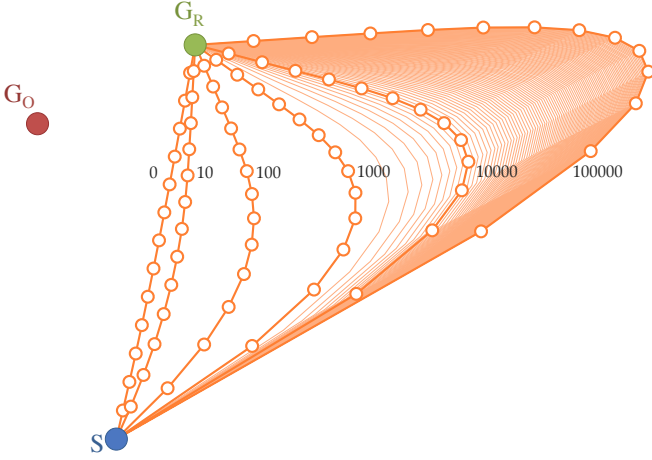


Fig. 4. The trajectory optimization process for legibility: the figure shows the trajectories across iterations for a situation with a start and two candidate goals.

of $\zeta_{S \rightarrow G}^*$ is also optimal (or near-optimal) for a different goal, making both goals almost equally likely along it. *This is why legibility does not also optimize C — rather than matching expectation, it manipulates it to convey intent.*

B. Legibility Score

A legible trajectory is one that enables quick and confident predictions. A score for legibility therefore tracks the probability assigned to the actual goal G^* 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 duration of the trajectory):

$$\text{LEGIBILITY}(\zeta) = \frac{\int P(G^* | \zeta_{S \rightarrow \zeta(t)}) f(t) dt}{\int f(t) dt} \quad (11)$$

with $P(G^* | \zeta_{S \rightarrow \zeta(t)})$ computed using C , as in (10).

C. Generating Legible Motion

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.4). This process is analogous to generating predictable motion, now with a different optimization criterion:

$$\begin{aligned} \zeta_{i+1} = \arg \max_{\zeta} & \text{LEGIBILITY}[\zeta_i] + \langle \bar{\nabla} \text{LEGIBILITY}, (\zeta - \zeta_i) \rangle \\ & - \frac{\eta}{2} \|\zeta - \zeta_i\|_M^2 \end{aligned} \quad (12)$$

By taking the functional gradient of (12) and setting it to 0, we obtain the following update rule for ζ_{i+1} :

$$\zeta_{i+1} = \zeta_i + \frac{1}{\eta} M^{-1} \bar{\nabla} \text{LEGIBILITY} \quad (13)$$

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

$$\text{LEGIBILITY}[\zeta] = K \int \mathcal{P}(\zeta(t), t) dt \quad (14)$$

and

$$\bar{\nabla} \text{LEGIBILITY} = K \left(\frac{\partial \mathcal{P}}{\partial \zeta} - \frac{d}{dt} \frac{\partial \mathcal{P}}{\partial \zeta'} \right) \quad (15)$$

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

$$\frac{\delta \mathcal{P}}{\delta \zeta}(\zeta(t), t) = \frac{g'h - h'g}{h^2} P(G_R) f(t) \quad (16)$$

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

$$\begin{aligned} \frac{\partial \mathcal{P}}{\partial \zeta}(\zeta(t), t) &= \frac{\exp(V_{G_R}(S) - V_{G_R}(\zeta(t)))}{(\sum_G \exp(V_G(S) - V_G(\zeta(t))))^2} \\ &\sum_G \left(\frac{\exp(-V_G(\zeta(t)))}{\exp(-V_G(S))} (V'_G(\zeta(t)) - V'_{G_R}(\zeta(t))) \right) P(G_R) f(t) \end{aligned} \quad (17)$$

Finally,

$$\bar{\nabla} \text{LEGIBILITY}(t) = K \frac{\partial \mathcal{P}}{\partial \zeta}(\zeta(t), t) \quad (18)$$

with $\frac{\partial \mathcal{P}}{\partial \zeta}(\zeta(t), t)$ from (17).

VI. EXAMPLE

Fig.4 portrays the functional optimization process for legibility for a point robot moving from a start location to one of two candidate goals.

A. Parameters

We detail our choice of parameters for creating this example below.

o **Efficiency cost C.** In this example, we use sum squared velocities as the cost functional C capturing the user’s expectation:

$$C[\xi] = \frac{1}{2} \int \xi'(t)^2 dt \quad (19)$$

This cost, frequently used to encourage trajectory smoothness [45], produces trajectories that reach directly toward the goal — this is the predictable trajectory, shown at iteration 0. Our experiments below show that this trajectory geometry is accurate for 2D robots like in Fig.4. This C also allows for an analytical V_G and its gradient.

o **Trajectory initialization.** We set $\xi_0 = \arg \min_{\xi} C[\xi]$: we initialize with the most predictable trajectory, treating C as a prior.

o **Trajectory parametrization.** We parametrize the trajectory as a vector of waypoint configurations.

o **Norms w.r.t. M .** We use the Hessian of C for M . As a result, the update rule in (4) and (13) propagates local gradient changes linearly to the rest of the trajectory.

