Assistive Planning in Complex, Dynamic Environments: a Probabilistic Approach

Pete Trautman

ABSTRACT

We explore the probabilistic foundations of onboard shared control in complex dynamic environments. In order to do this, we formulate shared control as a random process and describe the joint distribution that governs its behavior. This shared control model is a natural extension of the approach to fully autonomous navigation in human crowds detailed in [25, 23, 24]. We discuss how the probabilistic model relates to the current state of the art in shared control, called "linear blending", and prove the following. First, we recover linear blending as a special case of our probabilistic approach, and in the course of the proof, show what restrictions linear blending imposes. Second, linear blending treats the operator inputs deterministically and implementations can be brittle since they are prone to input device noise or operator jitter. Furthermore, since the operator model is limited to the present time, jointly reasoning over future operator-robot configurations is impossible. In general, we find that for complex dynamic environments linear blending is an inappropriate approach: it is prone to either overriding the operator (i.e., resorting to full autonomy), ceding complete control to the operator (i.e., resorting to full teleoperation), or generating unsafe shared control protocols from safe operator and safe autonomy input. Probabilistic shared control, in contrast, is able to find solutions that are consistently safe and in agreement with the operator. Furthermore, the probabilistic approach is able to achieve such superior solutions in a data driven way—that is, without employing anecdotal heuristics.

1. INTRODUCTION

Fusing human and machine capabilities has been an active research topic in computer science for decades. In robotics related applications, this line of research is often referred to as *shared autonomy* (although "flexible autonomy" and "blended autonomy" are often used interchangeably with shared autonomy). While shared autonomy has been addressed *in toto*, it can be broken down into more specialized areas of research. In particular, some researchers focus on the fusing of human and machine "perception" (or sensing). This line of inquiry has seen a great deal of interest in the human factors community [19], in particular. Similarly, a great deal of research has focused on fusing human and machine "decision making" in the machine learning [31, 17], control theory [12], and human robot interaction communities [14].

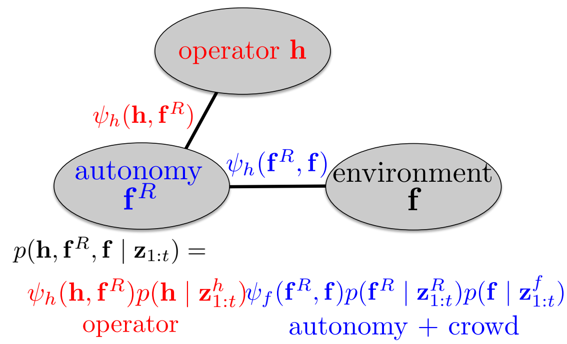

Figure 1: Diagram depicting the random processes and relationships between those random processes. Below the diagram is the shared control distribution, color coded to match the diagram.

Under the umbrella of shared decision making, an even more focussed line of research has emerged: shared control, whereby the moment to moment control commands sent to the platform motors are a synthesis of human input and autonomy input. Broadly speaking, shared control has been deployed in two cases: shared teleoperation (where the human operator is not co-located with the robot) and onboard shared control (where the human operator is physically on the robot). Shared teleoperation has seen activity in a number of application areas: search and rescue robots [4] and extraterrestrial robotics [2] are just two examples.

Onboard shared control has likewise enjoyed a wide variety of use cases: assistive wheelchair technology [13], assistive automobile driving, and assistive manufacturing vehicle operation (e.g., forklifts) are just a few of the examples. For any of these cases, shared control can be broken down into an autonomous modeling step, a human modeling step, and a human-machine arbitration step (see [10] for a compelling argument justifying this classification scheme).

In this paper, we explore the probabilistic foundations of onboard shared control in the presence of dynamic obstacles (e.g., the crowd)—see Figure 1 for an illustration of our model. In order to do this, we formulate shared control as a random process and describe the joint distribution (over the operator, autonomy, and crowd) that governs behavior. We then propose a tractable model of the full joint distribution and discuss how both this model and the full joint distribution relates to the current state of the art in shared control, called "linear blending." In particular, we show how linear blending can be interpreted as a special case of probabilistic shared control, and in the process, observe a number of things: first, in linear blending the operator is often not modeled, and so there is no flexibility in how operator commands are interpreted during optimization of the joint distribution. Thus, if the operator inputs are noisy or misleading, we have no way to account for this. Second, in linear blending, the autonomy reasons *independently* of the operator during optimization, so even if we had a precise operator model, the autonomy would not be informed of this information until after the optimization. Third, in linear blending, the autonomy is limited to a single optimal decision, which is then averaged with the operator input—this approach leaves no flexibility in how the autonomy might assist the user. Next, we present an extension of linear blending that can accommodate more than one statistic about the user (e.g., user inferred goal information, preferred trajectories, etc), and find that state of the art approaches to this problem are statistically unsound. To this end, we present a statistically valid model for how to properly condition the autonomy on user statistics. We conclude with an examination of linear blending in complex dynamic environments and find that it is inappropriate for such situations. We similarly examine the probabilistic model in complex dynamic environments and discover that it is flexible enough to maintain safety in crowds while acting agreeably with the operator.

2. RELATED WORK

We begin by defining the arbitration step of linear blending:

$$u^s_{LB}(t) = K_h u^h_t + K_R u^R_{t+1}, \quad (2.1)$$

where, at time t, $u^s_{LB}(t)$ is the linearly blended shared control command sent to the platform actuators, u^h_t is the human operator input (joystick deflections, keyboard inputs, etc.), u^R_{t+1} is the next autonomy command, and K_h, K_R are the operator and autonomy *arbitration parameters*, respectively. To ensure that the magnitude of $u^s_{LB}(t)$ does not exceed the magnitude of u^h_t or u^R_{t+1}, we require that $K_h + K_R = 1$.

This linear arbitration model has enjoyed wide adoption in the assistive wheelchair community ([29, 18, 28, 32, 20, 27, 16, 5]). More generally, shared control path planning researchers have widely adopted Equation 2.1 as a *de-facto* standard protocol—see [10, 9] for an extensive argument about how "linear policy blending can act as a common lens across a wide range of literature". Additionally, the work of [21, 30] advocates the broad adoption of a linear arbitration step for shared control.

For the purposes of this article, we describe how each quantity of Equation 2.1 is computed:

1) Collect the data at time t: $u^h_{1:t}$ are the historical operator inputs. $\mathbf{z}^R_{1:t}$ are the historical measurements of the state of the robot (such as odometry, localization, SLAM output, etc).

