

Breaking the Human-Robot Deadlock: Surpassing Shared Control Performance Limits with Sparse Human-Robot Interaction

Pete Trautman, Galois, Inc.

Abstract—Human machine teaming has, for decades, been conceptualized as a *function allocation* (FA) or *levels of autonomy* (LOA) process: the human is suited for some tasks, while the machine is suitable for others, and as machines improve they take over duties previously assigned to humans. A wide variety of methods—including adaptive, adjustable, blended, supervisory and mixed initiative control, implemented discretely or continuously, as potential fields, as virtual fixture interfaces, or haptic interfaces—are derivatives of FA/LOA. We formalize FA/LOA (and all their derivatives) under a single mathematical formulation called *classical shared control* (CSC). Despite the widespread adoption of CSC, we prove that it fails to optimize human and robot agreement and intent if either the human or robot model displays “intention ambiguity” (e.g., the human’s intended goal is unclear or the robot finds multiple viable solutions). Practically, this suboptimality can manifest as unnecessary and unresolvable disagreement (an unnecessary deadlock). For instance, if the robot chooses to go left around an obstacle and the human chooses to go right, CSC only provides two solutions: freeze in place or collide with the obstacle (we provide a wide variety of failure examples in [52], <https://arxiv.org/abs/1611.09490>). We find that CSC suboptimality stems from arbitrating over *model samples*, rather than over *models*. Our key insight is thus to arbitrate over human and robot distributions; we prove this method optimizes human and robot agreement and intent and resolves deadlocking. Our key contribution is computationally efficient distribution arbitration: if the human and robot carry N_t^h and N_t^R “intentions,” the joint (naively) has $N_t^h N_t^R$ intentions. In our approach, deadlock solutions have vanishingly small coefficients and only $N_{\min} = \arg \min\{N_t^h, N_t^R\}$ non-zero coefficients remain: *our joint has fewer modes than the individual agent models*. We call our approach N_{\min} -sparse generalized shared control.

I. INTRODUCTION

Leveraging the complementary capabilities of human and machine has long been a goal of computer science. Typically, researchers focus on a specific aspect of this broad vision. For instance, [42] focuses on fusing human and machine “perception.” Likewise, attempts to blend human and machine “decision making” occur in the machine learning [60], [36], control theory [29], and human robot interaction literature [32]. A special case of shared decision making is shared control: fuse human and robot platform commands. Examples of shared control abound; assistive wheelchair technology [31], brain-computer interface telemanipulation [43], surgical robots [45], search and rescue [9], extraterrestrial robotics [6], and assistive automobile driving [3] are just a few examples. For any of these cases, shared control can be broken down into a human and robot modeling step, a human and robot prediction step, and an arbitration step over the human and robot predictions (see Section II and [23] for justification of this schema).

In this paper we focus on the performance limits that arise as a consequence of the “classical” shared control architecture

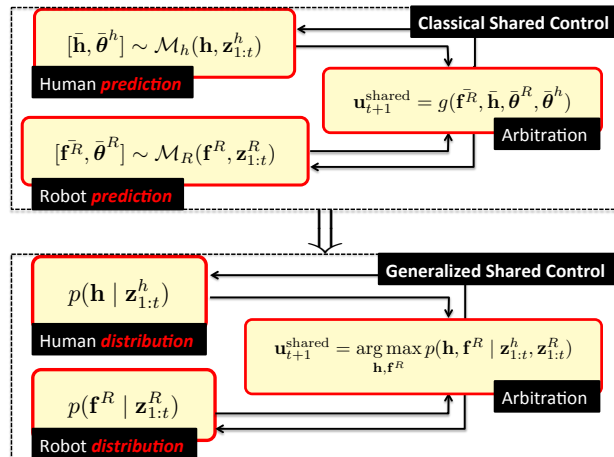


Fig. 1: Classical shared control (CSC) arbitrates robot and human *predictions*. We prove that no matter what predictions, models or arbitration function are used, CSC fails to optimize human-robot agreement and intent if intention ambiguity is present. Generalized shared control arbitrates *distributions*, optimizing human-robot agreement and intent, and is computationally efficient.

(top half of Figure 1). By identifying the cause of these classical performance limits—arbitrating over prediction samples—we gain insight into how to repair them. We thus introduce generalized shared control (GSC), where instead of arbitrating over *predictions* from the human and robot models, we arbitrate over the full human and robot *distributions* themselves (bottom half of Figure 1). Our approach is motivated by the following: distributions represented in a Gaussian process (GP) basis [49], [54], [55], [5], [30] explicitly capture multifaceted intentionality (which we call “intention ambiguity”) as well as an agent’s willingness to compromise within a given strategy (which we call “flexibility”). In Figure 2, we illustrate intention ambiguity and flexibility, and in Section III we define intention ambiguity and flexibility. Critically, finite human and robot samples fail to fully reflect intention ambiguity or flexibility, and so classical arbitration is unable to balance these nuanced attributes. For instance, suppose that we cannot determine if the human wishes to go to the left or right food station, and we represent this with a bimodal distribution: $p(\mathbf{h} | \mathbf{z}_{1:t}^h) = \sum_{i=L,R} w_i^h \mathcal{N}(\mathbf{h} | \boldsymbol{\mu}_i^h, \boldsymbol{\Sigma}_i^h)$. If $w_L^h > w_R^h$ then the most likely route for the human is left. Arbitration over this statistic can be misleading since the R mode is ignored.

Our key insight, then, is to simultaneously optimize the ambiguous and flexible intentions of the human and robot. Our key contribution is to provide a computational mechanism so that distribution arbitration is both well defined and computationally tractable. By formulating shared control as a special case of joint decision making (simultaneously optimizing human and robot agreement and intent), we can

leverage “negotiation” pruning: if a joint action is unresolvable (e.g., the human prefers left and the robot prefers right), then it is removed from arbitration. However, pruning is judicious: it enables arbitration over complementary joint human-robot actions for visiting both left and right. In general, the human and robot models carry N_t^h and N_t^R time varying possibilities in a GP basis. Naively, the human-robot joint distribution will thus have $N_t^h N_t^R$ modes. Pruning, however, eliminates many joint actions: if $N_{\min} = \arg \min\{N_t^h, N_t^R\}$, the joint distribution has only N_{\min} nontrivial components in the GP basis. Pruning in a statistically principled manner is a primary conceptual motivation for using a GP basis: sparsity quantifies the balance of model complexity against model fidelity. We call our approach sparse-GSC.

II. RELATED WORK

The *levels of autonomy* (LOA) paradigm [50] and the closely related *function allocation* (FA) approach [27], [46] are the conceptual bedrock of shared control: many, if not all, approaches can be traced back to these early theories of human-machine complementarity. According to [32], LOA/FA is typically implemented in three ways¹: adaptive autonomy, where the machine determines the level of assistance [39]; adjustable autonomy, where the human determines the level of assistance (sometimes called supervisory control [41]); or mixed-initiative control, where the human and machine deliberate to determine the level of assistance [44]. Sometimes, the term sliding autonomy is used to indicate one or a mixture of the three above [21]. Sliding scale autonomy [20] relaxes discrete LOA to a continuum of assistance. Additionally, prevailing shared control approaches such as potential field methods [2], virtual fixture based methods [25], [1], and haptic shared control [14] are ultimately rooted in LOA/FA. For example, in [28], the optimal controller of a cart-pendulum task was used in conjunction with a haptic interface to filter suboptimal user input (the machine decided when to give assistance). Some researchers, such as [17], [34], [18], infer user intent as a maximum likelihood, maximum a-posteriori, or *partially observable Markov decision process* estimation problem; in turn, shared control is computed as a reward function over the optimal user action and optimal robot action ($g(\bar{\mathbf{f}}^R, \bar{\mathbf{h}}, \bar{\boldsymbol{\theta}}^R, \bar{\boldsymbol{\theta}}^h)$ in Figure 1 could be a reward function). We observe that all of the above approaches are instances of classical shared control as described in Figure 1: the arbitration function $g(\mathbf{f}^R, \mathbf{h}, \boldsymbol{\theta}^R, \boldsymbol{\theta}^h)$ could be implemented as adaptive, adjustable, supervisory or mixed initiative control (by allowing the arguments $[\mathbf{f}^R, \mathbf{h}, \boldsymbol{\theta}^R, \boldsymbol{\theta}^h]$ to mediate who controls what and when); in a discrete or continuous fashion; or as a potential field, virtual fixture interface driver, or haptic interface driver. There is no restriction on the form of $g(\mathbf{f}^R, \mathbf{h}, \boldsymbol{\theta}^R, \boldsymbol{\theta}^h)$ other than that it accepts prediction samples, rather than distributions, as arguments. At its most abstract, any LOA or FA approach is just some function that actuates the human

¹*Caveat emptor*: the community has not reached consensus on terminology or definitions here; we choose [32] because it is a plausible representative.