B. Interpretation

The predictable trajectory (iteration 0) is efficient for achieving the goal, but it is not the best at helping an observer distinguish which goal the robot is targeting from the beginning. By exaggerating the motion to the right, the robot is increasing the probability of the goal on the right relative to the one on the left.

Exaggeration is one of the 12 principles of animation [51]. However, nowhere did we inform the robot of what exaggeration is and how it might be useful for legibility. The behavior elegantly emerged out of the optimization.

VII. THE UNPREDICTABILITY OF LEGIBILITY

The example leads to a surprising observation: in some cases, *by optimizing the LEGIBILITY functional, one can become arbitrarily unpredictable.*

Proof: Our gradient derivation in (17) enables us to construct cases in which this occurs. In a two-goal case like in Fig.4, with our example C from (19), the gradient for each trajectory configuration points in the direction $G_R - G_O$ and has positive magnitude everywhere but at ∞ , where $C[\xi] = \infty$. Fig.5 (red) plots C across iterations.

The reason for this peculiarity is that the model for how observers make inferences in (5) *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

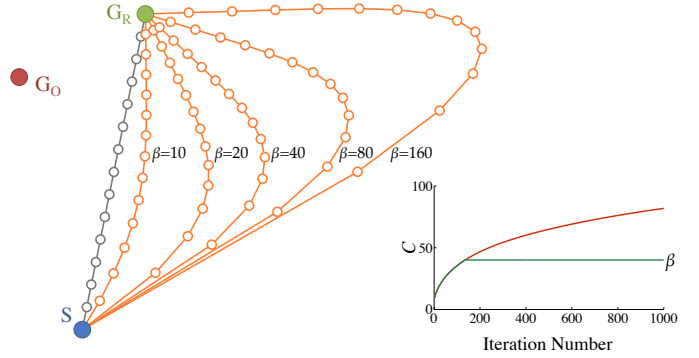


Fig. 5. The predictable trajectory in gray, and the legible trajectories for different trust region sizes in orange. On the right, the cost C over the iterations in the unconstrained case (red) and constrained case (green).

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 [46] or that it is not pursuing any of the goals and doing something else entirely — this is supported by our user study in Sec. X-B, which shows 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. X-B). Unpredictability above a certain level can also be detrimental to the collaboration process in general [2, 27, 41].

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 $C[\xi] \leq \beta$:

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.

VIII. 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 optimiza-

tion in (12):

$$\begin{aligned} \zeta_{i+1} &= \arg \max_{\zeta} \text{LEGIBILITY}[\zeta_i] + \bar{\nabla} \text{LEGIBILITY}^T(\zeta - \zeta_i) \\ &\quad - \frac{\eta}{2} \|\zeta - \zeta_i\|_M^2 \\ \text{s.t. } &C[\zeta] \leq \beta \end{aligned} \quad (20)$$

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

$$\begin{aligned} \mathcal{L}[\zeta, \lambda] &= \text{LEGIBILITY}[\zeta_i] + \bar{\nabla} \text{LEGIBILITY}^T(\zeta - \zeta_i) \\ &\quad - \frac{\eta}{2} \|\zeta - \zeta_i\|_M^2 + \lambda(\beta - \bar{\nabla} C^T(\zeta - \zeta_i) - C[\zeta_i]) \end{aligned} \quad (21)$$

with the following KKT conditions:

$$\bar{\nabla} \text{LEGIBILITY} - \eta M(\zeta - \zeta_i) - \bar{\nabla} C \lambda = 0 \quad (22)$$

$$\lambda(\beta - \bar{\nabla} C^T(\zeta - \zeta_i) - C[\zeta_i]) = 0 \quad (23)$$

$$\lambda \geq 0 \quad (24)$$

$$C[\zeta] \leq \beta \quad (25)$$

Inactive constraint: $\lambda = 0$ and

$$\zeta_{i+1} = \zeta_i + \frac{1}{\eta} M^{-1} \bar{\nabla} \text{LEGIBILITY} \quad (26)$$

Active constraint: The constraint becomes an equality constraint on the trajectory. The derivation for ζ_{i+1} is analogous to [17], using the LEGIBILITY functional as opposed to the classical cost used by the CHOMP motion planner[45]. From (22)

$$\zeta_{i+1} = \zeta_i + \frac{1}{\eta} M^{-1} \underbrace{(\bar{\nabla} \text{LEGIBILITY} - \lambda \bar{\nabla} C)}_{\bar{\nabla}(\text{LEGIBILITY} - \lambda C)} \quad (27)$$

