Less Is More: Mixed Initiative Model Predictive Control With Human Inputs

Rahul Chipalkatty, *Student Member, IEEE,* Greg Droge, *Student Member, IEEE,* and Magnus Egerstedt, *Fellow, IEEE*

***Abstract*—This paper presents a new method for injecting human inputs in mixed initiative interactions between humans and robots. The method is based on a model predictive control (MPC) formulation, which inevitably involves predicting the system (robot dynamics as well as human inputs) into the future. These predictions are complicated by the fact that the human is interacting with the robot, causing the prediction method itself to have an effect on future human inputs. We investigate and develop different prediction schemes, including fixed and variable horizon MPCs and human input estimators of different orders. Through a search-and-rescue-inspired human operator study, we arrive at the conclusion that the simplest prediction methods outperform the more complex ones, i.e., in this particular case, less is indeed more.**

***Index Terms*—human-robot interaction, model-predictive control, mixed initiative initeractions.**

I. Introduction

Despite advances in autonomous robotics and automation, some tasks still require human intervention or guidance to mediate uncertainties in the environment, or to manage the complexities of the task. In response to this, robot controllers have been designed that combine the strengths of both autonomous agents, which are adept at handling lower level, repetitive control tasks, and humans, whom are superior at handling higher-level cognitive tasks. Researchers in the Human-Robot Interaction field refer to this as *mixed initiative interactions*, or dynamic autonomy (e.g., [1], [2], [3], [4]).

Previous work on mixed initiative interactions has primarily focused on the development of effective graphical user interfaces or haptic feedback to relay task-dependent data to the human, and to relay human control information to the automatic controller or autonomous, robotic agent (e.g., [5], [6], [7], [8], [9], [10]). In this paper, we largely ignore this issue. Instead we focus on the design of the actual control laws and simply assume that the human operator already has some effective means of interacting with the system.

Connected to the notion of mixed initiative interactions is that of *sliding levels of autonomy*, e.g., [2], whereby a human operator may influence the system at varying degrees, typically as a function of the difficulty of the task. In [8], this concept is realized by allowing the user to set the autonomy level through a user interface in which the human commands are weighted. However, in order to set the human command weighting, this approach relies on the human operator to be constantly aware of the difficulty of the task and the capabilities of the robot, while simultaneously recognizing the state of the system with regards to the autonomous task-completion. In [11], the human input weighting is determined by a designed threat assessment function such that high threat levels lead to higher automation and lower threat levels lead to more human control. Sliding levels of autonomy will also be present in our formulation, albeit implicitly in that we will not insist on any formal assessment of the difficulty of the task. Rather, the autonomous controller will ensure that the task is completed and, beyond that, the human user may inject any input signals.

Task completion is central to the work in this paper and none of the previously mentioned references do indeed guarantee that the task actually gets done. For instance, [12], [13], [14], [15], [16], [17] present methods for composing autonomous control actions with teleoperation, but they do not contain formal guarantees for task completion.

Specifically, in [16], the authors present a method of combining user intent with obstacle avoidance through a fuzzy-logic scheme. The control is applied to an assistive walking device, where one of the two control sources are used based on a decision table of fuzzy inputs. The resulting behavior is a smoothed trading-off of control authority between user and obstacle avoidance algorithm and not concurrent control. As we know from hybrid control, switching between two stable controllers does not guarantee stability and this work neither discusses nor presents any such results.

Similarly, the authors in [13], present a traded and shared control where the shared control is actually a divided approach in that the automated control drives five degrees of freedom and the human operator fully controls the sixth degree of freedom. This essentially means that each is controlling decoupled aspects of the systems. The operator and automatic control never concurrently share control over the same aspects of the system. In [17], the authors propose a weighted sum of obstacle avoidance control and human input with the weights are determined by measuring smoothness, effectiveness, and safety of each control signal. However, this method gives no guarantees of obstacle avoidance in the face of human operator output.