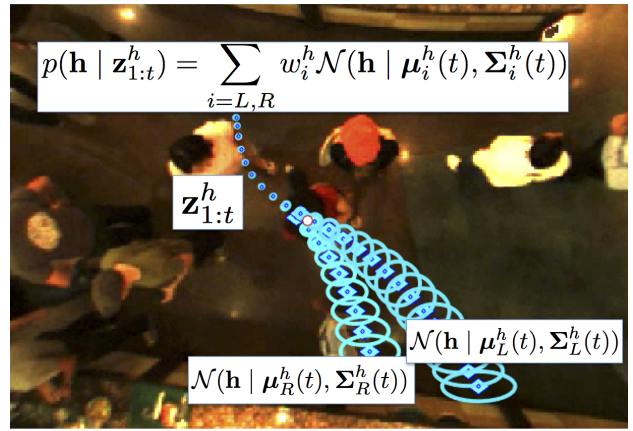


Fig. 2: Human exhibits ambiguity: $\mathcal{N}(\mathbf{h} | \boldsymbol{\mu}_R^h, \boldsymbol{\Sigma}_R^h)$ and $\mathcal{N}(\mathbf{h} | \boldsymbol{\mu}_L^h, \boldsymbol{\Sigma}_L^h)$ model intention to go in direction $\boldsymbol{\mu}_R^h$ or $\boldsymbol{\mu}_L^h$ (with preferences w_R^h and w_L^h). Within a mode, human has flexibility $\boldsymbol{\Sigma}_R^h$ about $\boldsymbol{\mu}_R^h$ and $\boldsymbol{\Sigma}_L^h$ about $\boldsymbol{\mu}_L^h$.

command, the machine command, or some combination of human and machine command. In Figure 1, we generically name that function $g(\mathbf{f}^R, \mathbf{h}, \boldsymbol{\theta}^R, \boldsymbol{\theta}^h)$.

LOA/FA is primarily appealing because it is straightforward: the tasks of human and machine are cleanly delineated into separate categories. However, for meaningful teaming to occur, the agents cannot be limited by dichotomization, since the complementary capabilities of human and machine typically overlap. As [16] puts it, the “interface between human and machine [in LOA/FA] is a trivial ‘you do this, I do that’ barter”. Importantly, GSC loosens these strict boundaries by mixing the human and robot distributions. We note that cognitive systems engineering [11], [16], [33], [15], team cognition theory [12], [40], [19], [13], [12] and human-agent-robot teamwork [7], [8] make important conceptual progress towards understanding such a nuanced human-machine negotiation. However, these approaches lack the mathematical precision, performance guarantees, and computational tractability provided by sparse-GSC. Importantly, the implementation of a cognitive systems engineering, team cognition, or human-agent-robot teamwork approach are left to the discretion of the engineer, because the concepts are not reduced to unambiguous mathematical guidance.

Similar to [23], we call special attention to the linear blending framework (a classical shared control framework):

$$\begin{aligned} u_{LB}^s(t+1) &= K_R u_t^R + K_h u_t^h \\ &= g(\bar{\mathbf{f}}^R = u_t^R, \bar{\mathbf{h}} = u_t^h, \bar{\boldsymbol{\theta}}^R = K_R, \bar{\boldsymbol{\theta}}^h = K_h), \end{aligned} \quad (\text{II.1})$$

where, at time t , $u_{LB}^s(t+1)$ is the linearly blended shared control command, u_t^h is the human input (joystick deflections, keyboard inputs, etc.), u_t^R is the autonomy command, and K_h, K_R are the operator and robot arbitration coefficients which range between 0 and 1 and are time varying ($\boldsymbol{\theta}^R, \boldsymbol{\theta}^h$ are the parameters of the CSC arbitration function). This approach can be used to describe a large swath of shared control approaches, as argued in [23]. For instance, any type of switching control (arguably the most broadly deployed type of shared control; for example, anti-lock braking systems in cars [4] and autopilot in commercial avionics [26] are

switching control systems), where either the human or machine is in complete control of the platform at time t , is just linear blending where $K_h, K_R = 0, 1$. Linear blending has enjoyed wide adoption in the assistive wheelchair community [10], [58], [38], [57], [61], [47], [56], [35]. In [48], [59], an even broader adoption of linear blending is advocated. These three examples are a small sample of the use cases of linear blending; a more comprehensive bibliographic accounting is found in [23], [22]. Importantly, [23] argues convincingly that most of shared control can be formalized as a block diagram very similar to the classical shared control block diagram in Figure 1. Also, [24] makes a similar formalization observation.

However, linear blending has fundamental performance limits: in [51] linear blending was proven suboptimal with respect to safety, efficiency and operator-robot agreement if intention ambiguity is present ([52] describes a number of concrete scenarios where the suboptimality causes failure states). For linear blending systems to be deployed safely, practitioners often use *ad-hoc* safeguard mechanisms. An example involves maneuvering around a static obstacle. If the robot chooses left and the human chooses right, linear blending can average the two inputs into a collision. Collision safeguarding iteratively discards unsafe solutions until a safe solution is found. However, as explained in Section 2.2 of [?], these safeguards have unintended consequences. In congested environments, collision safeguards can cause the shared control to freeze even if safe forward trajectories exist.

These shortcomings of linear blending beg the question: if we sample robot and human predictions $\mathbf{f}^R, \bar{\mathbf{h}}$ and robot and human parameters $\bar{\boldsymbol{\theta}}^R, \bar{\boldsymbol{\theta}}^h$ from any robot and human model, does some arbitration protocol $g(\bar{\mathbf{f}}^R, \bar{\mathbf{h}}, \bar{\boldsymbol{\theta}}^R, \bar{\boldsymbol{\theta}}^h)$ exist that optimizes human and robot agreement and intent? As we prove in the following, no matter how the human and robot models or arbitration function $g(\mathbf{f}^R, \mathbf{h}, \boldsymbol{\theta}^R, \boldsymbol{\theta}^h)$ are chosen, human and robot agreement and intent is not jointly satisfied if intention ambiguity is present. Additionally, we prove that sparse-GSC efficiently optimizes human and robot agreement and intent in the presence of intention ambiguity and flexibility.

III. TERMINOLOGY

We treat operator inputs u_t^h as *measurements* \mathbf{z}_t^h ($\mathbf{z}_t^h \equiv u_t^h$) of the operator *trajectory*, $\mathbf{h}: t \in \mathbb{R} \rightarrow \mathcal{X}$, where \mathcal{X} is the joint action space and \mathbf{h} is governed by $p(\mathbf{h} \mid \mathbf{z}_{1:t}^h)$. We collect measurements $\mathbf{z}_{1:t}^R$ of the robot trajectory \mathbf{f}^R , which is governed by $p(\mathbf{f}^R \mid \mathbf{z}_{1:t}^R)$. We do not assume that every measurement of $\mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^h$ in $1:t$ is present. Technically, then, both \mathbf{h} and \mathbf{f}^R are stochastic processes. Our human and robot models are represented in a GP basis (Figure 2):

$$\begin{aligned} p(\mathbf{h} \mid \mathbf{z}_{1:t}^h) &= \sum_{i=1}^{N_t^h} w_i^h \mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}_i^h, \boldsymbol{\Sigma}_i^h), \\ p(\mathbf{f}^R \mid \mathbf{z}_{1:t}^R) &= \sum_{j=1}^{N_t^R} w_j^R \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}_j^R, \boldsymbol{\Sigma}_j^R). \end{aligned} \quad (\text{III.1})$$

$\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}(t), \boldsymbol{\Sigma}(t))$ is a GP with mean function $\boldsymbol{\mu}(t)$ and covariance function $\boldsymbol{\Sigma}(t)$. We suppress time in the mean and

covariance functions for clarity. That is, in the following any $\boldsymbol{\mu} \equiv \boldsymbol{\mu}(t)$ and $\boldsymbol{\Sigma} \equiv \boldsymbol{\Sigma}(t)$.

Multi-modal models are common in real world applications. For instance, human goal inference from trajectory data often results in uncertainty over multiple goals; these possible multiple destinations can instigate N_t^h modes. For robots, congested environments often result in many N_t^R local maxima. We point the reader to [52] for a thorough discussion of how multi-modality can arise in human and robot models.

Definition 1 (Intent, intention ambiguity). We call $\boldsymbol{\mu}_i^h, \boldsymbol{\mu}_j^R$ the human and robot intentions. If $N_t^h > 1$ or $N_t^R > 1$, we say that intention ambiguity is present. We measure intention ambiguity through the weights w_i^h, w_j^R . For example, if one weight is very large, intention ambiguity is small. Conversely, if the weights are all equal, intention ambiguity is large.

Definition 2 (Flexibility). Flexibility is the willingness of an agent to compromise about an intention $\boldsymbol{\mu}_i^h$ or $\boldsymbol{\mu}_j^R$. Mathematically, the flexibility of intent $\boldsymbol{\mu}_i^h$ or $\boldsymbol{\mu}_j^R$ is $\boldsymbol{\Sigma}_i^h$ or $\boldsymbol{\Sigma}_j^R$.