Note that this is the functional gradient of LEGIBILITY with an additional (linear) *regularizer* λC penalizing *unpredictability*. Substituting in (23) to get the value for λ and using (22) again, we obtain a new update rule:

$$\begin{aligned} \zeta_{i+1} &= \zeta_i + \frac{1}{\eta} M^{-1} \bar{\nabla} \text{LEGIBILITY} - \\ &\quad \underbrace{\frac{1}{\eta} M^{-1} \bar{\nabla} C (\bar{\nabla} C^T M^{-1} \bar{\nabla} C)^{-1} \bar{\nabla} C^T M^{-1} \bar{\nabla} \text{LEGIBILITY} -}_{\text{projection on } \bar{\nabla} C^T(\zeta - \zeta_i) = 0} \\ &\quad \underbrace{M^{-1} \bar{\nabla} C (\bar{\nabla} C^T M^{-1} \bar{\nabla} C)^{-1} (C[\zeta_i] - \beta)}_{\text{offset correction to } \bar{\nabla} C^T(\zeta - \zeta_i) + C[\zeta_i] = \beta} \end{aligned} \quad (28)$$

Fig.5 shows the outcome of the optimization for various β values. In our second experiment below, we analyze what effect β has on the legibility of the trajectory in practice, as measured through users observing the robot's motion.

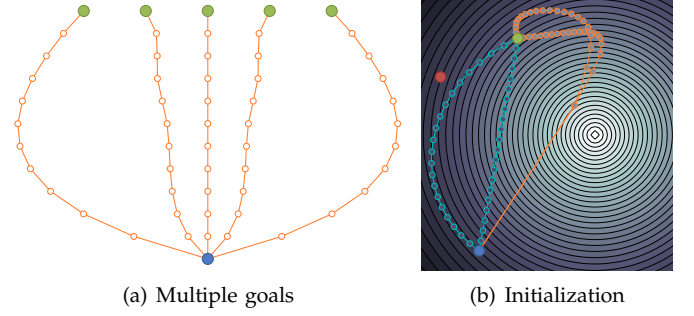


Fig. 8. (a) Legible trajectories for multiple goals. (b) Legibility is dependent on initialization.

IX. 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. VI.

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 (11) 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[45], which trades off between the sum-squared velocity cost we have been using thus far, and a cost penalizing the robot from coming too




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.

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_{GS} , $P(G_R|\xi_{S \rightarrow Q})$ 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.

X. FROM THEORY TO PRACTICE

Predictability and legibility are intrinsically properties that depend on the observer: a real user. Here, we go from the theory of the two properties to what happens in practice, when novice users observe motion.

We present two studies. The first tests our models for predictability and legibility: it tests whether a trajectory that is more legible but less predictable according to our theoretical scores is also more legible but less predictable to an observer.

The second study tests the motion generation: it tests the notion of optimizing for legibility within a trust

region of predictability.

A. The Contradiction Between Predictability and Legibility

The mathematics of predictability and legibility imply that being more legible can mean being less predictable and vice-versa. We set out to verify that this is also true in practice, when we expose subjects to robot motion. We ran an experiment in which we evaluated two trajectories – a theoretically more predictable one ξ_P and a theoretically more legible one ξ_L – in terms of how predictable and legible they are to novices.

Hypothesis. *There exist two trajectories ξ_L and ξ_P for the same task such that ξ_P is more predictable than ξ_L and ξ_L is more legible than ξ_P .*

Task. We chose a task like the one in Fig.1: reaching for one of two objects present in the scene. The objects were close together in order to make this an ambiguous task, in which we expect a larger difference between predictable and legible motion.

Manipulated Variables.

o Character: We chose to use three characters for this task – a simulated point robot, a bi-manual mobile manipulator named HERB [48], and a human – because we wanted to explore the difference between humans and robots, and between complex and simple characters.

o Trajectory: We designed (and recorded videos of) trajectories ξ_P and ξ_L for each of the characters such that




Fig. 9. The trajectories for each character.

$\text{PREDICTABILITY}(\xi_P) > \text{PREDICTABILITY}(\xi_L)$ according to (2), but $\text{LEGIBILITY}(\xi_P) < \text{LEGIBILITY}(\xi_L)$ according to (11) (evaluated based on the parameters from Sec. VI-A).

We describe below several steps we took to eliminate potential confounds and ensure that the effects we see are actually due to the theoretical difference in the score.

With the HERB character, we controlled for effects of timing, elbow location, hand aperture and finger motion by fixing them across both trajectories. For the orientation of the wrist, we chose to rotate the wrist according to a profile that matches studies on natural human motion [21, 36]), during which the wrist changes angle more quickly in the beginning than it does at the end of the trajectory. Fig.15 plots the end effector trace for the HERB trajectories: the gray one has a larger




Fig. 10. The end effector trace for the HERB predictable (gray) and legible (orange) trajectories.

predictability score ($0.54 > 0.42$), while the orange one has a higher legibility score ($0.67 > 0.63$).

With the human character, we used a natural reach for the predictable trajectory, and we used a reach that exaggerates the hand position to the right for the legible trajectory (much like with HERB or the point robot). We cropped the human’s head from the videos to control for gaze effects.

We slowed down the videos to control for timing effects.