2) Compute the *autonomous* input u^R_{t+1}: This quantity may be computed using nearly any off the shelf planning algorithm, and is dependent on application. The "Dynamic Window Approach" [11] and "Vector Field Histograms+" [26] are popular approaches to perform obstacle avoidance for wheelchairs. Sometimes, the autonomy is biased according to data about the operator—for instance, one might imagine an offline training phase where the robot is taught "how" to move through the space, and then this data could be agglomerated using, e.g., inverse optimal control. Alternatively, one might bias the autonomous decision making by conditioning the planner on the predicted or known human goal.

3) Compute the arbitration parameters K_h and K_R. A wide variety of heuristics have been adopted to compute this parameter: confidence in robot trajectory, smoothness, mitigating jerk, operator reliability, user desired trajectories, safeguarding against unsafe trajectories, etc. Indeed, much of the shared control literature is devoted to developing novel heuristics to compute this parameter.

4) Compute the shared control $u^s_{LB}(t)$ using Equation 2.1.

In the references above, the data $u^h_{1:t}$ are interpreted literally—no likelihood or predictive model filters this data stream. In other approaches, operator intention is modeled using a combination of dynamic Bayesian networks and/or Gaussian mixture models.

For this paper, we adopt the notation $\mathbf{z}^h_t \doteq u^h_t$, (we treat operator inputs as *measurements* of the operator *trajectory*, $\mathbf{h}\colon t \in \mathbb{R} \to \mathcal{X}$, where $\mathcal{X}$ is the action space). Similarly, we define measurements $\mathbf{z}^R_{1:t}$ of the robot trajectory $\mathbf{f}^R$ and measurements $\mathbf{z}^i_{1:t}$ of the i'th static or dynamic obstacle trajectory $\mathbf{f}^i$. We thus work in the space of distributions over the operator function $\mathbf{h}$, autonomy function $\mathbf{f}^R$, and crowd function $\mathbf{f} = (\mathbf{f}^1, \dots, \mathbf{f}^{n_t})$, measured through $\mathbf{z}^f_{1:t}$. The integer n_t is the number of people in the crowd at time t.

We comment on the work of [7, 15, 8]. In these papers, the authors construct a probabilistic model over user forward trajectories (i.e., a specific and personalized instantiation of $p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$). They then formulate shared control as a POMDP, in order to capture the effects of robot actions on the probabilistic model of the operator. However, for tractability of the POMDP, the autonomy is limited to only choosing from the next available state—this makes assistive navigation through crowds impossible (as shown in [24]). Further, the potential autonomous actions are limited by the input device; in one application, the autonomy was only able to reason over 9 directions. As we discuss in Section 6, limiting the autonomy to such a substantial degree can have substantial effects on performance.

3. FOUNDATIONS AND IMPLEMENTATION OF PROBABILISTIC SHARED CONTROL

As stated in Section 1, we seek to explore the probabilistic foundations of assistive onboard shared control; we thus posit that

$$\begin{aligned} u^s(t) &= \mathbf{f}^{R^*}_{t+1} \\ (\mathbf{h}, \mathbf{f}^R, \mathbf{f})^* &= \operatorname*{arg\,max}_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}^h_{1:t}, \mathbf{z}^R_{1:t}, \mathbf{z}^f_{1:t}). \end{aligned} \quad (3.1)$$

That is, assistive onboard shared control is the MAP value of the joint distribution over the operator, autonomy, and crowd.

We suggest this approach partly based on what we have learned from fully autonomous navigation in human crowds in [23], and partly based on the following: by formulating shared control as the MAP value of a joint probability distribution, we can directly explore fundamental modeling limitations imposed by the linear blending approach, and thereby explore what consequences these modeling assumptions have on performance.

In the remainder of this section, we explain how the approach in Equation 3.1 is a natural extension of our previous work in [25], and present a tractable model of the joint distribution over the operator, autonomy, and crowd. We use the probabilistic graphical model in Figure 2 to guide our derivation.

Definition 1. *A cooperative human crowd navigation model (as in [25]) is described by*

$$p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) = \psi(\mathbf{f}^R, \mathbf{f}) p(\mathbf{f}^R \mid \mathbf{z}^R_{1:t}) \prod_{i=1}^{n_t} p(\mathbf{f}^i \mid \mathbf{z}^i_{1:t}) \quad (3.2)$$

where $\bar{\mathbf{z}}_{1:t} = [\mathbf{z}^R_{1:t}, \mathbf{z}^f_{1:t}]$, $\psi(\mathbf{f}^R, \mathbf{f})$ *is the crowd potential function, and* $p(\mathbf{f}^R \mid \mathbf{z}^R_{1:t}), p(\mathbf{f}^i \mid \mathbf{z}^i_{1:t})$ *are the robot and crowd individual kinematic functions, respectively. The corresponding graphical model is presented on the left hand side of Figure 2.*

Definition 2. *A probabilistic shared control (PSC) model in the presence of static or dynamic obstacles is*

$$p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) \propto \psi(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) p(\mathbf{f}^R \mid \mathbf{z}^R_{1:t}) \prod_{i=1}^{n_t} p(\mathbf{f}^i \mid \mathbf{z}^i_{1:t})$$
$$= \psi(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) \quad (3.3)$$

where $p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$ *is the predictive distribution over the operator,* $\mathbf{z}^h_{1:t}$ *is data generated by the operator, and* $\psi(\mathbf{h}, \mathbf{f}^R)$ *is the potential function between the operator and the robot. For the purposes of this article, we choose*

$$\psi(\mathbf{h}, \mathbf{f}^R) = \exp(-\frac{1}{2\gamma}(\mathbf{h} - \mathbf{f}^R)(\mathbf{h} - \mathbf{f}^R)^\top).$$

The corresponding graphical model is presented on the right hand side of Figure 2. We note that the term γ *captures how "tightly" the autonomy is coupled to the operator.*

We note that this approach is an approximation to the full posterior $p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t})$; this approximation is motivated by our previous work in [24]. Additionally, we choose to not include the dynamic environment $\mathbf{f}$ in the potential function $\psi(\mathbf{h}, \mathbf{f}^R)$ because modeling operator to crowd interactions is not well understood at this point (e.g., an operator yelling at the crowd might be difficult to both sense and difficult for the autonomy to respond to; we defer such models to later work).