Our definition of flexibility is motivated by the following: suppose an agent is unimodal with model $\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})$. If the agent indicates intent $\boldsymbol{\mu}$ strongly (by providing substantial data supporting $\boldsymbol{\mu}$) then $\boldsymbol{\Sigma}$ will be small. In other words, the agent is $\boldsymbol{\Sigma}$ unwilling to compromise on $\boldsymbol{\mu}$. Conversely, if the agent has not provided a strong signal supporting $\boldsymbol{\mu}$ then the agent is $\boldsymbol{\Sigma}$ flexible about intent $\boldsymbol{\mu}$.

Definition 3 (Human-robot agreement). We use human intent and flexibility $\boldsymbol{\mu}^h, \boldsymbol{\Sigma}^h$ and robot intent and flexibility $\boldsymbol{\mu}^R, \boldsymbol{\Sigma}^R$ to discuss human-robot agreement. In particular

$$Z_{i,j}^{-1} = \exp \left[-\frac{1}{2} (\boldsymbol{\mu}_i^h - \boldsymbol{\mu}_j^R)^\top (\boldsymbol{\Sigma}_i^h + \boldsymbol{\Sigma}_j^R)^{-1} (\boldsymbol{\mu}_i^h - \boldsymbol{\mu}_j^R) \right]$$

quantifies human-robot agreement, as derived in Equation IV.5. Large $Z_{i,j}^{-1}$ indicates agreement between modes i, j , while small $Z_{i,j}^{-1}$ indicates disagreement between modes i, j .

Definition 4 (CSC). Classical shared control has three components: a human model $\mathcal{M}_h(\mathbf{h}, \mathbf{z}_{1:t}^h)$, a robot model $\mathcal{M}_R(\mathbf{f}^R, \mathbf{z}_{1:t}^R)$, and an arbitration function $g(\bar{\mathbf{f}}^R, \bar{\mathbf{h}}, \bar{\boldsymbol{\theta}}^R, \bar{\boldsymbol{\theta}}^h): [\bar{\mathbf{f}}^R, \bar{\mathbf{h}}, \bar{\boldsymbol{\theta}}^R, \bar{\boldsymbol{\theta}}^h] \rightarrow \mathcal{X}$, where \mathcal{X} is the joint action space. At time t , a finite set of samples

$$\begin{aligned} [\bar{\mathbf{h}}, \bar{\boldsymbol{\theta}}^h] &\sim \mathcal{M}_h(\mathbf{h}, \mathbf{z}_{1:t}^h) \\ [\bar{\mathbf{f}}^R, \bar{\boldsymbol{\theta}}^R] &\sim \mathcal{M}_R(\mathbf{f}^R, \mathbf{z}_{1:t}^R). \end{aligned}$$

is drawn from the models (where $\boldsymbol{\theta}^h, \boldsymbol{\theta}^R$ are parameters of the variables \mathbf{h}, \mathbf{f}^R). The classical shared control $u_{CSC}^s(t+1) = g(\bar{\mathbf{f}}^R, \bar{\mathbf{h}}, \bar{\boldsymbol{\theta}}^R, \bar{\boldsymbol{\theta}}^h)$ is then executed; the process repeats at $t+1$.

Example Let $\mathcal{M}_h(\mathbf{h}, \mathbf{z}_{1:t}^h) = \mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}^h, \boldsymbol{\Sigma}^h)$ and $\mathcal{M}_R(\mathbf{f}^R, \mathbf{z}_{1:t}^R) = \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}^R, \boldsymbol{\Sigma}^R)$. If $\bar{\mathbf{h}} = \boldsymbol{\mu}^h = \arg \max_{\mathbf{h}} \mathcal{N}(\mathbf{h} \mid \boldsymbol{\mu}^h, \boldsymbol{\Sigma}^h)$ and $\bar{\mathbf{f}}^R = \boldsymbol{\mu}^R = \arg \max_{\mathbf{f}^R} \mathcal{N}(\mathbf{f}^R \mid \boldsymbol{\mu}^R, \boldsymbol{\Sigma}^R)$, then the only statistically valid assignment is $[\bar{\mathbf{f}}^R, \bar{\mathbf{h}}, \bar{\boldsymbol{\theta}}^R, \bar{\boldsymbol{\theta}}^h] = [\boldsymbol{\mu}^R, \boldsymbol{\mu}^h, \boldsymbol{\Sigma}^R, \boldsymbol{\Sigma}^R]$. The arbitration function $g(\boldsymbol{\mu}^R, \boldsymbol{\mu}^h, \boldsymbol{\Sigma}^R, \boldsymbol{\Sigma}^R)$ produces the shared control $u_{CSC}^s(t+1)$.

The arbitration function is often chosen to be

$$g(\boldsymbol{\mu}^R, \boldsymbol{\mu}^h, \boldsymbol{\Sigma}^R, \boldsymbol{\Sigma}^h) = (\boldsymbol{\Sigma}^h)^{-1} \boldsymbol{\mu}^h + (\boldsymbol{\Sigma}^R)^{-1} \boldsymbol{\mu}^R,$$

which is just linear blending [23]. See Lemma IV.2.

Definition 5 (GSC). *Generalized shared control has three components: a human distribution $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$, a robot distribution $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ and an arbitration operator $\arg \max_{\mathbf{h}, \mathbf{f}^R}$. The human and robot distributions are composed into a joint $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ and the generalized shared control $u_{GSC}^s(t+1) = \arg \max_{\mathbf{h}, \mathbf{f}^R} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ is then executed ($\mathbf{z}_{1:t} = [\mathbf{z}_{1:t}^h, \mathbf{z}_{1:t}^R]$). The process repeats at $t+1$.*

Example If $p(\mathbf{h} | \mathbf{z}_{1:t}^h) = \mathcal{N}(\mathbf{h} | \boldsymbol{\mu}^h, \boldsymbol{\Sigma}^h)$ and $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) = \mathcal{N}(\mathbf{f}^R | \boldsymbol{\mu}^R, \boldsymbol{\Sigma}^R)$ and if we choose (see Section IV-C)

$$p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) = \delta(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} | \mathbf{z}_{1:t}^h) p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$$

then $u_{GSC}^s(t+1) = (\boldsymbol{\Sigma}^h)^{-1} \boldsymbol{\mu}^h + (\boldsymbol{\Sigma}^R)^{-1} \boldsymbol{\mu}^R$ (Lemma IV.2).

IV. HUMAN ROBOT JOINT DECISION MAKING

In this section, we discuss the origins of sparse-GSC to motivate our formulation. We start with an overview of how robot navigation in human crowds is best understood as a *joint decision making problem*, rather than as a traditional path planning problem—that is, optimization is best conducted over the *distributions* of the robot and the crowd, rather than over *samples* of the robot and crowd. We then discuss an extension of the navigation framework to shared control. If shared control is posed as a joint decision making problem (under a mild set of assumptions), we prove that $\delta(\mathbf{h}, \mathbf{f}^R)$ is the only statistically valid interaction function. We prove that the human robot joint distribution is sparse in GP basis.

A. Robot crowd navigation as a joint decision making problem

In [53], robot navigation in dense human crowds was explored. Typically, this is posed as a path planning problem. The authors took the perspective that it was instead a joint decision making problem: how should the robot move, in concert with the humans around it, so that the objectives of each participant were simultaneously optimized? As was proved mathematically, unless crowd navigation is treated as joint decision making, the robot suffers the *freezing robot problem* (FRP). The FRP has been experimentally observed in independent studies [54], [37]: beyond 0.55 people/m², the robot was unable to move, and a 3x improvement in safety was observed when crowd navigation was treated as joint decision making instead of as path planning.

The high level mathematics of this approach explains how crowd navigation is best understood as joint decision making and motivates the formulation of GSC. First, the joint predictive distribution over the robot model \mathbf{f}^R and the crowd model $\mathbf{f} = [\mathbf{f}^1, \dots, \mathbf{f}^{n_t}]$ was formulated (\mathbf{f}^i is a human in the crowd and n_t is the number of people), subject to data about the robot and crowd $[\mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f]$, represented as $p(\mathbf{f}^R, \mathbf{f} | \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f)$. The robot’s next action $u^s(t+1)$ —what the robot is predicted to do




Fig. 3: Illustration of robot navigation in crowds as joint decision making. The robot and crowd distributions are arbitrated over in the manner of GSC, thus producing the optimal joint decision for the robot and the crowd.

in order to integrate into the crowd—is then clear (Figure 3):

$$[\mathbf{f}^R, \mathbf{f}]^* = \arg \max_{\mathbf{f}^R, \mathbf{f}} p(\mathbf{f}^R, \mathbf{f} | \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f) \quad (\text{IV.1})$$

$$u^s(t+1) = \mathbf{f}^{R*}(t+1).$$

We label the robot action u^s precisely because shared control is occurring: $\psi(\mathbf{f}^R, \mathbf{f})$ models the “interacting” decision strategies of $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ and $p(\mathbf{f} | \mathbf{z}_{1:t}^f)$ and $\arg \max_{\mathbf{f}^R, \mathbf{f}}$ captures the optimal robot-crowd decision.