On the other hand, the approach in [15] allows users and automated behaviors concurrent control of a smart wheelchair to reach goals and avoid obstacles. The method projects the operator's intended velocity onto safe zones away from obstacles

Rahul Chipalkatty is with the Department of Mechanical Engineering, Georgia Institute of Technology, Atlanta, GA, 30332 USA e-mail: chipalkatty@gatech.edu.

Greg Droge and Magnus Egerstedt are with the Department of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, 30332 USA e-mail: gregdroge@gatech.edu, magnus@gatech.edu.

Manuscript received April 19, 2005; revised January 11, 2007.

and allows full operator control when safe, while modifying the commands if not safe and stopping the vehicle if need be. While effective in avoiding collisions, the resulting controller may prevent the user from ever reaching the intended goal. While the work in [12] contains a ”shared control” mode, it is really a supervisory control type where operators give high-level control commands through a collaborative control paradigm, e.g. giving human speech-like commands such as ”Drive left.” In [14], the authors take a unique approach in having the autonomy ”probe” the operator by intervening in the task when the autonomy decides there a discrepancy between what it senses about the task and what the human does. Like the previous methods described here, there are no guarantees that these autonomous interventions are accomplishing the task.

In contrast to these methods, [18] does indeed present a mixed initiative controller that guarantees a certain level of task completion. The strategy in [18] is based on navigation functions combined with human inputs to drive a differential drive robot around obstacles to a specific goal represented by the global minimum of the navigation function. The operator can drive the robot away from the planned navigation function path but, once the user stops issuing commands, the controller will drive the system towards the goal state again, with guaranteed task completion as long as the human operator stops issuing commands eventually.

In this paper, we frame the mixed initiative interaction problem as a model predictive control problem, following the initial work in [19]. The proposed approach makes a distinction between low-level (automatic controller) tasks and high-level (human) tasks, with completion guarantees associated with the low-level tasks without the need for strong assumptions on the human input signals. The technical difficulty associated with this approach is that it requires the prediction of human inputs into the future. And, said predictions will influence the performance of the robot which, in turn, will change the human inputs. So, care must be taken when constructing the prediction methods and one of the key investigations in this paper is how one should handle this “prediction-human” feedback loop. As such, we propose and develop a number of different strategies of increasing complexity and sophistication, and compare these strategies through a human operator study. The experimental human operator trial is based on a search-and-rescue inspired, cooperative human-robot navigation task. The main finding in this paper is that due to the complex connections between the human inputs and the prediction methods, the simpler methods outperform the more elaborate ones.

II. Mixed Initiative Interactions as an MPC

A. Problem Formulation

In this section, we develop the basic framework that will allow us to cast mixed initiative interactions in terms of model predictive control actions. Suppose first that the human operator is directly controlling the robot, with dynamics given by

$$x_{k+1} = f(x_k, v_k). \tag{1}$$

Here $f : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n$ represents the robot dynamics, with $x_k \in \mathbb{R}^n$ being the state of the robot, and where the human operator is issuing the commands $v_k \in \mathbb{R}^m$, as shown in Figure 1(a). We will formulate task completion in terms of making the system reach a set of target states $\mathbb{X}_f \subset \mathbb{R}^n$, and achieving this with direct human control may, for a various of reasons, not be feasible, desirable, or necessary. This part of the task will thus be offloaded to an automatic controller, as proposed in [20]. The problem we address in this paper is ultimately to devise a controller that drives the system in such a way that both the state constraints are satisfied (low-level task) and the human operator's “intentions” for the system behavior (high-level task) are respected as much as possible.

In order to preserve the human operator's intentions – without having to establish informed estimates of what these intentions might actually be – the control law will be designed in such a way that it minimizes deviations from the human input while also ensuring that the state constraints are satisfied. To this end, we replace v_k in (1) by a control input u_k, resulting in

$$x_{k+1} = f(x_k, u_k), \tag{2}$$

where the idea is that u_k should somehow be “close” to v_k (the human input), as shown in Figure 1.