Dependent Measures.

o Predictability: Predictable trajectories match the observer’s expectation. To measure how predictable a trajectory is, we showed subjects the character in the initial configuration and asked them to imagine the trajectory they expect the character will take to reach the goal. We then showed them the video of the trajectory and asked them to rate how much it matched the one they expected, on a 1-7 Likert scale. To ensure that they take the time to envision a trajectory, we also asked them to draw what they imagined on a two-dimensional representation of the scene before they saw the video. We further asked them to draw the trajectory they saw in the video as an additional comparison metric.

o Legibility: Legible trajectories enable quick and confident goal prediction. To measure how legible a trajectory is, we showed subjects the video of the trajectory and told them to stop the video as soon as they knew the goal of the character. We recorded the time taken and the prediction.

Subject Allocation. We split the experiment into two sub-experiments with different subjects: one about measuring predictability, and the other about measuring legibility.

For the predictability part, the character factor was between-subjects because seeing or even being asked about trajectories for one character can bias the expectation for another. However, the trajectory factor was within-subjects in order to enable relative comparisons on how much each trajectory matched expectation. This lead to three subject groups, one for each character. We

counter-balanced the order of the trajectories within a group to avoid ordering effects.

For the legibility part, both factors were between-subjects because the goal was the same (further, right) in all conditions. This leads to six subject groups.

We recruited a total of 432 subjects (distributed approximately evenly between groups) through Amazon’s Mechanical Turk, all from the United States and with approval rates higher than 95%. To eliminate users that do not pay attention to the task and provide random answers, we added a control question, e.g. "What was the color of the point robot?" and disregarded the users who gave wrong answers from the data set.

Analysis.

o Predictability: In line with our hypothesis, a factorial ANOVA revealed a significant main effect for the trajectory: subjects rated the predictable trajectory ζ_P as matching what they expected better than ζ_L , $F(1, 310) = 21.88, p < .001$, with a difference in the mean rating of 0.8, but small effect size ($\eta^2=.02$). The main effect of the character was only marginally significant, $F(2, 310) = 2, 91, p = .056$. The interaction effect was significant however, with $F(2, 310) = 10.24, p < .001$. The post-hoc analysis using Tukey corrections for multiple comparisons revealed, as Fig.11(a) shows, that our hypothesis holds for the point robot (adjusted $p < .001$) and for the human (adjusted $p = 0.28$), but not for HERB.

The trajectories the subjects drew confirm this (Fig.12): while for the point robot and the human the trajectory they expected is, much like the predictable one, a straight line, for HERB the trajectory they expected splits between straight lines and trajectories looking more like the legible one.

For HERB, ζ_L was just as (or even more) predictable than ζ_P . We conducted an exploratory follow-up study with novice subjects from a local pool (with no technical background) to help understand this phenomenon. We asked them to describe the trajectory they would expect HERB to take in the same scenario, and asked them to motivate it. Surprisingly, all 5 subjects imagined a different trajectory, motivating it with a different reason.

Two subjects thought HERB’s hand would reach from the right side because of the other object: one thought HERB’s hand is too big and would knock over the other object, and the other thought the robot would be more careful than a human. This brings up an interesting possible correlation between legibility and obstacle avoidance. However, as Fig.7 shows, a legible trajectory still exaggerates motion away from the other candidate objects even in if it means getting closer to a static obstacle.

Another subject expected HERB to not be flexible enough to reach straight towards the goal in a natural

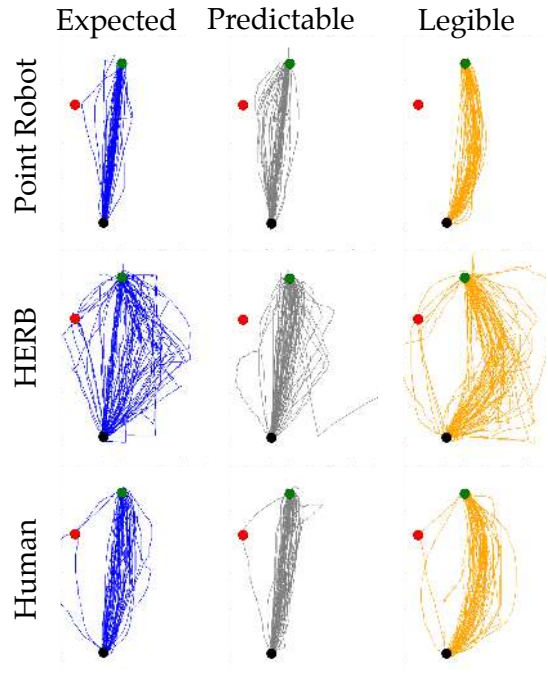


Fig. 12. The drawn trajectories for the expected motion, for ζ_P (predictable), and for ζ_L (legible).

way, like a human would, and thought HERB would follow a trajectory made out of two straight line segments joining on a point on the right. She expected HERB to move one joint at a time. We often saw this in the drawn trajectories with the original set of subjects as well (Fig.12, HERB, Expected).

The other subjects came up with interesting strategies: one thought HERB would grasp the bottle from above because that would work better for HERB’s hand, while the other thought HERB would use the other object as a prop and push against it in order to grasp the bottle.