B. Shared control as a joint decision making problem

In [51], the work in [53], [54] was extended to the case of a human sharing control with a robot for navigation through human crowds (e.g., shared control wheelchairs in crowds). Instead of choosing the standard shared control paradigm, the problem was treated as one of optimal joint decision making: how can human, robot and crowd objectives be simultaneously optimized? The formulation extends Equation IV.1:

$$[\mathbf{h}, \mathbf{f}^R, \mathbf{f}]^* = \arg \max_{\mathbf{h}, \mathbf{f}^R, \mathbf{f}} p(\mathbf{h}, \mathbf{f}^R, \mathbf{f} | \mathbf{z}_{1:t}^h, \mathbf{z}_{1:t}^R, \mathbf{z}_{1:t}^f)$$

In this paper, we are not interested in the interplay between the environment \mathbf{f} and the human and robot. Instead, we specifically explore arbitration between the intention ambiguous and flexible human and robot. Dropping the notation for \mathbf{f} and letting $\mathbf{z}_{1:t} = [\mathbf{z}_{1:t}^h, \mathbf{z}_{1:t}^R]$, our GSC equation is:

$$[\mathbf{h}, \mathbf{f}^R]^* = \arg \max_{\mathbf{h}, \mathbf{f}^R} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) \quad (\text{IV.2})$$

Unlike Equation IV.1, it is unclear which $\arg \max$ component to choose for the GSC shared control u_{GSC}^s : \mathbf{h}^* or \mathbf{f}^{R*} ? Typically, this question is resolved using heuristic arguments: follow an LOA/FA approach, such as “the human retains control under these circumstances, and the machine retains control under other circumstances”. The LOA/FA paradigm is insufficiently nuanced; we search instead for a mathematical approach inspired by the teaming literature [16], [12], [7]. To do so, we show how the question— \mathbf{h}^* or \mathbf{f}^{R*} ?—is resolved by requiring that $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ respects the intent and flexibility information contained in $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ and $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$.

C. Valid human-robot interaction functions

To gain insight about whether \mathbf{h}^* or \mathbf{f}^{R*} is our shared control, we begin with the decomposition (similar to [51])

$$\begin{aligned} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}^h, \mathbf{z}_{1:t}^R) &= p(\mathbf{f}^R | \mathbf{h}, \mathbf{z}_{1:t}^R) p(\mathbf{h} | \mathbf{z}_{1:t}^h) \\ &= \psi(\mathbf{h}, \mathbf{f}^R) p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) p(\mathbf{h} | \mathbf{z}_{1:t}^h), \end{aligned} \quad (\text{IV.3})$$

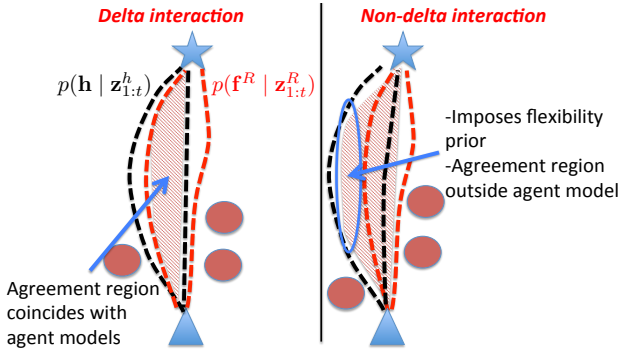


Fig. 4: Human model $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ in black and robot model $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ in red (dotted lines indicate one standard deviation); $\delta(\mathbf{h}, \mathbf{f}^R)$ limits “agreement” region to overlap between human and robot models. Any $\psi(\mathbf{h}, \mathbf{f}^R) \neq \delta(\mathbf{h}, \mathbf{f}^R)$ acts as a flexibility prior on $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ and $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$.

where we use the chain rule of probability, and the assumption that $p(\mathbf{f}^R | \mathbf{h}, \mathbf{z}_{1:t}^R) = \psi(\mathbf{h}, \mathbf{f}^R)p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$. What should the interaction function $\psi(\mathbf{h}, \mathbf{f}^R)$ be? First, let $\mathbf{h} - \mathbf{f}^R \equiv [\mathbf{h}(1), \dots, \mathbf{h}(T)] - [\mathbf{f}^R(1), \dots, \mathbf{f}^R(T)]$, where $\mathbf{h}(t), \mathbf{f}^R(t) \in \mathcal{X}$ and we discretize the functions \mathbf{h}, \mathbf{f}^R by $1 : T$ to define subtraction. A plausible choice for the interaction function is $\psi(\mathbf{h}, \mathbf{f}^R) = \exp(-\frac{1}{2\gamma}(\mathbf{h} - \mathbf{f}^R)^2)$ since this function attracts \mathbf{h} and \mathbf{f}^R . Recall that $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ and $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ encode all available intention and flexibility information (Equation III.1). Unfortunately, if we choose $\gamma > 0$ we impose flexibility on $p(\mathbf{h} | \mathbf{z}_{1:t}^h), p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$. In Figure 4, we illustrate how the “agreement region” of $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ is enlarged beyond what either $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ or $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ allow. Thus, we let $\gamma \rightarrow 0$:

$$\delta(\mathbf{h}, \mathbf{f}^R) \equiv \lim_{\gamma \rightarrow 0} \left[\exp\left(-\frac{1}{2\gamma}(\mathbf{h} - \mathbf{f}^R)^2\right) \right] = \begin{cases} 1 & \text{if } \mathbf{h} = \mathbf{f}^R, \\ 0 & \text{if } \mathbf{h} \neq \mathbf{f}^R. \end{cases}$$

The purpose of an interaction function is to place the two agents in a common frame; however, we cannot pollute $p(\mathbf{h} | \mathbf{z}_{1:t}^h), p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ with extra information. Relatedly, $\delta(\mathbf{h}, \mathbf{f}^R)$ can be interpreted as a transform $\mathbf{h} \rightarrow \mathbf{f}^R$ or $\mathbf{f}^R \rightarrow \mathbf{h}$. We make this precise:

Theorem IV.1 (Delta only choice). *If $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ and $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ contain all flexibility information and if we decompose the joint as in Equation IV.3, then $\psi(\mathbf{h}, \mathbf{f}^R) = \delta(\mathbf{h}, \mathbf{f}^R)$.*

Proof: Let $\psi(\mathbf{h}, \mathbf{f}^R) \neq \delta(\mathbf{h}, \mathbf{f}^R)$ and suppose the human intends $\mathbf{x} \in \mathcal{X}$ with $\Sigma \rightarrow 0$. Then $p(\mathbf{h} | \mathbf{z}_{1:t}^h) = \delta(\mathbf{h}, \mathbf{x})$, and

$$\begin{aligned} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) &= \psi(\mathbf{h}, \mathbf{f}^R)\delta(\mathbf{h}, \mathbf{x})p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) \\ &= \psi(\mathbf{x}, \mathbf{f}^R)p(\mathbf{f}^R | \mathbf{z}_{1:t}^R). \end{aligned}$$

Since $\psi(\mathbf{h}, \mathbf{f}^R)$ has finite support, $\psi(\mathbf{x}, \mathbf{f}^R)$ could be large when $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ is small, thus inflating the flexibility of $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ about \mathbf{x} . Similarly, $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ could be large when $\psi(\mathbf{x}, \mathbf{f}^R)$ is small, thus diminishing the flexibility of $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ about \mathbf{x} . Choosing $\psi(\mathbf{h}, \mathbf{f}^R) \neq \delta(\mathbf{h}, \mathbf{f}^R)$ thus causes the joint distribution in Equation IV.3 to inaccurately reflect the flexibility of $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$. The same argument can be made about incorrectly inflating or diminishing the human’s flexibility about a robot action.