4. LINEAR BLENDING AS A PROBABILISTIC SHARED CONTROL MODEL

In order to understand linear blending in the context of Equation 3.1, we begin by considering the conditioning relationships of the full joint distribution:

$$p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}^R_{1:t}, \mathbf{z}^f_{1:t}, \mathbf{h}) p(\mathbf{h} \mid \mathbf{z}^h_{1:t}).$$

To understand how the linear abitration modeling assumptions effect the full joint, we first insert the linear blending

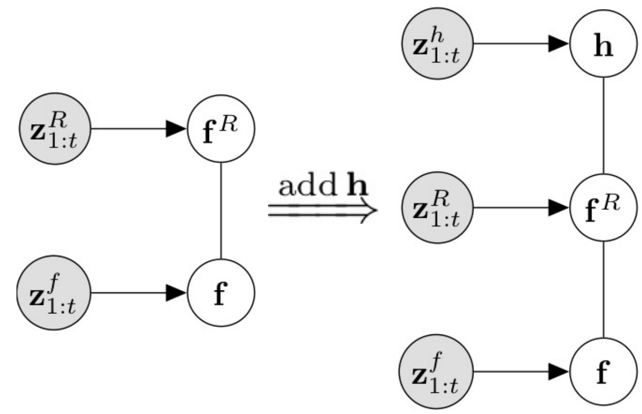

Figure 2: The correspondence between the fully autonomous cooperative collision avoidance model of [25] and Equation 3.2 (left hand side), and the model of probabilistic shared control in Equation 3.3.

operator model: $p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) = \delta(\mathbf{h} - \mathbf{z}^h_t)$. Thus,

$$p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}^f_{1:t}, \mathbf{z}^R_{1:t}, \mathbf{h}) \delta(\mathbf{h} - \mathbf{z}^h_t)$$
$$= p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}^f_{1:t}, \mathbf{z}^R_{1:t}, \mathbf{z}^h_t).$$

In this case,

$$\arg\max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = \arg\max_{\mathbf{f}^R, \mathbf{f}} p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}^f_{1:t}, \mathbf{z}^R_{1:t}, \mathbf{z}^h_t).$$

Because this distribution is already conditioned on $\mathbf{z}^h_t$, there is no need for a linear arbitration step (we will see this explicitly in Theorem 4.1). Furthermore, by fixing the operator state at $\mathbf{z}^h_t$, we are no longer *jointly* optimizing over the autonomy and the operator.

Theorem 4.1 (**Equation 3.1 generalizes linear blending**). *Let*

$$p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}^f_{1:t}, \mathbf{z}^R_{1:t}, \mathbf{z}^h_t) = \psi_h(\mathbf{z}^h_t, \mathbf{f}^R_{t+1}) p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$$

and suppose that we use Laplace's Approximation [3] to model the autonomous part of the distribution:

$$p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) = \mathcal{N}(\mathbf{f}^R_{t+1} \mid \mathbf{f}^{\bar{R}}{}_{t+1}, \boldsymbol{\sigma}^R).$$

where $\mathbf{f}^{\bar{R}}{}_{t+1}$ *is a mode of the distribution. Then the probabilistic shared control is*

$$u^s_{PSC}(t) = \boldsymbol{\sigma}(\frac{1}{\gamma}\mathbf{z}^h_t + \frac{1}{\boldsymbol{\sigma}^R}\mathbf{f}^{\bar{R}})$$

where $\boldsymbol{\sigma}^{-1} = (\gamma^{-1} + (\boldsymbol{\sigma}^R)^{-1})$ *and* γ *is the operator-autonomy attraction parameter.*

Proof. We compute

$$u^s(t) = \arg\max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t})$$
$$\propto \arg\max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} \mathcal{N}(\mathbf{f}^R_{t+1} \mid \mathbf{z}^h_t, \gamma) \mathcal{N}(\mathbf{f}^R_{t+1} \mid \mathbf{f}^{\bar{R}}{}_{t+1}, \boldsymbol{\sigma}^R)$$
$$= \arg\max_{\mathbf{f}^R_{t+1}} \mathcal{N}(\mathbf{f}^R_{t+1} \mid \boldsymbol{\mu}, \boldsymbol{\sigma}) \quad (4.1)$$
$$= \boldsymbol{\sigma}(\frac{1}{\gamma}\mathbf{z}^h_t + \frac{1}{\boldsymbol{\sigma}^R}\mathbf{f}^{\bar{R}})$$

where

$$\boldsymbol{\mu} = \boldsymbol{\sigma}(\frac{1}{\gamma}\mathbf{z}_t^h + \frac{1}{\boldsymbol{\sigma}^R}\bar{\mathbf{f}^R})$$
$$\boldsymbol{\sigma}^{-1} = (\gamma^{-1} + (\boldsymbol{\sigma}^R)^{-1}).$$

□

We now examine the how the assumptions of linear blending can effect performance. First, the constraint $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) = \delta(\mathbf{h} - \mathbf{z}_t^h)$ can lead to very brittle implementations. In the presence of noisy input, for instance, using the unfiltered data $\mathbf{z}_t^h$ can cause the shared control to oscillate with the input noise (or the jitter of the operator). Furthermore, since the operator model is limited to the present time, jointly reasoning over future operator-robot configurations is impossible. For many applications, blending the autonomy with the human's *trajectory* can be critically important to instantiate fluid sharing.

Further, the assumption that the joint robot-crowd distribution is a unimodal Gaussian $p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f) = \mathcal{N}(\mathbf{f}_{t+1}^R \mid \bar{\mathbf{f}^R}_{t+1}, \boldsymbol{\sigma}_{t+1}^R)$ can lead to severe restrictions in shared control capability; even though a locally optimal autonomy strategy $\bar{\mathbf{f}^R}$ is included in the linear blend $u_{LB}^s \propto K_h\mathbf{z}_t^h + K_R\bar{\mathbf{f}^R}_{t+1}$, there is no guarantee that the human operator will choose this optima, or even a nearby optimum.

By maintaining the multitude of hypotheses inherent to $p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$, we greatly increase the possibility that the autonomy will be able to assist the operator in a way that is both desirable and safe. As an example, linear trajectory blending over a *safe* operator input and a *safe* autonomous input can result in an unsafe shared trajectory. That is, the weighted average of two safe trajectories can be averaged into an unsafe trajectory.

Importantly, the issue of averaging two safe trajectories into an unsafe one has been addressed in the literature using a heuristic called "safeguarding" (e.g., discard shared trajectories that are unsafe, and recompute a safe trajectory—perhaps stopping the platform entirely—as in [6]). However, as with any heuristic, safeguarding begs the question: can we develop an arbitration protocol that avoids this emergency protocol? Furthermore, safeguarding has the potential for causing *frozen operator-autonomy* states: suppose the safeguard chooses to stop the platform, but the operator does not understand why the platform has stopped. If the platform is stopped and the environment is not moving (either a standing crowd or a cluttered obstacle field), then the autonomy will continue to compute the same optima, and so long as the operator continues to insist on the same input that triggered the safeguard, the platform will not move. In this case, a forward shared control solution exists, but because linear blending only considers a single autonomy strategy, the discord between the operator and the autonomy cannot be resolved.

5. CONDITIONAL TRAJECTORY BLENDING

In this section, we extend our definition of linear blending to capture the salient characteristics of recent approaches; broadly speaking, this line of work seeks to include additional data about the operator during the arbitration step. Further, we introduce a probabilistic model, called conditional trajectory blending, that is able to incorporate any information about the operator in a statistically principled way.