Overall, that ζ_P was not more predictable than ζ_L despite what the theory suggested because the cost function we assumed did not correlate to the cost function the subjects actually expected. What is more, every subject expected a different cost function, indicating that a predictable robot would have to adapt to the particulars of a human observer.

o Legibility: We collected from each subject the time at which they stopped the trajectory and their guess of the goal. Fig.11(b) (above) shows the cumulative percent of the total number of subjects assigned to each condition that made a correct prediction as a function of time along the trajectory. With the legible trajectories, more of the subjects tend to make correct predictions faster.

To compare the trajectories statistically, we unified time and correctness into a typical score inspired by the Guttman structure (e.g. [7]): guessing wrong gets a score of 0, and guessing right gets a higher score if it

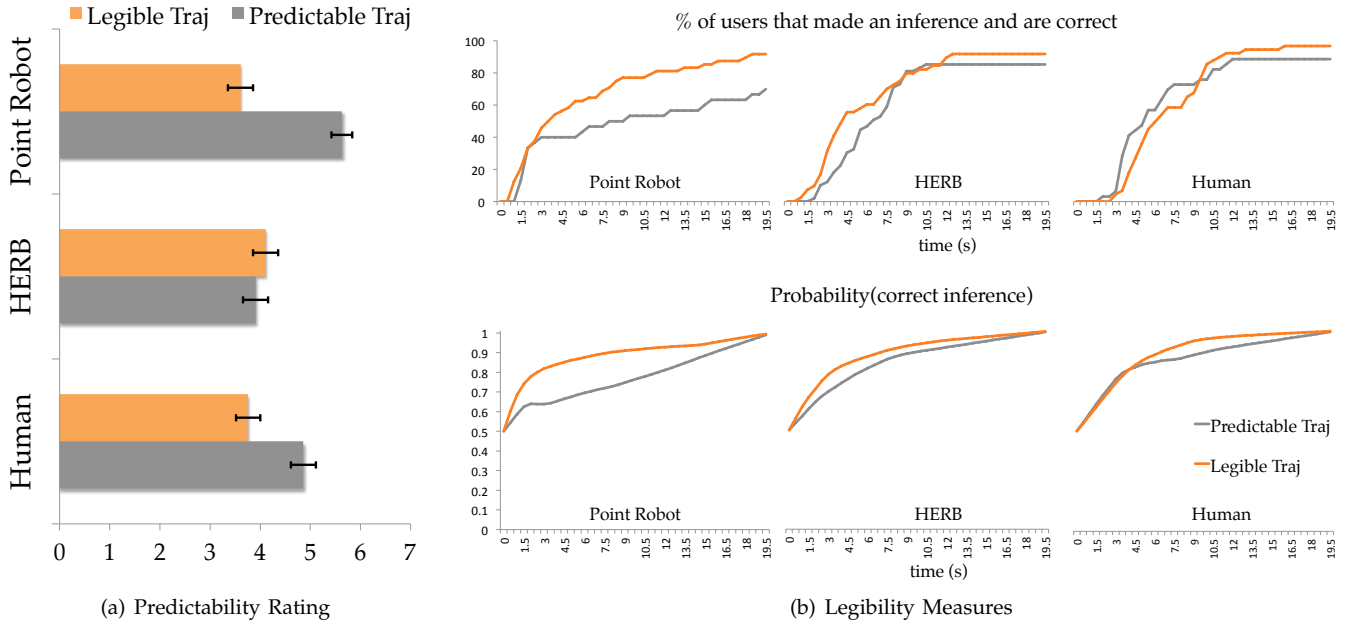


Fig. 11. (a) Ratings (on Likert 1-7) of how much the trajectory matched the one the subject expected. Error bars represent standard error on the mean. (b) Cumulative number of users that responded and were correct (above) and the approximate probability of being correct (below).

happens earlier. A factorial ANOVA predicting this score revealed, in line with our hypothesis, a significant effect for trajectory: the legible trajectory had a higher score than the predictable one, $F(1, 241) = 5.62, p = .019$. The means were 6.75 and 5.73, much higher than a random baseline of making a guess independent of the trajectory at uniformly distributed time, which would result in a mean of 2.5 – the subjects did not act randomly. The effect size was small, $\eta^2 = .02$. No other effect in the model was significant.

Although a standard way to combine timing and correctness information, this score rewards subjects that gave an incorrect answer 0 reward. This is equivalent to assuming that the subject would keep making the incorrect prediction. However, we know this not to be the case. We know that at the end (time T), every subject would know the correct answer. We also know that at time 0, subjects have a probability of 0.5 of guessing correctly. To account for that, we computed an approximate probability of guessing correctly given the trajectory so far as a function of time – see Fig. 11(b)(below). Each subject’s contribution propagates (linearly) to 0.5 at time 0 and 1 at time T . The result shows that indeed, the probability of making a correct inference is higher for the legible trajectory at all times.

This effect is strong for the point robot and for HERB, and not as strong for the human character. We believe that this might be a consequence of the strong bias humans have about human motion – when a human moves even a little unpredictably, confidence in goal prediction drops. This is justified by the fact that subjects did have

high accuracy when they responded, but responded later compared to other conditions. Thus, legible human trajectories would need a stronger emphasis on optimality w.r.t. C (i.e. smaller trust region parameter β in (20)).