Conversely, let $\psi(\mathbf{h}, \mathbf{f}^R) = \delta(\mathbf{h}, \mathbf{f}^R)$ and suppose the human intends $\mathbf{x} \in \mathcal{X}$. Then $p(\mathbf{h} | \mathbf{z}_{1:t}^h) = \delta(\mathbf{h}, \mathbf{x})$, and

$$\begin{aligned} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) &= \delta(\mathbf{h}, \mathbf{f}^R)\delta(\mathbf{h}, \mathbf{x})p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) \\ &= \delta(\mathbf{x}, \mathbf{f}^R)p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) \\ &= p(\mathbf{f}^R = \mathbf{x} | \mathbf{z}_{1:t}^R). \end{aligned}$$

Thus, the robot’s flexibility about action \mathbf{x} is accurately reflected (the same argument can be made for the human). The choice $\psi(\mathbf{h}, \mathbf{f}^R) = \delta(\mathbf{h}, \mathbf{f}^R)$ thus respects the flexibility information contained in the human and robot distributions. ■

Since $\delta(\mathbf{h}, \mathbf{f}^R)\mathcal{N}(\mathbf{h} | \boldsymbol{\mu}_i^h, \boldsymbol{\Sigma}_i^h) = \mathcal{N}(\mathbf{f}^R | \boldsymbol{\mu}_i^h, \boldsymbol{\Sigma}_i^h)$, we have

$$\begin{aligned} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) &= \delta(\mathbf{h} - \mathbf{f}^R)p(\mathbf{h} | \mathbf{z}_{1:t}^h)p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) \\ &= \delta(\mathbf{h} - \mathbf{f}^R) \sum_{i=1}^{N_t^h} w_i^h \mathcal{N}(\mathbf{h} | \boldsymbol{\mu}_i^h, \boldsymbol{\Sigma}_i^h) \\ &\quad \times \sum_{j=1}^{N_t^R} w_j^R \mathcal{N}(\mathbf{f}^R | \boldsymbol{\mu}_j^R, \boldsymbol{\Sigma}_j^R) \\ &= \sum_{i=1}^{N_t^h} \sum_{j=1}^{N_t^R} Z_{i,j}^{-1} w_i^h w_j^R \mathcal{N}(\mathbf{f}^R | \boldsymbol{\mu}_{i,j}, \boldsymbol{\Sigma}_{i,j}) \end{aligned} \quad (\text{IV.4})$$

where

$$\begin{aligned} Z_{i,j}^{-1} &= \exp\left[-\frac{1}{2}(\boldsymbol{\mu}_i^h - \boldsymbol{\mu}_j^R)^\top (\boldsymbol{\Sigma}_i^h + \boldsymbol{\Sigma}_j^R)^{-1} (\boldsymbol{\mu}_i^h - \boldsymbol{\mu}_j^R)\right], \\ \boldsymbol{\mu}_{i,j} &= \boldsymbol{\Sigma}_{i,j} \left((\boldsymbol{\Sigma}_i^h)^{-1} \boldsymbol{\mu}_i^h + (\boldsymbol{\Sigma}_j^R)^{-1} \boldsymbol{\mu}_j^R \right), \\ \boldsymbol{\Sigma}_{i,j}^{-1} &= (\boldsymbol{\Sigma}_i^h)^{-1} + (\boldsymbol{\Sigma}_j^R)^{-1}. \end{aligned} \quad (\text{IV.5})$$

Since $\delta(\mathbf{h}, \mathbf{f}^R)\mathcal{N}(\mathbf{f}^R | \boldsymbol{\mu}_j^R, \boldsymbol{\Sigma}_j^R) = \mathcal{N}(\mathbf{h} | \boldsymbol{\mu}_i^R, \boldsymbol{\Sigma}_i^R)$, we have

$$p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) = \sum_{i=1}^{N_t^h} \sum_{j=1}^{N_t^R} Z_{i,j}^{-1} w_i^h w_j^R \mathcal{N}(\mathbf{h} | \boldsymbol{\mu}_{i,j}, \boldsymbol{\Sigma}_{i,j}). \quad (\text{IV.6})$$

Because of the equivalence of Equations IV.4 and IV.6, our initial question— \mathbf{h}^* or \mathbf{f}^{R*} ?—is easily resolved:

$$\begin{aligned} [\mathbf{h}, \mathbf{f}^R]^* &= \arg \max_{\mathbf{h}, \mathbf{f}^R} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) \\ u_{GSC}^s &= \mathbf{h}^* = \mathbf{f}^{R*}. \end{aligned} \quad (\text{IV.7})$$

We immediately note that although the arguments in Equations IV.4 and IV.6 reduce to either \mathbf{f}^R or \mathbf{h} , each expansion still carries all the information of both human and robot models in $Z_{i,j}^{-1}, w_i^h, w_j^R, \boldsymbol{\mu}_{i,j}$ and $\boldsymbol{\Sigma}_{i,j}$. By enforcing $\psi(\mathbf{h}, \mathbf{f}^R) = \delta(\mathbf{h}, \mathbf{f}^R)$ we recover a principle of joint decision making: the human and robot models must reach consensus. Thus, we do not have to rely on heuristics to determine who decides what and when. GSC provides a statistical mechanism to extract the most meaningful information from both decision makers.

Lemma IV.2 (Recovering linear blending). *In Equation IV.4, if $N_t^h = N_t^R = 1$ then linear blending is recovered.*

Proof: If $N_t^h = N_t^R = 1$, then $\arg \max_{\mathbf{h}} p(\mathbf{h} | \mathbf{z}_{1:t}^h) = \boldsymbol{\mu}_1^h$ and $\arg \max_{\mathbf{f}^R} p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) = \boldsymbol{\mu}_1^R$. The linear blending

solution is $u_{LB}^s = K_h \mu_1^h + K_R \mu_1^R$. Further, $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) \propto \mathcal{N}(\mathbf{f}^R | \mu_{1,1}, \Sigma_{1,1})$, so

$$\begin{aligned} \arg \max_{\mathbf{h}, \mathbf{f}^R} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) &= \mu_{1,1} \\ &= \Sigma_{1,1}^{-1} [(\Sigma_1^h)^{-1} \mu_1^h + (\Sigma_1^R)^{-1} \mu_1^R]. \end{aligned}$$

By choosing $K_h = \Sigma_{1,1}^{-1} (\Sigma_1^h)^{-1}$ and $K_R = \Sigma_{1,1}^{-1} (\Sigma_1^R)^{-1}$, we recover linear blending. If the human and robot distributions are unimodal normal, K_h, K_R must be the covariances. ■

GSC thus generalizes linear blending; Lemma IV.3 shows that GSC is lower bounded in performance by linear blending. Further, this lemma prompts the following observation: $\delta(\mathbf{h}, \mathbf{f}^R)$ partitions $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ into a mixture of linear blends, since each $\mu_{i,j}$ is a linear blend. We explore this idea:

Definition 6 (Mixture based linear blending). *Suppose that $p(\mathbf{h} | \mathbf{z}_{1:t}^h) = \sum_{i=1}^{N^h} w_i^h \mathcal{N}(\mathbf{h} | \mu_i^h, \Sigma_i^h)$ and $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) = \sum_{j=1}^{N^R} w_j^R \mathcal{N}(\mathbf{f}^R | \mu_j^R, \Sigma_j^R)$. Then let*

$$\begin{aligned} \mu^{h*} &= \arg \max_{\mathbf{h}} p(\mathbf{h} | \mathbf{z}_{1:t}^h) \\ \mu^{R*} &= \arg \max_{\mathbf{f}^R} p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) \end{aligned}$$

with Σ^{h*}, Σ^{R*} corresponding to μ^{h*}, μ^{R*} . Mixture based linear blending is $u_{MBLB}^s = (\Sigma^{h*})^{-1} \mu^{h*} + (\Sigma^{R*})^{-1} \mu^{R*}$.

Mixture based linear blending is an obvious generalization of linear blending in Equation II.1. Our motivation is to make explicit what computations are required for classical shared control: to determine $\arg \max_{\mathbf{h}} p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ or $\arg \max_{\mathbf{f}^R} p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ as in Definition 6, we must find all the modes in the GP mixture.

Lemma IV.3. *GSC is performance lower bounded by mixture based linear blending.*

Proof: Suppose that μ^{h*}, μ^{R*} in Definition 6 are in agreement (e.g., Z_{h^*,R^*}^{-1} is large). Since w^{h*}, w^{R*} are the largest coefficients in the human and robot GP mixtures, $\arg \max_{\mathbf{h}, \mathbf{f}^R} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ recovers u_{MBLB}^s .

Conversely, assume that μ^{h*}, μ^{R*} do not agree ($Z_{i,j}^{-1}$ is small). Then some other mixture component in $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ optimizes human and robot agreement and intent better. ■

V. SPARSITY AND OPTIMALITY THEOREMS OF GENERALIZED AND CLASSICAL SHARED CONTROL

To illustrate the differences between classical and generalized shared control, we begin with an example (Figure 5): the human wants to go left with $w_L^h = 0.8$ and right with $w_R^h = 0.2$, and the robot prefers left with $w_L^R = 0.4$ and right with $w_R^R = 0.6$. That is,

$$\begin{aligned} p(\mathbf{h} | \mathbf{z}_{1:t}^h) &= \sum_{i=L,R} w_i^h \mathcal{N}(\mathbf{h} | \mu_i^h, \Sigma_i^h) \\ p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) &= \sum_{j=L,R} w_j^R \mathcal{N}(\mathbf{f}^R | \mu_j^R, \Sigma_j^R). \end{aligned}$$

Sampling the arg max preferences of the individual agents— $\mu_L^h = \arg \max_{\mathbf{h}} p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ and $\mu_R^R = \arg \max_{\mathbf{f}^R} p(\mathbf{f}^R |$

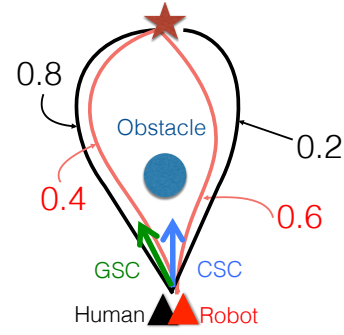


Fig. 5: The human prefers left with weight 0.8 and right with weight 0.2; The robot prefers left with weight 0.4 and right with weight 0.6. Classical shared control (CSC) results in a collision with the obstacle that both agents were trying to avoid. GSC efficiently finds the optimal solution.