Definition 3 (**Extending Equation 2.1**). *Let*
1) $\bar{\mathbf{h}} \sim p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ *be defined through time* $T > t$, *such that* $\bar{\mathbf{h}}\colon [1, T] \subset \mathbb{R} \mapsto \mathcal{X}$,
2) $\bar{\mathbf{f}^R} \sim p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f)$, *with* $\bar{\mathbf{f}^R}\colon [1, T] \subset \mathbb{R} \mapsto \mathcal{X}$,
3) u_{LTB}^s *be the shared control* $u_{LTB}^s\colon [1, T] \subset \mathbb{R} \to \mathcal{X}$,
4) and K_h, K_R *be the arbitration parameters, similar to the arbitration parameter of Section 2.*

Definition 4 (**LTB**). *Linear Trajectory Blending (LTB) is the process*
1) Sample the autonomy $\bar{\mathbf{f}^R} = \arg\max_{\mathbf{f}^R, \mathbf{f}} p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f)$.
2) Sample the operator $\bar{\mathbf{h}} = \arg\max_{\mathbf{h}} p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$.
3) Construct the shared control

$$u_{LTB}^s = K(\frac{1}{K_h}\bar{\mathbf{h}} + \frac{1}{K_R}\bar{\mathbf{f}^R}). \tag{5.1}$$

By extending to trajectories, we can more easily incorporate operator information into the arbitration step. In this vein, we now define an extension of linear trajectory blending that biases the autonomous decision making on operator data (motivated by the approach in [10]).

Definition 5 (**LTBo**). *Let* $p(G \mid \mathbf{z}^h)$ *be a distribution about the operator, where* $\mathbf{z}^h \subseteq \mathbf{z}_{1:t}^h$. *Then* operator biased linear trajectory blending *(LTBo) is the process*
1) Sample $G_1 = \arg\max_G p(G \mid \mathbf{z}^h)$. G_1 *could be a full predicted trajectory, a waypoint, a goal, or any other appropriate quantity.*
2) Sample the operator biased *distribution*

$$\bar{\mathbf{f}^R}_{\mathbf{h}} = \arg\max_{\mathbf{f}^R, \mathbf{f}} p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f, G_1).$$

3) Sample the operator trajectory $\bar{\mathbf{h}} = \arg\max_{\mathbf{h}} p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$.
4) Construct the shared control

$$u_{LTBo}^s = K(\frac{1}{K_h}\bar{\mathbf{h}} + \frac{1}{K_R}\bar{\mathbf{f}^R}_{\mathbf{h}}).$$

With these definitions, we extend the approach of Equation **??** to include 1) a model of the operator $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$ and 2) to enable seeding the autonomy with information about the operator. By modeling the operator with a distribution, we potentially bypass the issue of noisy inputs leading to "jittery" linear blends. Also, by seeding the autonomy with operator statistics, we can potentially drive the autonomy towards solutions that are more closely aligned with user desire.

While this approach is sensible, it is not immediately clear why we should extend linear blending in such a way. We thus pause to examine this approach in the context of the full joint distribution. In particular, we note that

$$p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) = p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{h}) p(\mathbf{h} \mid \mathbf{z}_{1:t}^h). \tag{5.2}$$

If we were to only sample a single statistic of the operator distribution $\bar{\mathbf{h}} = \arg\max_{\mathbf{h}} p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$, then we are implicitly making the modeling assumption that $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) = \delta(\mathbf{h} - \bar{\mathbf{h}})$; thus we have

$$p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) \propto p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \mathbf{h})\delta(\mathbf{h} - \bar{\mathbf{h}})$$
$$= p(\mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}^f, \mathbf{z}_{1:t}^R, \bar{\mathbf{h}}).$$

In this case,

$$\arg\max_{\mathbf{h},\mathbf{f}^R,\mathbf{f}} p(\mathbf{h},\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}_{1:t}) = \arg\max_{\mathbf{f}^R,\mathbf{f}} p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^f_{1:t},\mathbf{z}^R_{1:t},\bar{\mathbf{h}}),$$

and so (as before) there is no *need* for a linear arbitration step—we have already conditioned the autonomy on the operator, and we can recover linear blending by choosing $p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^f_{1:t},\mathbf{z}^R_{1:t},\bar{\mathbf{h}})$ as in Theorem 4.1. If we have a *separate* model $p(G \mid \mathbf{z}^h)$ about the operator, we know that it does not contain information beyond what is available in $p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$, since $\mathbf{z}^h \subseteq \mathbf{z}^h_{1:t}$ (see the discussion on the "data processing inequality" in [22]). With the definition below, we show how to combine multiple operator data points in a statistically sound manner.

Definition 6 (**Conditional Trajectory Blending**). *Assume that we have the distributions $p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$ and $p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^f_{1:t},\mathbf{z}^R_{1:t},\mathbf{h}_b)$. Then let $\{\mathbf{h}^b\}_{b=1}^{N_h} \sim p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$ be a collection of N_h samples of the operator model. If we take the model of the operator to be $p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) = \sum_{b=1}^{N_h} w^b\delta(\mathbf{h}-\mathbf{h}^b)$, then*

$$\begin{aligned}
p(\mathbf{h},\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}_{1:t}) &= p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^f_{1:t},\mathbf{z}^R_{1:t},\mathbf{h})p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) \\
&= p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^f_{1:t},\mathbf{z}^R_{1:t},\mathbf{h})\sum_{b=1}^{N_h} w^b\delta(\mathbf{h}-\mathbf{h}^b) \\
&= \sum_{b=1}^{N_h} w^b p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^f_{1:t},\mathbf{z}^R_{1:t},\mathbf{h}^b) \qquad (5.3)
\end{aligned}$$

where $w^b = p(\mathbf{h}^b = \mathbf{h} \mid \mathbf{z}^h_{1:t})$ is the probability of sample $\mathbf{h}^b$. We interpret the shared control to be

$$u^s_{CTB} = \arg\max_{\mathbf{f}^R,\mathbf{f}} \sum_{b=1}^{N_h} w^b p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^f_{1:t},\mathbf{z}^R_{1:t},\mathbf{h}^b).$$

In particular, we revisit the case of LTBo: suppose that we sample $G_1 \sim p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$—$G_1$ is present in this distribution since it contains all the data—and then sample $\bar{\mathbf{h}} \sim p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$. Then conditional trajectory blending tells us that we should find the arg max of the distribution

$$u^s_{CTB} = \arg\max_{\mathbf{f}^R,\mathbf{f}} \Big[w^1 p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^f_{1:t},\mathbf{z}^R_{1:t},G_1) + w^2 p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^f_{1:t},\mathbf{z}^R_{1:t},\bar{\mathbf{h}})\Big].$$

In general, then, $u^s_{CTB} \neq u^s_{LTBo}$; however, since conditional trajectory blending is derived directly from the full joint we know that it is combining the data $G_1,\bar{\mathbf{h}}$ in a statistically sound manner. The next theorem provides information about the limiting behavior of conditional trajectory blending.