Interpretation. Overall, the results do come in partial support of formalism: legible trajectories were more legible for all characters, and predictable trajectories were more predictable for two out of the three characters (not for HERB). However, the effect sizes were small, mainly pointing to the challenge of finding the right cost C for each observer.

B. The Trust Region

So far, we manipulated legibility using two levels. In this section, we test our legibility motion planner, as well as our theoretical notion of a trust region, by analyzing the legibility in practice as the trajectory becomes more and more legible according to our formalism. 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.

Hypotheses.

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

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




Fig. 13. 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. 14. 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.

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 for the point robot character.

Dependent Measures. We measured the legibility of the seven trajectories in the same way as before, combining the time and correctness into a Guttman score as in the Analysis for the previous experiment. 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 ($F(6, 290) = 12.57, p < 0.001$), with a medium-large effect size, $\eta^2 = .2$. Fig.13(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 $p < 0.001$, 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 post-hoc analysis also revealed that the trajectories with $\beta = 20, 40, 80$, or 320 were significantly more legible than the trajectory with $\beta = 10$ ($p = .003, p < .001, p = .004$, and $p = .002$ respectively).

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.14) for different β values, we observed that for higher β s, the wide majority

of users stopped the trajectory in the beginning. The consequence is that *our legibility measure failed to capture whether the mid-part of the trajectory becomes illegible*: the end of the trajectory conveys the goal because it reaches it, but what happens in between the beginning and end? Thus, we ran a follow-up study to verify that legibility in this region does decrease at $\beta = 320$ as compared to our $\beta^* = 40$, in which we explicitly measured the legibility in the middle of the trajectory.

Follow-Up. 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.13 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.13). In an analogous analysis, we found that users in the $\beta = 40$ condition believed this significantly less than users in the $\beta = 320$ condition ($t(57) = 5.7, p < 0.001$).

Interpretation. Overall, the results suggest 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 is 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.*

XI. DISCUSSION

This paper studied motion planning in the presence of an observer who is watching the motion and making inferences about it.

We first formalized predictability and legibility based on the inferences that the observer makes, which have opposing directionality. We then proposed mathematical




Fig. 15. 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.

models for these inferences in order to arrive at predictability and legibility scores that a robot can evaluate. Finally, we derived functional gradient optimization methods for generating predictable or legible motion, as well as a constrained optimization method for optimizing for legibility in a trust region of predictability.

Our studies on novice users provided some support for our models — trajectories more predictable according to the model were overall more predictable in practice, and trajectories more legible according to the model were overall more legible in practice while inside a trust region of predictability. However, three main challenges remain in planning such motions for complex cases (i.e. high-dimensional spaces and non-convex C):

Challenge I: Finding C . If the human observer expects *human-like* motion, cues from animation or biomechanics [22, 25, 37, 57] can help provide good approximations for C . However, our studies suggest that efficiency of *robot* motion has different meanings for different observers (see follow-up experiment in Sec. X-A).

A possibility is to learn from demonstrations provided by the observer. Here, the robot can learn a C that explains the demonstrations[4], using tools like Inverse Optimal Control (IOC) [1, 44, 60]. However, extending these tools to higher dimensions is an open problem [44], and recent work focuses on learning costs that make the

demonstrations locally optimal [32, 38], or on restricting the space of trajectories to one in which optimization is tractable [30].

Aside from investigating the extension of IOC to high-dimensional spaces, we also propose a second thread of research: the idea of familiarizing users to robot behavior. Can users *be taught* a particular C over time? Our preliminary results [16] suggest that familiarization helps for the motion generated by the C from (19), but that it suffers from severe limitations, especially on less natural choices of C . One possibility is that motion for non-anthropomorphic arms is complex enough that we cannot rely solely on the user to do all the learning, suggesting that the two threads of research – familiarization and learning from demonstration – are complementary.

Challenge II: Computing V_G . 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. Legibility optimization happens then in real-time even for high-dimensional cases, as shown in Fig.15.

But this is not the case, for instance, for non-convex functions that require obstacle avoidance, when the robot has many degrees of freedom. In such cases, finding good approximations for V becomes crucial, and many techniques value function approximation techniques could 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.

Challenge III: Finding β . The final challenge is finding how unpredictable the trajectory can become in a given situation. This too can be learned based on the user, or set based on the ambiguity level of the situation (as measured by the legibility score of the predictable trajectory).

Other Future Work Directions. Even though this paper was about goal-directed motion, the formalism for legibility can be applied more generally to transform an efficiency cost C into a legible one. We are excited to investigate this formalism with other channels of communication and other robot morphologies. Recently, we showed the formalism’s applicability to pointing gestures [20], Tellex et al. [50] used the same underlying mathematics to generate legible natural language requests, and Szafir et al. [?] used our results to generate legible quadrotor flight.

Still for goal-directed motion, we are interested in the concept of legibility when each goal could be achieved

in multiple different robot configurations, as it happens in most manipulation tasks [17]. In the example from Fig.1, the added flexibility of goal sets could enable the robot to grasp the object from the side when legible, and closer to the front when predictable.