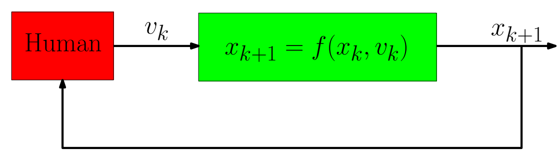

(a) Direct Human Control

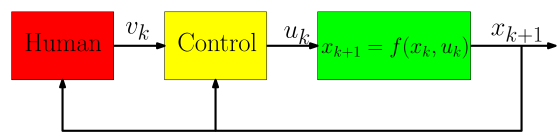

(b) Human-in-the-Loop Control

Fig. 1. Control methodology for the proposed MPC-based human-in-the-loop controller: The proposed controller will stay close to the human input signal while completing a lower-level task.

Ensuring that a system will reach some target set while staying close to the human inputs requires some way of predicting where the human operator intends to drive the system which, in turn, requires a prediction of future human operator inputs. In Section III-A, two such methods of increasing levels of complexity (zero-order hold and system identification) are presented. For the purpose of the current discussion, we simply assume that we somehow, at every time instant k, have a predicted sequence of human input values, denoted by $\mathcal{V}_k = \{v_k, \ldots, v_{k+N_k-1}\}$, where $N_k \in \mathbb{N}$ is the prediction horizon.[1]

We now need to find the control sequence $\mathcal{U}_k = \{u_k, \ldots, u_{k+N_k-1}\}$, that minimizes its deviations from $\mathcal{V}_k$, while ensuring that the state, at the end of the sequence, satisfies the state constraint, i.e., that $x_{k+N_k} \in \mathbb{X}_f$. The only

[1]We use the subscript k in the horizon N_k to allow for the prediction horizon to vary and, thus, depend on k.

input in $\mathcal{U}_k$ that will actually be applied to the robot is the first one, i.e., u_k. After that, a new predicted human input sequence is found and a new input sequence is computed. The receding horizon optimal control problem that is solved at each time instant k is thus given by

$$\min_{\mathcal{U}_k} V_{N_k}(\mathcal{V}_k, \mathcal{U}_k), \tag{3}$$

where the cost is given by a sum of instantaneous stage costs

$$V_{N_k}(\mathcal{V}_k, \mathcal{U}_k) = \sum_{i=k}^{k+N_k-1} L(v_i, u_i), \tag{4}$$

subject to the constraints

$$x_{k+1} = f(x_k, u_k), \tag{5}$$
$$x_{k+N_k} \in \mathbb{X}_f. \tag{6}$$

We will refer to this problem as $\mathcal{P}_{MPC}$, and we note that the cost (4) is designed to penalize deviations from the human command in order to preserve human intent while the terminal state constraint (6) guarantees that the state constraint, which is required for the lower level task, is enforced at the end of the time horizon. Without this terminal constraint, the control would simply equal the predicted human input. However, the presence of the terminal constraint will most likely cause the control to deviate from the human input.

The choice of horizon, N_k, is critical in that a large N_k requires that the future prediction of the human input be accurate over a long time horizon. Otherwise the computed control will not accurately reflect the intent of the user – it will still reach the target state though since this is a hard constraint. If N_k is too small, the control effort attempts to reach the constraint set within a small amount of time, making the deviations from the human input potentially quite large. As such, N_k must be short enough so that the prediction of future user inputs is valid yet long enough to ensure that the user intentions are respected.

Dynamic, sliding autonomy is achieved by this controller in that any human operator input (the control sequence $\mathcal{V}_k$), that drives the system to the goal set at the end of the control horizon will be used as given (i.e. fully manual control). If this is not the case, the system will seek to correct the command while trying to also respect the operator intent. In this way, the human operator can be a detached supervisor or have a more active role depending on how close the human commands come to carrying out the lower-level task. Additionally, if the human provides no inputs whatsoever, the robot will simply satisfy the low-level task without any human intervention (full autonomy). We thus have a control architecture that supports varying levels of autonomy without any need to explicitly specify the levels of autonomy.