$\mathbf{z}_{1:t}^R$ —requires fusing μ_L^h with flexibility Σ_L^h and μ_R^R with flexibility Σ_R^R . CSC directs the platform to head for the obstacle (blue arrow in Figure 5); by reasoning over independent samples, we have created an unresolvable disagreement between the two agents. In other words, no function $g(\mu_R^R, \mu_L^h, \Sigma^R, \Sigma^h)$ exists that can resolve this disagreement is unavailable. We make this argument precise:

Theorem V.1 (CSC suboptimal). *For any robot model $\mathcal{M}_h(\mathbf{h}, \mathbf{z}_{1:t}^h)$ and any human model $\mathcal{M}_R(\mathbf{f}^R, \mathbf{z}_{1:t}^R)$, classical arbitration $g(\mathbf{f}^R, \mathbf{h}, \theta^R, \theta^h)$, which reasons over finite sets of independent model predictions $[\bar{\mathbf{h}}_i, \bar{\theta}_i^h]_{i=1}^N \sim \mathcal{M}_h(\mathbf{h}, \mathbf{z}_{1:t}^h)$ and $[\bar{\mathbf{f}}_i^R, \bar{\theta}_i^R]_{j=1}^M \sim \mathcal{M}_R(\mathbf{f}^R, \mathbf{z}_{1:t}^R)$, optimizes human and robot agreement and intent only if intention ambiguity is absent.*

Proof: If we only finitely sample from independent human $[\bar{\mathbf{h}}_i, \bar{\theta}_i^h]_{i=1}^N$ and robot $[\bar{\mathbf{f}}_i^R, \bar{\theta}_i^R]_{j=1}^M$ models, we cannot guarantee that agreement and intent will be optimized (contrast with Theorem V.3). In the language of multi-objective optimization, finding the Pareto front is inefficient. If $N, M \rightarrow \infty$ (or, equivalently, if the samples capture all the information in the parameters of the human and robot GP expansions), then the optimal human and robot agreement and intent solution can be recovered. Importantly, as $N, M \rightarrow \infty$, classical shared control becomes generalized shared control. ■

This theorem is a strong negative result: independent agent sampling results in shared controllers that are fragile to human or robot ambiguity; classical shared control is unsuitable for real world application. We now provide a positive result: GSC is human and robot agreement and intent optimal and sparse.

We revisit Figure 5: how should this situation be resolved? The human-robot system should either go left or right; going straight is not preferred by either agent; the robot views left and right as having similar value; and the human strongly prefers left. The optimal shared control is to proceed left. Consider Equation IV.4, adapted for our scenario:

$$p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) = \sum_{i,j=L,R} Z_{i,j}^{-1} w_i^h w_j^R \mathcal{N}(\mathbf{f}^R | \mu_{i,j}^R, \Sigma_{i,j}^R),$$

and note that $Z_{L,R}^{-1} \approx Z_{R,L}^{-1} \approx 0$ by Equation IV.5. This is an example of **sparsity** (illustrated in Figure 6); naively,




Fig. 6: Transforming $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ into a GP basis for the example in Figure 5. **L**: “native” coordinates model all joint actions as equivalent (blue arrows). On transforming into the GP basis, interactions are ranked according to joint optimality: unresolvable actions have trivial weight; viable joint actions are weighted heavily. GSC is $N_{\min} = \arg \min\{N_t^h, N_t^R\}$ sparse.

$p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ has four components. However, the normalization coefficient $Z_{i,j}^{-1}$ prunes those actions that cannot be resolved. GSC is well approximated by 2 modes:

$$p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) \approx Z_{L,L}^{-1} w_L^h w_L^R \mathcal{N}(\mathbf{f}^R | \boldsymbol{\mu}_{L,L}, \boldsymbol{\Sigma}_{L,L}) + Z_{R,R}^{-1} w_R^h w_R^R \mathcal{N}(\mathbf{f}^R | \boldsymbol{\mu}_{R,R}, \boldsymbol{\Sigma}_{R,R}).$$

Note that $N_{\min} = 2 = \arg \min\{N_t^h = 2, N_t^R = 2\}$. The $\boldsymbol{\mu}_{L,L}$ and $\boldsymbol{\mu}_{R,R}$ modes are far enough apart that $\arg \max_{\mathbf{h}, \mathbf{f}^R} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ can be determined by coefficient inspection and the fact that $Z_{L,L}^{-1} \approx Z_{R,R}^{-1}$. Thus,

$$u_{GSC}^s = \boldsymbol{\mu}_{L,L} = \boldsymbol{\Sigma}_{L,L} \left((\boldsymbol{\Sigma}_L^h)^{-1} \boldsymbol{\mu}_L^h + (\boldsymbol{\Sigma}_L^R)^{-1} \boldsymbol{\mu}_L^R \right).$$

N_{\min} -sparse GSC simultaneously optimizes human and robot agreement and intent in the presence of intent ambiguity and flexibility. We make this observation precise and general:

Theorem V.2 (GSC is N_{\min} -Sparse). *Given $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ and $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ as in Equation III.1 and $N_{\min} = \arg \min\{N_t^h, N_t^R\}$, then $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$, as derived in Equation IV.4, has N_{\min} non-trivially weighted components.*

Proof: Recall that

$$Z_{i,j}^{-1} = \exp \left[-\frac{1}{2} (\boldsymbol{\mu}_i^h - \boldsymbol{\mu}_j^R)^\top (\boldsymbol{\Sigma}_i^h + \boldsymbol{\Sigma}_j^R)^{-1} (\boldsymbol{\mu}_i^h - \boldsymbol{\mu}_j^R) \right].$$

Thus, $Z_{i,j}^{-1}$ dominates w_i^h, w_j^R , since $0 \leq w_i^h, w_j^R \leq 1$. By computing $Z_{i,j}^{-1}$ for all i, j , we can efficiently rank the joint modes in terms of importance to the joint $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$: in order for a $Z_{i,j}^{-1}$ to be exponentially small, $(\boldsymbol{\mu}_i^h - \boldsymbol{\mu}_j^R)^2$ must be large (e.g., the human and robot disagree). Further, only N_{\min} modes can have non-trivial weight: if the human or robot holds N_{\min} intentions, they cannot agree with more than N_{\min} intentions of the other agent. ■

Theorem V.3 (GSC optimal). *For the human model $p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ and the robot model $p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$ as in Equation III.1, determining the $\arg \max_{\mathbf{h}, \mathbf{f}^R}$ of the N_{\min} -sparse distribution $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ optimizes human and robot agreement and intent.*

Proof: Consider

$$\begin{aligned} u_{GSC}^s &= \arg \max_{\mathbf{h}, \mathbf{f}^R} p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t}) \\ &= \arg \max_{\mathbf{h}, \mathbf{f}^R} \delta(\mathbf{h}, \mathbf{f}^R) p(\mathbf{h} | \mathbf{z}_{1:t}^h) p(\mathbf{f}^R | \mathbf{z}_{1:t}^R) \\ &\approx \arg \max_{\mathbf{f}^R} \sum_{k=1}^{N_{\min}} Z_k^{-1} w_k^h w_k^R \mathcal{N}(\mathbf{f}^R | \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k). \end{aligned}$$

The coefficient $Z_k^{-1} w_k^h w_k^R$ models human and robot agreement Z_k^{-1} , human intent preference w_k^h , and robot intent preference w_k^R . Since the discarded $Z_{i,j}^{-1}$'s correspond to disagreement modes, the solution that optimizes human and robot agreement and intent is the $\arg \max_{\mathbf{f}^R}$ solution of the N_{\min} -sparse mixture. In the language of multi-objective optimization, the GP basis renders Pareto front discovery efficient. ■

Returning to Figure 5: what if the human only wants to go left and the robot only wants to go right ($w_L^h = w_R^R = 1$)? Conceptually, we have a fundamental human-robot disagreement; GSC reveals this disagreement, since the joint distribution $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ has zero probability everywhere (the left and right modes have coefficients $w_L^h = w_R^R = 0$). Since there is no solution for u_{GSC}^s , the platform stops. Critically, stopping is the only reasonable solution: the two decision makers have taken irreconcilable positions with respect to communication via trajectory data. A standard fix would be to place a prior on the human model or instantiate rules of control dominance. Both approaches are fundamentally flawed: a prior violates Theorem IV.1, and rules of dominance are a form of $g(\mathbf{f}^R, \mathbf{h}, \boldsymbol{\theta}^R, \boldsymbol{\theta}^h)$, albeit at another level of information, and would instigate a suboptimality with respect to “negotiation:” since only trajectory data is available, precedence cannot be resolved without additional information.