Another avenue of further research is investigating the role that perceived capability plays. Our follow-up users in the first experiment had different expectations about what HERB is capable of, which shaped their expectations about the motion. Here, familiarization to the robot can potentially be useful in adjusting the perceived capability to better match the real capability. Furthermore, the perceived capability plays a role in what goals the observer might attribute to the robot, which can be captured in our formalism as the prior $P(G)$ over the candidate goals.

Limitations. Our work is limited in many ways. As the previous section discussed, in generating predictable or legible motion, 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 assumptions stop holding. It does not, however, fix the model itself.

Despite its limitations and remaining challenges, this work integrates the idea of an observer and the inferences that he makes directly into motion planning, paving the road to more seamless human-robot collaborations.

ACKNOWLEDGEMENTS

We thank Geoff Gordon, Jodi Forlizzi, Hendrik Christiansen, Kenton Lee, Chris Dellin, Alberto Rodriguez, and the members of the Personal Robotics Lab for fruitful discussion and advice. This material is based upon work supported by NSF-IIS-0916557, NSF-EEC-0540865, ONR-YIP 2012, the Intel Embedded Computing ISTC, and the Intel PhD Fellowship.

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] R. Alami, A. Clodic, V. Montreuil, E. A. Sisbot, and R. Chatila. Toward human-aware robot task planning. In *AAAI Spring Symposium*, pages 39–46, 2006.
- [4] B. Argall, S. Chernova, M. Veloso, and B. Browning. A survey of robot learning from demonstration. *RAS*, 57(5):469 – 483, 2009.
- [5] C. L. Baker, R. Saxe, and J. B. Tenenbaum. Action understanding as inverse planning appendix. *Cognition*, 2009.
- [6] 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.