B. Task Completion

A number of results have been obtained establishing asymptotic stability for MPCs (e.g., [21], [22], [23], [24]). The stability arguments typically rely on a terminal constraint or cost that penalizes deviations from an equilibrium point. The formulation in this paper differs from previous formulations in that the prediction horizon is allowed to vary and that the terminal constraint is given in terms of a set that may not contain any equilibrium points. However, convergence results – and thus task completion guarantees – can still be given. To this end, we will rely on the so-called dual-mode MPC technique from [21] and [22], where the system is driven to the constraint set by the MPC, and then another (locally invariant) controller is employed within the constraint set. This second mode is technically needed for the convergence results to hold, and we will simply assume that, when in the terminal constraint set, the human input will be used as the ”locally invariant” controller. Thus, we require that the human operator must not be “stupid” in the sense that it will not actively force the system away from the goal set once it has been reached. Phrased in other words, the operator is capable to keeping the system in the constraint set and we let the set of the corresponding admissible human inputs be denoted by $v \in \mathbb{V}(x)$, where $x \in \mathbb{X}_f$.

We must moreover make some (mild) assumptions on the stage cost L, namely that it is bounded below by a K-function, which gives us a positive and increasing cost with respect to the norm of the difference between the human inputs and the control inputs. The cost is furthermore assumed to be zero when in the terminal constraint set, which is consistent with the human operator providing the invariant control action in that set. The final assumption needed to ensure task completion is that the goal set can indeed be reached. This is particularly important when the prediction horizon is allowed to vary, since too short a horizon may otherwise prove problematic. We here summarize these four assumptions:

- A1 $L(v_k, u_k) \geq \gamma(\|(u_k - v_k)\|)$ for some K-function γ, with $L(0, 0) = 0$.
- A2 $L(v_k, v_k)$=0 for all $v_k \in \mathbb{V}(x)$.
- A3 The set $\mathbb{X}_f$ is positively invariant under control $v \in \mathbb{V}(x)$ in the sense that $f(x, v) \in \mathbb{X}_f,\ \forall x \in \mathbb{X}_f,\ \forall v \in \mathbb{V}(x)$.
- A4 There exists a lower bound $M \geq 1$ such that the goal set is reachable from all states over any time horizon $N \geq M$.

Theorem 2.1 (Task Completion): Under the assumptions A1-A4, the state will converge to the constraint set, $\mathbb{X}_f$, as $k \to \infty$, when the first element in the optimal solution to $\mathcal{P}_{MPC}$ is applied at each iteration.

The proof of this theorem is given in Appendix A, but before we move on to the different, candidate prediction methods under consideration, some words should be said about the reasonableness of the assumptions. Assumption A1 is clearly not restrictive and the particular choice of stage cost that we will use in the operator trials will simply be given by $L(u, v) = \|u - v\|^2$.

Assumptions A2 and A3 ensure that once the system reaches the terminal set, the stage cost is zeroed and the system will not be driven out of the constraint set, i.e. $u = v$ (the applied control is the human input) and $x_{k+1} = f(x_k, u_k) \in \mathbb{X}_f\ \forall\ x_k \in \mathbb{X}_f$, with $u \in \mathbb{V}(x)$. These two conditions imply a “strong” assumption in that we assume that the bounds on the human operator control and the ability of the operator is sufficient for keeping the state within the constraint set once this set has been reached. In other words, the human operator is trusted with the control to make $\mathbb{X}_f$ invariant. We argue that this is a reasonable assumption because once the system has converged to the state constraint set, it should be obvious to the human operator that large incorrect command inputs will not be beneficial. Finally, Assumption A4 is in essence a controllability assumption, which may or may not hold, depending on the robot dynamics.