Theorem 5.1 (**CTB approximates PSC**). *As the number of operator samples tends to infinity, probabilistic shared control (Equation 3.1) is recovered.*

Proof. Representing $p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) = \sum_{b=1}^{\infty} w^b\delta(\mathbf{h}-\mathbf{h}^b)$,

$$\begin{aligned}
p(\mathbf{h},\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}_{1:t}) &= p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^R_{1:t},\mathbf{z}^f_{1:t},\mathbf{h})p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) \\
&= p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^R_{1:t},\mathbf{z}^f_{1:t},\mathbf{h})\sum_{b=1}^{\infty} w^b\delta(\mathbf{h}-\mathbf{h}^b) \\
&= \sum_{b=1}^{\infty} w^b p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^R_{1:t},\mathbf{z}^f_{1:t},\mathbf{h}^b).
\end{aligned}$$

□

It is important to emphasize that conditioning the autonomy on operator statistics and then averaging $\bar{\mathbf{f}^R}_{\mathbf{h}}$ with a separate operator statistic $\bar{\mathbf{h}}$ is not just unnecessary, but potentially statistically unsound as well, since it is unclear how such an approach originates from Equation 5.2.

Lemma 5.2 (**LTB statistically unsound**). *LTB is not guaranteed to incorporate data in a statistically sound manner.*

Proof. Suppose that one were to sample $\bar{\mathbf{h}} \sim p(G \mid \mathbf{z}^h)$, then sample $\bar{\mathbf{h}} = \arg\max_{\mathbf{h}} p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$, then compute $\bar{\mathbf{f}^R}_{\mathbf{h}}$, and then compute $K_h\bar{\mathbf{h}} + K_R\bar{\mathbf{f}^R}_{\mathbf{h}}$. Since we have incorporated $\bar{\mathbf{h}}$ *twice* in the linear blend, the data has been overused. One could potentially compensate for this by "removing" the effect of double usage of $\bar{\mathbf{h}}$ in K_h, but it is unclear how to do this in a statistically sound manner. □

Thus, if we have a model of $p(\mathbf{f}^R,\mathbf{f} \mid \mathbf{z}^R_{1:t},\mathbf{z}^f_{1:t},\mathbf{h}^b)$ and a model of the operator $p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$, we have a clear mandate (under the probabilistic line of reasoning) for how to correctly formulate our shared control algorithm.

6. OPTIMALITY OF SHARED CONTROL

Consider the following Gaussian sum approximations:

$$p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) = \sum_{m=1}^{N_h} \mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}_m,\boldsymbol{\Sigma}_m)$$

and

$$p(\mathbf{f}^R,\mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) = \sum_{n=1}^{N_R} \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_n,\boldsymbol{\Sigma}_n).$$

Then we have that

$$\begin{aligned}
&\psi_h(\mathbf{h},\mathbf{f}^R)p(\mathbf{h} \mid \mathbf{z}^h_{1:t})p(\mathbf{f}^R,\mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) \\
&\approx \psi_h(\mathbf{h},\mathbf{f}^R)\sum_{m=1}^{N_h}\alpha_m\mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}_m,\boldsymbol{\Sigma}_m)\sum_{n=1}^{N_R}\beta_n\mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_n,\boldsymbol{\Sigma}_n), \qquad (6.1)
\end{aligned}$$

and as $N_h, N_R \to \infty$ and the covariances $\boldsymbol{\Sigma}_m,\boldsymbol{\Sigma}_n$ approach the zero matrix (e.g., as the Gaussian sums approach sums of Dirac samples), we recover the densities $p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$ and $p(\mathbf{f}^R,\mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$—see [1]. We note that the approximation $\sum_{n=1}^{N_R}\beta_n\mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_n,\boldsymbol{\Sigma}_n)$ consists of only safe modes, since $\psi_f(\mathbf{f}^R,\mathbf{f})$ will assign near zero β to any modes that are on a collision course. This is not true for the Gaussian mixture assigned to the operator, since the operator may very well make choices which place the platform on a collision course (resolving such situations are a key functionality of assistive shared control).

From Theorem 4.1, we know that linear trajectory blending makes the assumption that

$$\begin{aligned}
&\psi_h(\mathbf{h},\mathbf{f}^R)p(\mathbf{h} \mid \mathbf{z}^h_{1:t})p(\mathbf{f}^R,\mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) \\
&= \psi_h(\mathbf{h},\mathbf{f}^R)\delta(\mathbf{h}-\bar{\mathbf{h}})\mathcal{N}(\mathbf{f}^R \mid \bar{\mathbf{f}^R},\boldsymbol{\Sigma}) \\
&= \psi_h(\bar{\mathbf{h}},\mathbf{f}^R)\mathcal{N}(\mathbf{f}^R \mid \bar{\mathbf{f}^R},\boldsymbol{\Sigma}).
\end{aligned}$$

That is, $N_R = 1$, with mean chosen at the largest mode of the Gaussian sum, and $N_h = 1$ with mean chosen at the largest mode of the Gaussian sum, with zero covariance. If these conditions are true (one mode of the autonomy-crowd distribution dominates, and one mode of the operator distribution dominates and has small covariance), then u^s_{LTB} is the optimal solution.

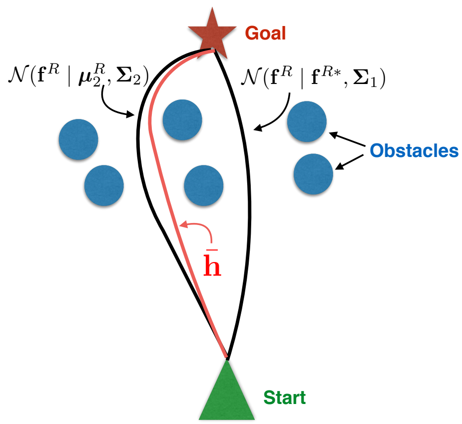

Figure 3: One global autonomy optima at $\mathbf{f}^{R*}$ and a safe but suboptimal autonomy mode at $\boldsymbol{\mu}^R_2$ through some obstacle field (additional autonomous modes exist but we leave them off for clarity). The operator's unimodal predicted trajectory at $\bar{\mathbf{h}}$ is *safe*. Covariance functions removed for clarity.