- [7] G. Bergersen, J. Hannay, D. Sjöberg, 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] 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.
- [11] G. Csibra and G. Gergely. The teleological origins of mentalistic action explanations: A developmental hypothesis. *Developmental Science*, 1:255–259, 1998.
- [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] D. Dey, T. Y. Liu, M. Hebert, and J. A. Bagnell. Contextual sequence prediction with application to control library optimization. In *R:SS*, July 2012.
- [14] A. Dragan, G. Gordon, and S. Srinivasa. Learning from experience in manipulation planning: Setting the right goals. In *ISRR*, 2011.
- [15] A. Dragan, K. Lee, and S. Srinivasa. Legibility and predictability of robot motion. In *Human-Robot Interaction*, 2013.
- [16] A. Dragan, K. Lee, and S. Srinivasa. Familiarization to robot motion. In *Human-Robot Interaction*, 2014.
- [17] A. Dragan, N. Ratliff, and S. Srinivasa. Manipulation planning with goal sets using constrained trajectory optimization. In *ICRA*, May 2011.
- [18] A. Dragan and S. Srinivasa. Generating legible motion. In *Robotics: Science and Systems*, 2013.
- [19] A. Dragan and S. S. Srinivasa. Formalizing assistive teleoperation. In *R:SS*, 2012.
- [20] R. Holladay, A. Dragan and S. S. Srinivasa. Legible Robot Pointing. In *RO-MAN*, 2014.
- [21] J. Fan, J. He, and S. Tillery. Control of hand orientation and arm movement during reach and grasp. *Experimental Brain Research*, 171:283–296, 2006.
- [22] T. Flash and N. Hogan. The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci.*, 5:1688–1703, July 1985.
- [23] G. Gergely, Z. Nadasdy, G. Csibra, and S. Biro. Taking the intentional stance at 12 months of age. *Cognition*, 56(2):165 – 193, 1995.
- [24] M. Gielniak and A. Thomaz. Generating anticipation in robot motion. In *RO-MAN*, 2011.
- [25] M. Gielniak and A. L. Thomaz. Spatiotemporal correspondence as a metric for human-like robot motion. In *ACM/IEEE HRI*, 2011.
- [26] P. Hauf and W. Prinz. The understanding of own and others actions during infancy: You-like-me or me-like-you? *Interaction Studies*, 6(3):429–445, 2005.
- [27] J. Heinzmann and A. Zelinsky. The safe control of human-friendly robots. In *IEEE/RSJ IROS*, 1999.
- [28] C. Igel, M. Toussaint, and W. Weishui. Rprop using the natural gradient. *Trends and Applications in Constructive Approximation*, pages 259–272, 2005.
- [29] T. S. Jim Mainprice, E. Akin Sisbot and R. Alami. Planning safe and legible hand-over motions for human-robot interaction. In *LARP Workshop on Technical Challenges for Dependable Robots in Human Environments*, 2010.
- [30] A. Jain, B. Wojcik, T. Joachims and A. Saxena. Learning Trajectory Preferences for Manipulators via Iterative Improvement. In *NIPS*, 2013.
- [31] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal. STOMP: Stochastic trajectory optimization for motion planning. In *IEEE ICRA*, 2011.
- [32] M. Kalakrishnan, P. Pastor, L. Righetti, and S. Schaal. Learning Objective Functions for Manipulation. In *IEEE ICRA*, 2013.
- [33] K. Kamewari, M. Kato, T. Kanda, H. Ishiguro, and K. Hiraki. Six-and-a-half-month-old children positively attribute goals to human action and to humanoid-robot motion. *Cognitive Development*, 20(2):303 – 320, 2005.
- [34] G. Klien, D. Woods, J. Bradshaw, R. Hoffman, and P. Feltych. Ten challenges for making automation a "team player" in joint human-agent activity. *Intelligent Systems*, nov.-dec. 2004.
- [35] T. Kruse, P. Basili, S. Glasauer, and A. Kirsch. Legible robot navigation in the proximity of moving humans. In *Advanced Robotics and its Social Impacts (ARSO)*, 2012.
- [36] F. Lacquaniti and J. Soechting. Coordination of arm and wrist motion during a reaching task. *J Neurosci.*, 2:399–408, April 1982.
- [37] J. Lasseter. Principles of traditional animation applied to 3d computer animation. In *SIGGRAPH*, 1987.
- [38] S. Levine and V. Koltun. Continuous inverse optimal control with locally optimal examples. In *ICML '12: Proceedings of the 29th International Conference on Machine Learning*, 2012.
- [39] C. Lichtenthaler, T. Lorenz, and A. Kirsch. Towards a legibility metric: How to measure the perceived value of a robot. In *ICSR Work-In-Progress-Track*, 2011.
- [40] C. Lichtenthaler and A. Kirsch. Towards Legible Robot Navigation - How to Increase the Intend Expressiveness of Robot Navigation Behavior. In *International Conference on Social Robotics - Workshop Embodied Communication of Goals and Intentions*, 2013.
- [41] S. Nikolaidis and J. Shah. Human-robot teaming using shared mental models. In *ACM/IEEE HRI*, 2012.
- [42] A. T. Phillips and H. M. Wellman. Infants' understanding of object-directed action. *Cognition*, 98(2):137 – 155, 2005.
- [43] S. Quinlan. *The Real-Time Modification of Collision-Free Paths*. PhD thesis, Stanford University, 1994.
- [44] N. Ratliff, J. A. Bagnell, and M. Zinkevich. Maximum margin planning. In *ICML*, 2006.
- [45] N. Ratliff, M. Zucker, J. A. D. Bagnell, and S. Srinivasa. Chomp: Gradient optimization techniques for efficient motion planning. In *ICRA*, May 2009.
- [46] E. Short, J. Hart, M. Vu, and B. Scassellati. No fair!! an interaction with a cheating robot. In *ACM/IEEE HRI*, 2010.
- [47] B. Sodian and C. Thoermer. Infants' understanding of looking, pointing, and reaching as cues to goal-directed action. *Journal of Cognition and Development*, 5(3):289–316, 2004.
- [48] S. Srinivasa, D. Berenson, M. Cakmak, A. Collet, M. Dogar, A. Dragan, R. Knepper, T. Niemueller, K. Strabala, M. V. Weghe, and J. Ziegler. Herb 2.0: Lessons learned from developing a mobile manipulator for the home. *Proc. of the IEEE, Special Issue on Quality of Life Technology*, 2012.
- [49] D. Szafir, B. Mutlu, and T. Fong. Communication of Intent in Assistive Free Flyers. In *HRI*, 2014.
- [50] L. Takayama, D. Dooley, and W. Ju. Expressing thought: improving robot readability with animation principles. In *HRI*, 2011.
- [51] S. Tellex, R. Knepper, A. Li, D. Rus, and N. Roy. Asking for Help Using Inverse Semantics. In *RSS*, 2014.
- [52] F. Thomas, O. Johnston, and F. Thomas. *The illusion of life: Disney animation*. Hyperion New York, 1995.
- [53] M. A. Tinker. *Legibility of Print*. 1963.
- [54] E. Todorov and W. Li. A generalized iterative lqg method for locally-optimal feedback control of constrained nonlinear stochastic systems. In *ACC*, 2005.
- [55] M. Tomasello, M. Carpenter, J. Call, T. Behne, and H. Moll. Understanding and sharing intentions: the origins of cultural cognition. *Behavioral and Brain Sciences (in press)*, 2004.
- [56] M. Toussaint. Robot trajectory optimization using approximate inference. In *International Conference on Machine Learning*, 2009.
- [57] C. Vesper, S. Butterfill, G. Knoblich, and N. Sebanz. A Minimal Architecture for Joint Action. In *Neural Networks*, 2010.
- [58] A. Witkin and M. Kass. Spacetime constraints. In *SIGGRAPH*, 1988.
- [59] A. L. Woodward. Infants selectively encode the goal object of an actor's reach. *Cognition*, 69(1):1 – 34, 1998.
- [60] E. Wiese, A. Wykowska, J. Zwickel and H. M. Áijller. I see what you mean: how attentional selection is shaped by ascribing intentions to others. In *Journal PLoS ONE*, 2012.
- [61] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, 2008.
- [62] M. Zucker, N. Ratliff, A. Dragan, M. Pivtoraiko, M. Klingensmith, C. Dellin, J. Bagnell, and S. Srinivasa. Covariant hamiltonian optimization for motion planning. *International Journal of Robotics Research (IJRR)*, 2013.