III. Design Choices

The strength of Theorem 2.1 is that it is quite general and it allows us to combine human inputs with guaranteed task completion for a large class of robotic systems. However, when actually deploying it, a number of design choices must be made. This section focuses on these choices. And, as the overall aim is to understand what constitutes good such choices, a variety of different methods must be investigated. We do not intend to cover all possible such methods, but rather derive methods with increasing levels of complexity since an explicit aim is to investigate if humans prefer simpler or more involved prediction methods when interacting with robotic systems. We will thus start with a discussion of the prediction methods used, followed by a way of updating the prediction horizon in an adaptive manner. The last design consideration to be discussed involves how to actually solve the mixed initiative MPC in an effective manner for linear systems with quadratic costs.

A. Human Input Prediction Methods

As already noted, we need to be able to predict human inputs in order to solve $\mathcal{P}_{MPC}$. In this section, we discuss two such prediction methods, although we note that other such methods are conceivable. The first is Zero-Order Hold (ZOH), where only the current human input is needed to make the prediction, which constitutes the simplest possible prediction method. The second method is prediction by Least Squares System Identification (SID), which requires that we store a certain number of past human inputs, which thus represents a more complex (higher-order) prediction method.

1) Zero-Order Hold Prediction (ZOH): The ZOH prediction method simply says that the future human inputs will all be the same as the current human input. And, as such, this method represents one extreme of the complexity spectrum in that a less involved prediction method can hardly be envisioned. Given v_k, the predicted human input sequence to be used in $\mathcal{P}_{MPC}$, is thus given by

$$\mathcal{V}_k = \{v_k, v_k, \ldots, v_k\}.$$

Although this prediction is simple (and inaccurate), it will be shown to be surprisingly effective in experimentation.

2) Least-Squares System Identification (SID): Human input predictions using linear least-squares system identification allows us to make predictions that reflect longer-term trends in the human inputs. For instance, a system identification approach would make better predictions of periodic human inputs, and this method not only serves as a way to use past information to predict future human inputs, but we will be able to gauge the performance of this prediction and, as a consequence, update the prediction horizon accordingly.

At time k, let the $N_s \in \mathbb{N}$ past and current human input values be denoted by $s_k = [v_{k-N_s}, \ldots, v_k]^T$, and

$$H_k = \begin{bmatrix} v_{k-N_s-N_s} & \cdots & v_{k-N_s} \\ \vdots & \vdots & \vdots \\ v_{k-2-N_s} & \cdots & v_{k-2} \\ v_{k-1-N_s} & \cdots & v_{k-1} \end{bmatrix},$$

with model parameters $\phi = [\phi_{N_s}, \ldots, \phi_1]$ that must be determined from the past data.

The least squares problem, $s_k = H_k \phi^T$, is then solved, giving the parameters

$$\phi = (H_k^T H_k)^{-1}(s_k^T H_k). \tag{7}$$

However, the quantity, $(H_k^T H_k)$ is potentially singular, so the Levenberg-Marquardt procedure [25] is used to regularize this matrix, resulting in

$$\phi = (H_k^T H_k + \delta I)^{-1}(s_k^T H_k), \tag{8}$$

for some small $\delta \in \mathbb{R}_+$. The collection of past, current, and one time step in the future human inputs is given by $s_{k+1} = [v_{k-N_s+1}, \ldots, v_k + 1]^T = G s_k$, where

$$G = \begin{bmatrix} 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \\ \phi_{N_s} & \phi_{N_s-1} & \phi_{N_s-2} & \cdots & \phi_1 \end{bmatrix} \in \mathbb{R}^{N_s \times N_s}.$$

From the last row, we have the predicted human input one time step into the future, i.e., $v_{k+1} = \phi s_k$.