However, if these conditions are not met—in crowds, for instance, many modes of $p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t})$ have significant weight—then u^s_{LTB} is not the optimal solution. Further, for situations in which the operator trajectory may exhibit multimodality (Figure 5), u^s_{LTB} is again not the optimal solution. Importantly, then, linear trajectory blending insufficiently addresses scenarios in which either the autonomous distribution or the operator distribution is multimodal.

To understand why this is important, first consider the illustration in Figure 3. Suppose that there are two nontrivial safe modes that pass through the obstacle field: $\mathcal{N}(\mathbf{f}^R \mid \mathbf{f}^{R*}, \mathbf{\Sigma}_1)$ and $\mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}^R_2, \mathbf{\Sigma}_2)$. Thus

$$
\begin{aligned}
&\psi_h(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) \\
&= \psi_h(\bar{\mathbf{h}}, \mathbf{f}^R)[\beta_1 \mathcal{N}(\mathbf{f}^R \mid \mathbf{f}^{R*}, \mathbf{\Sigma}_1) + \beta_2 \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}^R_2, \mathbf{\Sigma}_2)] \\
&= \frac{\beta_1}{Z_1} \mathcal{N}(\mathbf{f}^R \mid \bar{\mathbf{f}}^{R*}, \bar{\mathbf{\Sigma}}_1) + \frac{\beta_2}{Z_2} \mathcal{N}(\mathbf{f}^R \mid \bar{\boldsymbol{\mu}}^R_2, \bar{\mathbf{\Sigma}}_2)
\end{aligned}
$$

where $\beta_1 > \beta 2$ (since the first mode is the global optima), and the means and covariances in the last line are of the form of Equation 5.1. However

$$
\begin{aligned}
\frac{1}{Z_1} &\propto \exp\left(-\frac{1}{2}\left(\bar{\mathbf{h}} - \mathbf{f}^{R*}\right)^\top (\gamma + \mathbf{\Sigma}_1)^{-1} \left(\bar{\mathbf{h}} - \mathbf{f}^{R*}\right)\right) \\
\frac{1}{Z_2} &\propto \exp\left(-\frac{1}{2}\left(\bar{\mathbf{h}} - \boldsymbol{\mu}_2\right)^\top (\gamma + \mathbf{\Sigma}_2)^{-1} \left(\bar{\mathbf{h}} - \boldsymbol{\mu}_2\right)\right),
\end{aligned}
$$

and so $1/Z_1$ is exponentially smaller than $1/Z_2$ since $\bar{\mathbf{h}} - \mathbf{f}^{R*}$ is much larger than $\bar{\mathbf{h}} - \boldsymbol{\mu}_2$. Thus, the probabilistic shared control in this situation is very close to both $\boldsymbol{\mu}^R_2$ and $\bar{\mathbf{h}}$.

Conversely, $u^s_{LTB} \propto K_h \bar{\mathbf{h}} + K_R \mathbf{f}^{R*}$. Here, if K_h is close to K_R, then u^s_{LTB} is unsafe and thus safeguarding has to be employed. If $K_R \gg K_h$, then the autonomy unnecessarily overrides the operator's (safe) choice. If $K_h \gg K_R$, then the operator is controlling the platform, and is thus not not being assisted. The question underlying these three cases is the following: what heuristic should be employed to choose K_h and K_R? By carrying multiple modes (as in probabilistic shared control), we bypass this dilemma, since heuristics are never invoked: basic rules of probability theory (namely, the normalizing factor) determine the best choice. In other words, probabilistic shared control is able to determine the shared control in a data driven way rather than through anecdote.

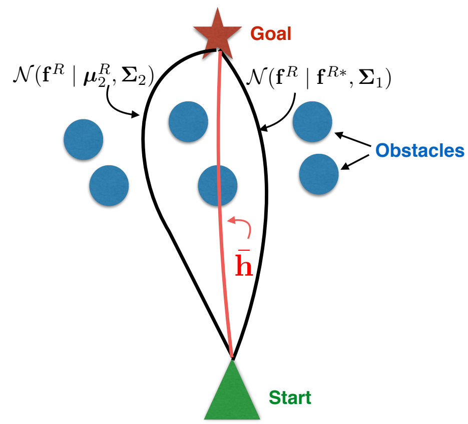

Figure 4: One global autonomy optima at $\mathbf{f}^{R*}$ and a safe but suboptimal autonomy mode at $\boldsymbol{\mu}^R_2$ through some obstacle field (additional autonomous modes exist but we leave them off for clarity). The operator's unimodal predicted trajectory at $\bar{\mathbf{h}}$ is *unsafe*. Covariance functions removed for clarity.

In Figure 4, we illustrate a very similar situation, except this time, the operator has chosen an *unsafe* trajectory. For linear trajectory blending, $u^s_{LTB} \propto K_h \bar{\mathbf{h}} + K_R \mathbf{f}^{R*}$, and so we must choose our heuristics such that $K_R \gg K_h$ in order to avoid collision—that is, we must insert logic that overrides the operator when the operator makes unsafe decisions.

In contrast, consider

$$
\begin{aligned}
&\psi_h(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) p(\mathbf{f}^R, \mathbf{f} \mid \bar{\mathbf{z}}_{1:t}) \\
&= \psi_h(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) \times \\
&[\beta_1 \mathcal{N}(\mathbf{f}^R \mid \mathbf{f}^{R*}, \mathbf{\Sigma}_1) + \beta_2 \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}^R_2, \mathbf{\Sigma}_2)] \\
&= \psi_h(\mathbf{h}, \mathbf{f}^R) \mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}_h, \mathbf{\Sigma}_h) \times \\
&[\beta_1 \mathcal{N}(\mathbf{f}^R \mid \mathbf{f}^{R*}, \mathbf{\Sigma}_1) + \beta_2 \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}^R_2, \mathbf{\Sigma}_2)]
\end{aligned}
$$

where we have maintained both modes and also maintained a unimodal distribution over $\mathbf{h}$ (rather than choosing the delta approximation). In this situation, β_1 is close to β_2, and the difference between the operator mean and either of the autonomy means are nearly the same, so Z_1 is close to Z_2. However, both $\mathbf{\Sigma}_1$ and $\mathbf{\Sigma}_2$ are both fairly narrow (otherwise, they are not safe modes), and because there is flexibility in $p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$, the MAP value is close to either $\mathbf{f}^{R*}$ or $\boldsymbol{\mu}^R_2$. Because the operator is treated probabilistically, there is no need to employ heuristics to detect poor operator choices: the autonomy assists the operator by blending near the tails of $p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$.