This presents an opportunity: the derivation of GSC provides a template for systematically incorporating additional human-robot interaction functionality. For instance, models of negotiation should be equipped with the right amount of data and the models should match that amount of data. Rather than blindly equipping our robots with elaborate capabilities, we are better advised to examine the limitations of available data sources. As we incorporate additional information sources—speech, gesture, etc—following the GSC derivation prevents us from building in architectural shortcomings, or, perhaps worse, using unnecessary data sources. By examining the mathematical foundations of human machine interaction, we can gain insight into more efficient design mechanisms.

Thus, N_{\min} -sparse GSC provides an optimal and computationally efficient solution to the shared control problem, and generalizes existing shared control approaches; [52], provides numerous examples showing the superiority of GSC over CSC.

VI. EVALUATION: SHARED CONTROL IN CROWDS

The variety of shared control applications is too large to discuss here. In [52] (<https://arxiv.org/pdf/1611.09490.pdf>), we explore the performance of GSC against CSC for numerous scenarios including remote tele-operation over unreliable networks, task handoffs, assistive driving in congestion, mode

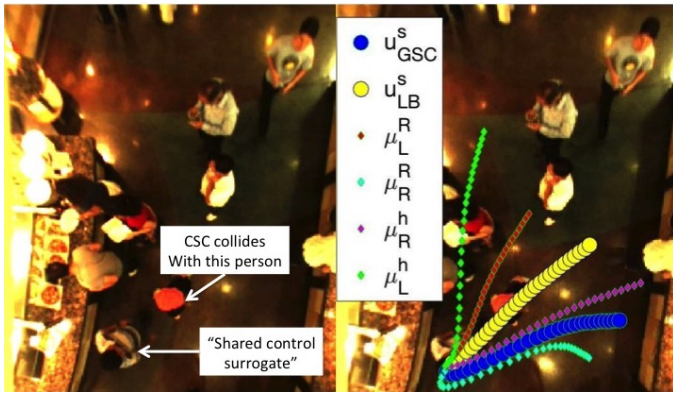


Fig. 7: **L**: Overview of scenario: “Shared control surrogate” can choose to go left or right. **R**: Human chooses robot suboptimal direction for social reasons. CSC blends left and right, collision results. GSC optimally arbitrates.

switching in commercial avionics, and explanation interfaces. Here, we evaluate GSC against CSC in a pedestrian crowd, where the choice to proceed left or right happens often. We compare the *architectures* of GSC and CSC when each framework is supplied with the same robot model and the same human model. Following Figure 1, we let $\mathcal{M}_h(\mathbf{h}, \mathbf{z}_{1:t}^h) = p(\mathbf{h} | \mathbf{z}_{1:t}^h)$ and $\mathcal{M}_R(\mathbf{f}^R, \mathbf{z}_{1:t}^R) = p(\mathbf{f}^R | \mathbf{z}_{1:t}^R)$, and assume these models are known. Because we only have access to overhead video data, we simulate a human robot team, i.e., we treat human movements as inputs to a hypothetical shared platform and run a multimodal path planner starting from the same position. We then combine the human inputs and the robot inputs using GSC or CSC. As illustrated in Figure 7, both the human and robot models are bi-modal so arbitration ambiguity occurs (since we do not know if the human will go to the left or right goal, the human model is bi-modal). We observe some results from our experiments:

1. When the robot optima μ^{R*} disagrees with the human optima μ^{h*} , CSC fails severely, often with a collision. See Figure 7 and Table I.

2. Computational costs for GSC and CSC are nearly the same because GSC is sparse (Δt in Table I is the computational overhead of GSC as compared to CSC). In our experiments, the modes in $p(\mathbf{h}, \mathbf{f}^R | \mathbf{z}_{1:t})$ are far enough apart that the $\arg \max$ is the largest coefficient in $\{Z_{i,j}^{-1} w_i^h w_j^R\}_{i,j=1}^{N_t^h, N_t^R}$. In general, this is not true, and finding the $\arg \max$ requires gradient or expectation-maximization (EM) methods. Even if the modes in Equation IV.4 are not far apart, however, the sparsity of GSC allows us to eliminate large portions of the joint search space, so gradient/EM computation is improved.

3. Testing shared control on static data can tell us only so much; in the real world, the human reacts to platform movements, interfaces, etc. Thus, our experiments are not exhaustive. These experiments serve to validate the above theorems, and demonstrate the utility of the GSC approach.

4. Since humans are naturally adept at maneuvering through crowds, it is difficult to draw conclusions such as “the shared controller improves the trajectory of the human.” That is, the humans almost always choose safe and efficient paths, once a direction is chosen (i.e., going left might have been more safe

C_{CSC}	$\mu_{u_{CSC}^s}^h$	C_{GSC}	$\mu_{u_{GSC}^s}^h$	$\mu(Z_{i \neq j}^{-1} / Z_{i=j}^{-1})$	Δt
Y	0.12m	N	0.015m	5×10^{-5}	5×10^{-5} s
Y	0.11m	N	0.013m	1×10^{-13}	2×10^{-4} s
Y	0.24m	N	0.07m	8×10^{-4}	2×10^{-4} s
Y	0.32m	N	0.1m	3×10^{-8}	3×10^{-5} s
N	0.03m	N	0.03m	3×10^{-7}	3×10^{-4} s

TABLE I: Rows report results from the runs. C_{CSC} , C_{GSC} is Y if CSC/GSC induced a collision/deadlock and N otherwise. The values $\mu_{u_{CSC}^s}^h$, $\mu_{u_{GSC}^s}^h$ are the average distance between the human trajectory and u_{CSC}^s , u_{GSC}^s respectively. The value $\mu(Z_{i \neq j}^{-1} / Z_{i=j}^{-1})$ computes the run average of “disagreeable” solutions divided by “agreeable” solutions; it quantifies the difference in “importance” value of disagreeable versus agreeable solutions. The value Δt is the additional compute time required by GSC compared to CSC. The last row is a run with no ambiguity; GSC recovers CSC.

and efficient, but along the right route the human stays safe and efficient). Thus, we compare how well the shared controller compares to the human trajectory: how close does GSC or CSC come to what the human does (variables $\mu_{u_{CSC}^s}^h$, $\mu_{u_{GSC}^s}^h$ in Table I)? Overall, CSC strays quite far from the human’s trajectory, while GSC stays very close.

Test sequence identification We tested sequences where the human pursued a robot suboptimal but still safe and efficient path (e.g., overriding the human was unnecessary; typically, the human followed social conventions, as in Figure 7). We chose these sequences in order to test Theorems V.3 (GSC optimal) and V.1 (CSC suboptimal), and we ran computational efficiency tests to validate Theorem V.2 (GSC sparse) against a baseline of mixture based linear blending. We also tested a sequence where the human chose the robot optimal strategy in order to test that GSC is indeed lower bounded by linear blending, as in Lemma IV.3 (last row in Table I).

Testing procedures In all sequences, we tested mixture based linear blending (Definition 6) against GSC (Equation IV.7). For GSC, we leveraged sparsity by computing values $Z_{i,j}^{-1}$ for all i, j , identified the i^*, j^* that corresponded to the largest value of $Z_{i,j}^{-1}$, and then computed

$$u_{GSC}^s = \mu_{i^*, j^*}^s = \Sigma_{i^*, j^*} \left((\Sigma_{i^*}^h)^{-1} \mu_{i^*}^h + (\Sigma_{j^*}^R)^{-1} \mu_{j^*}^R \right).$$

We catalog our findings in Table I.

VII. CONCLUSIONS AND FUTURE WORK

We formalized classical shared control and proved that regardless of human model, robot model, or arbitration function, CSC is not guaranteed to optimize human and robot agreement and intent unless intention ambiguity is absent. We introduced generalized shared control, and proved that it optimizes human and robot agreement and intent under arbitrary ambiguity in a computationally efficient manner. After Theorem V.3, we discussed the limitations of GSC; for example, trajectory data does not provide enough information to manage a negotiation with appropriate guarantees. We must be careful however, since Theorem V.1 still applies: if the models and data are extended in a CSC architecture, negotiation will be suboptimal. However, the derivation of GSC provides a template for how to accommodate additional functionalities—like negotiation—without building in suboptimalities.