Similarly, the predicted human input two time steps in the future is obtained using $s_{k+2} = G s_{k+1}$, with $v_{k+2} = \phi s_{k+1}$. Repeating this procedure $N_k - 1$ times, the future human input sequence is obtained; $\mathcal{V}_k = \{v_k, v_{k+1}, \ldots, v_{k+N_k}\} = \{v_k, \phi s_k, \ldots, \phi s_{k+N_k-1}\}$.[2]

B. Adaptive Prediction Horizons

As there is an inherent trade-off between prediction horizon and prediction quality, we may want to be able to adaptively adjust the horizon as a function of the prediction quality. And, in order to measure the performance of the human prediction at time k, we propose to utilize the system identification model obtained at the current time and produce a human input signal backwards in time for the length of the current control horizon. The performance measure is a cost on the deviation from this signal to the actual human input signal recorded over that time. If the deviations are large, then the predicted model is not accurate (i.e. not performing well) and the horizon should be shortened. On the other hand, if the deviations are small, then

[2] When using SID, there is always an initialization period required to accumulate past human input values before the actual system identification can commence. We use ZOH for this initial period in the experiments.

the human input model is performing well and the horizon can be increased.

In light of this discussion, we will try to find the horizon that minimizes the deviations in the predicted and actual human input signals based on past data. For more details on this method for choosing control horizons, see [26]. However, as the particulars of this design method are not fundamental to the developments in this paper, we simply refer the reader to [26]. What is, however, of fundamental importance is the manner in which the human operator responds to varying time horizons. And, although they are designed to improve the performance of the MPC controller, it will in fact turn out that the human operator inputs change significantly in response to varying horizons, making the usefulness of this method somewhat dubious in contexts where the "environment", i.e., the human inputs, is coupled to the prediction horizon.

C. Solving the MPC

The framework for mixed initiative interactions proposed in this paper utilizes model predictive control and this involves solving an optimal control problem at each time instant. This is potentially computationally quite intensive, and in some instances even infeasible. Therefore, if a closed-form solution could be found this would significantly improve the usability of the proposed methodology. The following section details such a closed-form controller for the specific case of norm-squared stage costs, linear dynamical systems, and linear state constraints, which corresponds to the scenario under investigation in the user-studies, detailed in the following sections.

If the robot dynamics are given by a linear, completely controllable system ($x_{k+1} = Ax_k + Bu_k$), the cost is given by

$$V_{N_k} = \sum_{i=k}^{k+N_k-1} \|v_i - u_i\|^2,$$

and the task is modeled by the constraint, $x(k+N) \in \mathbb{X}_f = \{x \mid Mx = b\}$, one can solve $\mathcal{P}_{MPC}$ analytically as long as the human input sequences are sufficiently regular. In particular, if the human input sequences belong to the Hilbert space of square-summable sequences[3] (denoted by $\mathcal{H}$), in [19], $\mathcal{P}_{MPC}$ was solved as a direct application of Hilbert's Projection Theorem, following the procedure in [27]. In particular, the optimal u_k^{opt} at time k is given by

$$u_k^{opt} = L_k^* M^T(MLL^*M^T)^{-1}(b - MA^{N_k}x_k - ML\mathcal{V}_k) + v_k, \tag{9}$$

where the linear operator, $L : \mathcal{H} \to \mathbb{R}^n$, and the adjoint operator, $L^* : \mathbb{R}^n \to \mathcal{H}$, are given by

$$\begin{aligned} L\mathcal{U}_k &= \sum_{i=k}^{k+N_k-1} A^{k+N_k-1-i}Bu_i, \\ L^* &= \{B^T(A^{N_k-1})^T, B^T(A^{N_k-2})^T, \ldots, B^T\}, \end{aligned}$$

and $L_k^* = B^T(A^{N_k-1})^T$.[4]

As such, we have a closed-form solution to a version of the optimal control problem that needs to be solved every time step, instead of having to numerically solve a potentially computationally intensive constrained quadratic program. In the subsequent section, we will use the closed-form solution when evaluating the different design choices in the experimental operator trial.