Finally, consider the situation in Figure 5. In this case, the operator has generated ambiguous data about how he wishes to move between the start and the goal, and so two modes carrying similar weights is the proper representation (although the mode centered at $\bar{\mathbf{h}}$ is taken to be near the MAP of $p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$). For linear blending, we have $u^s_{LTB} \propto K_h\bar{\mathbf{h}}+K_R\mathbf{f}^{R*}$ and so we end up with a situation very similar to that discussed in Figure 3—an autonomy and an operator that are needlessly in disagreement, and thus difficult to disambiguate with the heuristics K_h and K_R. To be fair, if $\bar{\mathbf{h}}$ and $\mathbf{f}^{R*}$ happen to lie close to one another, then linear blending provides the optimal solution—but for multimodal autonomous and operator distributions, such a situation is the exception rather than the rule (this exception requiring that the operator make globally optimal decisions).

For probabilistic shared control, the shared control is likely a trajectory near $\bar{\mathbf{h}}$ and $\boldsymbol{\mu}^R_2$—a solution that is both safe and preserves the operator's desires. However, depending on the weights α_m and β_n the solution may end up as a blend of $\boldsymbol{\mu}^h_2$ and $\boldsymbol{\mu}^R_3$—again, a solution that reflects the operator's desires and is still safe.

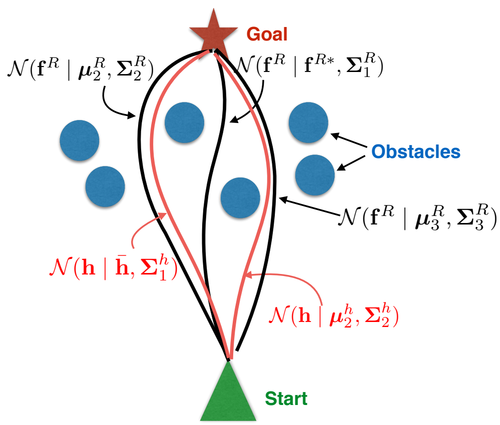

Figure 5: One global autonomy optima at $\mathbf{f}^{R*}$ and two safe but suboptimal autonomy modes at $\boldsymbol{\mu}^R_2$ and $\boldsymbol{\mu}^R_3$ through some obstacle field. The operator's bimodal predicted trajectory centered at $\bar{\mathbf{h}}$ and $\boldsymbol{\mu}^h_2$. Covariance functions removed for clarity.

The lesson of these three Figures is that for complex scenarios, linear blending can easily end up in situations that cannot be disambiguated with heuristics. This is a result of restricting the blend to the optimal autonomous decision and the most likely operator desires, which can easily be in conflict when nontrivial modes are not reasoned over. Probabilistic shared control, in contrast, maintains these suboptimal modes, and is thus able to reach a solution that is statistically principled, and is therefore safer and more likely in agreement with the operator.

7. APPROXIMATE INFERENCE FOR PROBABILISTIC SHARED CONTROL

We now describe how we use Monte Carlo sampling to estimate the MAP value of our model

$$p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} \mid \mathbf{z}_{1:t}) \propto \psi_h(\mathbf{h}, \mathbf{f}^R)\psi_f(\mathbf{f}^R, \mathbf{f})p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) \prod_{i=R}^{n_t} p(\mathbf{f}^i \mid \mathbf{z}^i_{1:t}).$$

First, we draw a joint sample $(\mathbf{h}, \mathbf{f}^R, \mathbf{f})_k = \{\mathbf{h}_k, \mathbf{f}^R_k, \mathbf{f}_k\}$ from each of the independent agent models:

- The operator model: $\mathbf{h}_k \sim p(\mathbf{h} \mid \mathbf{z}^h_{1:t})$
- The robot dynamics model: $\mathbf{f}^R_k \sim p(\mathbf{f}^R \mid \mathbf{z}^R_{1:t})$
- For each agent i in the crowd, draw $\mathbf{f}^i_k \sim p(\mathbf{f}^i \mid \mathbf{z}^i_{1:t})$.

We say that $(\mathbf{h}, \mathbf{f}^R, \mathbf{f})_k \sim p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) \prod_{i=R}^{n_t} p(\mathbf{f}^i \mid \mathbf{z}^i_{1:t})$, and we have that

$$p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) \prod_{i=R}^{n_t} p(\mathbf{f}^i \mid \mathbf{z}^i_{1:t}) \approx \sum_{k=1}^{N_s} \bar{w}_k \delta[(\mathbf{h}, \mathbf{f}^R, \mathbf{f}) - (\mathbf{h}, \mathbf{f}^R, \mathbf{f})_k]$$

where

$$\bar{w}_k = p(\mathbf{h}_k = \mathbf{h} \mid \mathbf{z}^h_{1:t}) \prod_{i=R}^{n_t} p(\mathbf{f}^i_k = \mathbf{f}^i \mid \mathbf{z}^i_{1:t}).$$

We therefore have the approximation

$$\begin{aligned} &\psi_h(\mathbf{h}, \mathbf{f}^R)\psi_f(\mathbf{f}^R, \mathbf{f})p(\mathbf{h} \mid \mathbf{z}^h_{1:t}) \prod_{i=R}^{n_t} p(\mathbf{f}^i \mid \mathbf{z}^i_{1:t}) \\ &= \psi_h(\mathbf{h}, \mathbf{f}^R)\psi_f(\mathbf{f}^R, \mathbf{f}) \sum_{k=1}^{N_s} \bar{w}_k \delta[(\mathbf{h}, \mathbf{f}^R, \mathbf{f}) - (\mathbf{h}, \mathbf{f}^R, \mathbf{f})_k] \\ &= \sum_{k=1}^{N_s} \bar{w}_k \psi_h(\mathbf{h}_k, \mathbf{f}^R_k)\psi_f(\mathbf{f}^R_k, \mathbf{f}_k)\delta[(\mathbf{h}, \mathbf{f}^R, \mathbf{f}) - (\mathbf{h}, \mathbf{f}^R, \mathbf{f})_k]. \end{aligned} \tag{7.1}$$

In particular, we take the MAP value $(\mathbf{h}, \mathbf{f}^R, \mathbf{f})^*$ to be the sample with the largest weight $\bar{w}_k\psi_h(\mathbf{h}_k, \mathbf{f}^R_k)\psi_f(\mathbf{f}^R_k, \mathbf{f}_k)$. Importantly, this approximate inference technique converges to the distribution as $N_s \to \infty$.

We point out that Gaussian sums and Monte Carlo inference are closely related, and so the arguments in Section 6 are applicable here. Specifically, the convergence criteria of Gaussian sums requires the covariance matrices of the mixture components to go to zero—that is, each mixture component becomes a sample. Furthermore, Monte Carlo sampling is well suited to this kind of inference, as was validated in the real world implementation in [25]. For practical reasons, the Gaussian sum approximation may very well be better suited to this problem, and we intend on exploring different approximate inference techniques in future papers. In this paper, however, Monte Carlo sampling has the appropriate properties to demonstrate our results.