REFERENCES

- [1] D. Aarno, S. Ekvall, and D. Kragic. Adaptive virtual fixtures for machine-assisted teleoperation tasks. In *ICRA*, 2005.
- [2] P. Aigner and B. McCarragher. Human integration into robot control utilising potential fields. In *ICRA*, 1997.
- [3] J. Alonso-Mora, P. Gohl, S. Watson, R. Siegwart, and R. Beardsley. Shared control of autonomous vehicles based on velocity space optimization. *ICRA*, 2014.
- [4] A. Aly, E. Zeidan, A. Hamed, and F. Salem. An antilock-braking systems control: A technical review. In *Intelligent Control and Automation*, 2011.
- [5] G. Aoude, B. Luders, J. Joseph, N. Roy, and J. How. Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns. *Autonomous Robots*, 2011.
- [6] A. Bejczy. Recent developments in advanced teleoperation at jpl. In *IEEE International Conference on Advanced Robotics*, 1991.
- [7] J. Bradshaw, V. Dignum, and C. J. an M. Sierhuis. Human-agent-robot teamwork. In *ACM/IEEE International Conference on Human-Robot Interaction*, 2012.
- [8] J. Bradshaw, P. Feltovich, M. Johnson, L. Bunch, M. Breedy, T. Eskridge, H. Jung, J. Lott, and A. Uszok. Coordination in human-agent-robot teamwork. In *International Symposium on Collaborative Technologies and Systems*, 2008.
- [9] J. Burke and R. Murphy. From remote tool to shared roles. In *IEEE Robotics and Automation Magazine*, 2008.
- [10] T. Carlson and Y. Demiris. Collaborative control for a robotic wheelchair: Collaborative control for a robotic wheelchair: Evaluation of performance, attention and workload. *IEEE Transaction on Systems, Man, and Cybernetics—Part B: Cybernetics*, 2012.
- [11] K. Christoffersen and D. Woods. How to make automated systems team players. In *Advances in Human Performance and Cognitive Engineering Research*. JAI Press/Elsevier, 2002.
- [12] N. Cooke, J. Gorman, C. W. Myers, and J. L. Durand. Interactive team cognition. *Cognitive Science*, 2013.
- [13] H. M. Cuevas, S. M. Fiore, B. S. Caldwell, and L. Strater. Augmenting team cognition in human-automation teams performing in complex operational environments. In *Aviation Space and Environmental Medicine*, 2007.
- [14] D.A. Abbink, M. Mulder, and E.R. Boer. Haptic shared control: smoothly shifting control authority. *Cognition, Technology and Work*, 2012.
- [15] S. Dekker and E. Hollnagel. Human factors and folk models. *Cognition, Technology and Work*, 2003.
- [16] S. Dekker and D. Woods. MABA-MABA or abracadabra? progress on human-automation coordination. *Cognition, Technology and Work*, 2002.
- [17] 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.
- [18] 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.
- [19] M. Demir, N. J. McNeese, and N. J. Cooke. Team communication behaviors of the human-automation teaming. In *IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support*, 2016.
- [20] M. Desai and H. Yanco. Blending human and robot inputs for sliding scale autonomy. In *IEEE ROMAN*, 2005.
- [21] M. B. Dias, B. Kannan, B. Browning, E. Jones, B. Argall, M. F. Dias, M. B. Zinck, M. Veloso, and A. T. Stentz. Sliding autonomy for peer-to-peer human-robot teams. Technical Report CMU-RI-TR-08-16, Robotics Institute, 2008.
- [22] A. Dragan and S. Srinivasa. Formalizing assistive teleoperation. *Robotics: Science and Systems*, 2012.
- [23] A. Dragan and S. Srinivasa. A policy blending formalism for shared control. *International Journal of Robotics Research*, 2013.
- [24] A. Enes and W. Book. Blended shared control of Zermelo’s navigation problem. *American Control Conference*, 2010.
- [25] S. Fang, M. Peshkin, and M. Maclver. Human-in-the-loop active electrosense. *Bioinspiration and Biomimetics*, 2016.
- [26] Federal Aviation Administration. *Advanced Avionics Handbook: Chapter 4, Automated Flight Controls*. Federal Aviation Administration, 2014.
- [27] P. Fitts. Human engineering for an effective air navigation and traffic control system. Technical report, National Research Council, Washington, D.C., 1951.
- [28] K. Fitzsimmons, E. Tzorakoleftherakis, and T.D. Murphey. Optimal human-in-the-loop interfaces based on Maxwell’s demon. In *American Control Conference*, 2016.
- [29] J. Fu and U. Topcu. Probably approximately correct MDP learning and control with temporal logic constraints. In *Robotics: Science and Systems*, 2014.
- [30] C. Fulgenzi, C. Tay, A. Spalanzani, and C. Laugier. Probabilistic navigation in dynamic environment using rapidly-exploring random trees and Gaussian processes. In *IEEE International Conference on Intelligent Robots and Systems*, 2008.
- [31] 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.
- [32] 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.
- [33] E. Hollnagel and D. Woods. *Joint cognitive systems: patterns in cognitive systems engineering*. CRC Press, 2005.
- [34] 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.
- [35] 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.
- [36] W. B. Knox and P. Stone. TAMER: Training an agent manually via evaluative reinforcement. In *IEEE International Conference on Development and Learning*, 2008.
- [37] H. Kretzschmar, M. Spies, C. Sprunk, and W. Burgard. Socially compliant mobile robot navigation via inverse reinforcement learning. In *The International Journal of Robotics Research*, 2016.
- [38] 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.
- [39] C. Martin, R. Macfadzean, and K. Barber. Supporting dynamic adaptive autonomy for agent-based systems. In *AAAI Workshop on AI and Manufacturing Research Planning*, 1996.
- [40] N. McNeese, N. Cooke, R. Gray, and M. Fedele. Knowledge elicitation methods for developing insights into team cognition during team sports. *Advances in Human Factors in Sports and Outdoor Recreation*, 2017.
- [41] C. Miller. Frameworks for supervisory control: Characterizing relationships with uninhabited vehicles. *Journal of Human Robot Interaction*, 2012.
- [42] 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, 2015.
- [43] K. Muelling, A. Venkatraman, J.-S. Valois, J. Downey, J. Weiss, S. Javdani, M. Hebert, A. Schwartz, J. Collinger, and A. Bagnell. Autonomy infused teleoperation with application to BCI manipulation. In *Proceedings of Robotics: Science and Systems*, 2015.
- [44] R. Murphy, J. Casper, M. Micire, and J. Hyams. Mixed initiative control of multiple heterogeneous mixed initiative control of multiple heterogeneous robots for urban search and rescue. In *Robotics and Automation*, 2000.
- [45] A. Okamura. Methods for haptic feedback in teleoperated robot-assisted surgery. *The Industrial Robot*, 2004.
- [46] R. Parasuraman, T. Sheridan, and C. Wickens. A model for types and levels of human interaction with automation. *IEEE Transaction on Systems, Man, and Cybernetics—Part A: Systems and Humans*, 2000.
- [47] 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.
- [48] 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.
- [49] C. E. Rasmussen and C. Williams. *Gaussian Processes for Machine Learning*. MIT Press, 2006.

- [50] T. Sheridan and W. Verplank. Human and computer control of undersea teleoperators. Technical report, MIT, 1978.
- [51] P. Trautman. Assistive planning in complex, dynamic environments. In *IEEE Systems, Man, and Cybernetics* (<http://arxiv.org/abs/1506.06784>), 2015.
- [52] P. Trautman. Generalized shared control versus classical shared control: Illustrative examples (<https://arxiv.org/pdf/1611.09490.pdf>). Technical report, Galois Inc., 2017.
- [53] P. Trautman and A. Krause. Unfreezing the robot: Navigation in dense interacting crowds. In *IROS*, 2010.
- [54] P. Trautman, J. Ma, A. Krause, and R. M. Murray. Robot navigation in dense crowds: the case for cooperation. In *ICRA*, 2013.
- [55] P. Trautman, J. Ma, R. M. Murray, and A. Krause. Robot navigation in dense human crowds: Statistical models and experimental studies of human robot cooperation. *International Journal of Robotics Research*, 2015.
- [56] C. Urdiales, E. Pérez, G. Peinado, M. Fdez-Carmona, J. Peula, R. Annichiarico, F. Sandoval, and C. Caltagirone. On the construction of a skill-based wheelchair navigation profile. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 2013.
- [57] C. Urdiales, J. Peula, M. Fdez-Carmona, C. Barrué, E. Pérez, I. Sánchez-Tato, J. del Toro, F. Galluppi, U. Cortés, R. Annichiarico, C. Caltagirone, and F. Sandoval. A new multi-criteria optimization strategy for shared control in wheelchair assisted navigation. *Autonomous Robots*, 2011.
- [58] H. Wang and X. Liu. Adaptive shared control for a novel mobile assistive robot. *IEEE/ASME Transaction on Mechatronics*, 2014.
- [59] M. Wang and J. Liu. Interactive control for internet-based mobile robot teleoperation. *Robotics and Autonomous Systems*, 2005.
- [60] P. Welinder, S. Branson, S. Belongie, and P. Perona. The multidimensional wisdom of crowds. In *NIPS*, 2010.
- [61] H. Yu, M. Spenko, and S. Dubowsky. An adaptive shared control system for an intelligent mobility aid for the elderly. *Autonomous Robots*, 2003.