IV. Human Operator Study

A. Experimental Considerations

The purpose of the human operator studies is multi-faceted in that we not only want to gauge the effectiveness of the lower-level control in an experimental setting, but, more importantly, we investigate whether or not human operators are afforded the freedom required to accomplish higher-level tasks. In addition, we would like to measure overall task performance of different versions of the controller versus manual control as well as any operator workload differences. And, as already hinted at, it turns out that due to the complex and unknown coupling between the user behaviors and the prediction methods used, the most effective mixed initiative MPC strategy corresponds to the simplest one. In fact, the versions of the prediction methods used are Zero-Order Hold with Fixed Prediction Horizon (ZOH), Least Squares System Identification with Fixed Prediction Horizon (FSID), Least Squares System Identification with Variable Prediction Horizon (VSID), and pure, Manual Control (Manual).

The experimental scenario under consideration is inspired by a search and rescue operation navigation task where three points of interest are given (where potential victims may be) before the task begins. The automatic controller is commanded to drive the robot to any one of these points (low-level task), while the human operator is to guide the robot to the points in the order deemed appropriate by the operator. The human also has the power to influence the path taken by the robot to each of these points. During the task, the human is also asked to identify a possible, new area of interest (where there could be additional victims) on the way between two of the predefined points. The human must then actively alter the robot path to visit this point. This scenario requires sliding autonomy in that the human involvement ranges from a little to a lot of operator interaction with the automatic controller.

The search and rescue-inspired scenario is depicted in Figure 2, where the predefined goal points are labeled as Goal 1, Goal 2, and Goal 4, whereas the mid-task goal-point is labeled Point 3. The low-level task is to ensure that the robot does indeed reach one of the predefined goal points (Goals 1, 2, and 4). The high-level task consists of choosing in which order the goals are visited as well as visiting the goal point not predefined as a goal (Point 3) in-between Goal 2 and Goal 4. The operator is situated in the same room as the task environment and has full view of work environment, as shown in Figure 3. The laboratory environment does not simulate

[3]Note that technically speaking, we need different Hilbert spaces for different prediction horizons, but this does not change anything about the solution since the optimization problem is resolved at every time instant.

[4]Note that the key feature of $\mathcal{P}_{MPC}$ that enables this closed-form solution is that the stage cost L does not depend on x, which is why no Riccati Equations need to be solved.

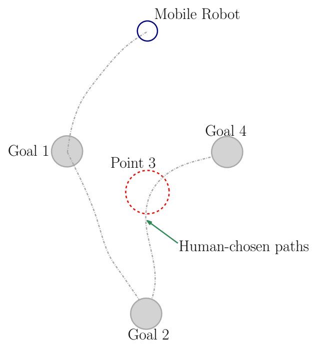

Fig. 2. An example of the shared control of a mobile robot navigation task. The automated controller drives the robot to the goal points (Goal 1, Goal 2, Goal 3), while the user specifies the order the goals are visited and influences the robot path so as to visit Point 3 on the way to the final goal point.

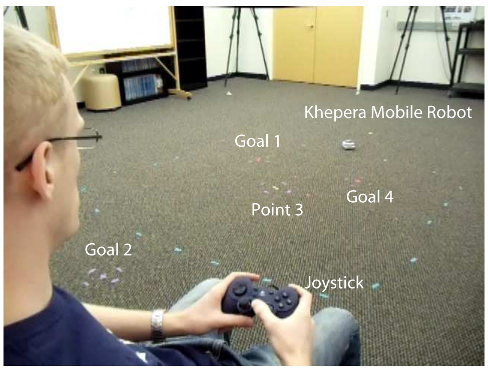

Fig. 3. Navigation Task Evironment with Khepera Mobile Robot. The joystick shown is used as the operator interface.

the same robot mobility, situational awareness, and workload challenges found in actual search and rescue operations, but it successfully serves the purpose of requiring the human-robot team to operate at different levels of autonomy, as well as provide information about the efficacy of the different prediction schemes.