8. CONCLUSIONS

In this paper, we presented an alternative formalism for shared control. We showed how the state of the art in shared control, linear blending, and some relevant generalizations (linear trajectory blending) are both special cases of our probabilistic approach. Further, we showed that linear blending is prone to statistical inconsistencies, but that the probabilistic approach is not. Finally, we explored some illustrative scenarios showing how linear blending is inappropriate for all but the most benign of circumstances, while probabilistic shared control is able to properly model both the operator and the autonomy in complex, dynamic environments.

9. REFERENCES

[1] D. Alspach and H. Sorenson. Nonlinear bayesian estimation using gaussian sum approximations. *IEEE Transactions on Automatic Control*, 1972.

[2] A. Bejczey. Recent developments in advanced teleoperation at jpl. In *IEEE International Conference on Advanced Robotics*, 1991.

[3] C. Bishop. *Pattern Recognition and Machine Learning.* Springer Science+Business Media, LLC, New York, NY, 2006.

[4] J. Burke and R. Murphy. From remote tool to shared roles. In *IEEE Robotics and Automation Magazine*, 2008.

[5] T. Carlson and Y. Demiris. Collaborative control for a robotic wheelchair: Collaborative control for a robotic wheelchair: Evaluation of performance, attention and workload. *IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS—PART B: CYBERNETICS*, 2012.

[6] J. Crandall and M. Goodrich. Characterizing efficiency of human robot interaction: A case study of shared-control teleoperation. In *IEEE International Conference on Intelligent Robots and Systems*, 2002.

[7] E. Demeester, A. Huntemann, E. V. Poorten, and J. D. Schutter. Ml, map and greedy pomdp shared control: comparison of wheelchair navigation assistance for switch interfaces. In *International Symposium on Robotics*, 2012.

[8] E. Demeester, A. Huntemann, D. Vanhooydonck, F. Vanacker, H. V. Brussel, and M. Nuttin. User-adapted plan recognition and user-adapted shared control: A bayesian approach to semi-autonomous wheelchair driving. In *Autonomous Robots*, 2008.

[9] A. Dragan and S. Srinivasa. Formalizing assistive teleoperation. *Robotics: Science and Systems*, 2012.

[10] A. Dragan and S. Srinivasa. A policy blending formalism for shared control. *International Journal of Robotics Research*, 2013.

[11] D. Fox, W. Burgard, and S. Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics and Automation Magazine*, 1997.

[12] J. Fu and U. Topcu. Probably approximately correct mdp learning and control with temporal logic constraints. In *Robotics: Science and Systems*, 2014.

[13] S. Gulati, C. Jhurani, B. Kuipers, and R. Longoria. A framework for planning comfortable and customizable motion of an assistive mobile robot. In *IEEE International Conference on Intelligent Robots and Systems*, 2009.

[14] B. Hardin and M. Goodrich. On using mixed-initiative control: A perspective for managing large-scale robotic teams. In *ACM/IEEE International Conference on Human-Robot Interaction*, 2009.

[15] A. Huntemann, E. Demeester, and E. V. Poorten. Probabilistic approach to recognize local navigation plans by fusing past driving information with a personalized user model. In *IEEE International Conference on Robotics and Automation*, 2013.

[16] P. Inigo-Blasco, F. D. del Rio, S. V. Diaz, and D. C. Muniz. The shared control dynamic window approach for non-holonomic semi-autonomous robots. In *International Symposium on Robotics/Robotik*, 2014.

[17] W. B. Knox and P. Stone. Tamer: Training an agent manually via evaluative reinforcement. In *IEEE International Conference on Development and Learning*, 2008.

[18] A. Lopes, U. Nunes, and L. Vaz. Assisted navigation based on shared-control, using discrete and sparse human-machine interfaces. In *IEEE Engineering in Medicine and Biology Society*, 2010.

[19] A. M. Morison, D. D. Woods, and T. Murphy. Human-robot interaction as extending human perception to new scales. In *Handbook of Applied Perception.* Cambridge University Press, in press.

[20] G. Peinado, C. Urdiales, J. Peula, M. Fdez-Carmona, R. Annicchiarico, F. Sandoval, and C. Caltagirone. Navigation skills based profiling for collaborative wheelchair control. In *IEEE International Conference on Robotics and Automation*, 2011.

[21] A. Poncela, C. Urdiales, E. Perez, and F. Sandoval. A new efficiency-weighted strategy for continuous human/robot cooperation in navigation. *IEEE Transaction on Systems, Man, and Cybernetics—Part A: Systems and Humans*, 2009.

[22] S. Soatto. *Steps Toward a Theory of Visual Information.* http://arxiv.org/abs/1110.2053, Technical Report UCLA-CSD100028, September 13, 2010 2010.

[23] P. Trautman. *Robot Navigation in Dense Crowds: Statistical Models and Experimental Studies of Human Robot Cooperation.* PhD thesis, California Institute of Technology, 2013.

[24] P. Trautman and A. Krause. Unfreezing the robot: Navigation in dense interacting crowds. In *IEEE International Conference on Intelligent Robots and Systems*, 2010.

[25] P. Trautman, J. Ma, R. Murray, and A. Krause. Robot navigation in dense crowds: the case for cooperation. In *IEEE International Conference on Robotics and Automation*, 2013.

[26] I. Ulrich and J. Borenstein. Vfh+: Reliable obstacle avoidance for fast mobile robots. In *IEEE International Conference on Robotics and Automation*, 1998.

[27] C. Urdiales, E. Pérez, G. Peinado, M. Fdez-Carmona, J. Peula, R. Annicchiarico, F. Sandoval, and C. Caltagirone. On the construction of a skill-based wheelchair navigation profile. *IEEE TRANSACTIONS ON NEURAL SYSTEMS AND REHABILITATION ENGINEERING*, 2013.

[28] C. Urdiales, J. Peula, M. Fdez-Carmona, C. Barrué, E. Pérez, I. Sánchez-Tato, J. del Toro, F. Galluppi, U. Cortés, R. Annichiaricco, C. Caltagirone, and F. Sandoval. A new multi-criteria optimization strategy for shared control in wheelchair assisted navigation. *Autonomous Robots*, 2011.

[29] H. Wang and X. Liu. Adaptive shared control for a novel mobile assistive robot. *IEEE/ASME TRANSACTIONS ON MECHATRONICS*, 2014.

[30] M. Wang and J. Liu. Interactive control for internet-based mobile robot teleoperation. *Robotics and Autonomous Systems*, 2005.

[31] P. Welinder, S. Branson, S. Belongie, and P. Perona. The multidimensional wisdom of crowds. In *Neural Information Processing Systems*, 2010.

[32] H. Yu, M. Spenko, and S. Dubowsky. An adaptive shared control system for an intelligent mobility aid for the elderly. *Autonomous Robots*, 2003.