The experiments aree conducted using a differential drive mobile robot (Khepera III) wirelessly receiving commands from a Ubuntu PC running the Robotic Operating System ([28]). Localization information is supplied by a VICON motion capture system giving planar position and orientation.

A low-level unicycle controller (for example see [29]) that takes in planar change-in-position commands and outputs velocity and angular velocity commands allows us to use a discrete-time, linear control system model to command the mobile robot:

$$x_{k+1} = x_k + u_k \tag{10}$$

where $x = (x_1, x_2)^T$ are the planar Cartesian coordinates of the robot, which provides the discrete system dynamics needed to formulate $\mathcal{P}_{MPC}$.

The low-level task that must be solved by the controller corresponds to ensuring that the robot reaches one of the three goals, each of which is modeled as a linear constraint, $x_k = b_i$, for $i = 1, 2, 4$, where b_i is the planar goal location for goal $i,\ i = 1, 2, 4$. (Note that we index these 1,2,4 since those are the four predefined goal points.) Hence, for $\mathbb{X}_i = \{x \mid x = b_i\}$, the constraint set in the optimal control problem is

$$\mathbb{X}_f = \mathbb{X}_1 \cup \mathbb{X}_2 \cup \mathbb{X}_4. \tag{11}$$

This goal set is a union of three linear sets and is, as such, not a linear set. However, what we do in practice is solve three optimal control problems analytically – resulting in three different candidate control values – and then select the control signal that best matches the human input. And, once a goal point is reached, the operator is given full manual control to stay at this goal point or move on to another goal point. These human issued commands are supplied by way of a video game-like gamepad with joysticks, shown in Figure 3. The joystick allows the operator to issue change-in-position commands in the global frame without regard to the orientation of the robot.

B. Results

The user study procedure employed is detailed in Appendix B and in this section, we report on the findings from said user study. The robot is defined as successfully reaching a goal point if its position is recorded as being within a 10 cm ball around the goal. A total of 40 trials were conducted, and in all 40 trials, the robot successfully reached all three of the goal points, demonstrating low-level task completion. More importantly, the operators in every trial were able to guide the robot in the specified goal order given by the test administrator as well as visiting Point 3. As a result, it can be concluded that operators and the robot were able to complete both low-level and high-level functions with the manual control as well as complete high-level functions while the three different versions of the mixed-initiative controller ensured low-level task completion. The repeated-measures approach to these experiments isolates the effects of the different controllers without the effects of operator-to-operator variability.

In order to evaluate any performance and workload differences between manual control and the three versions of the proposed controller, we analyzed the total task completion time (i.e. time to reach Goal 4) and NASA TLX workload data. Figure 4 shows the task completion times for each participant with each of the four controllers. The ZOH controller resulted in the fastest completion times across *all* participants. The mean completion times for each controller are as follows: 46.5 sec for ZOH, 60.8 sec for Manual, 76.3 sec for FSID, and 77.1 sec for VSID. As such, it is clear that the least complex of the mixed initiative MPC outperformed the more complex variants as well as the pure, manual controller. The resulting t and p-values for each pairwise test were ZOH versus Manual ($t(10)= 3.071$, $p=0.0133$), ZOH versus FSID ($t(10)= 5.798$, $p=0.0002$), and ZOH versus VSID ($t(10)= 5.056$, $p=0.0007$). Using the convention that p-values less than 0.05 are deemed statistically significant, we can see that the ZOH controller statistically significantly out-performs the other controllers.

Workload is measured using the NASA TLX survey and our analysis is carried out over the total raw NASA TLX scores. The raw total scores for all participants are plotted in Figure 5 where we see a trend that the workload scores for the ZOH controller tend to be less compared to the other controllers. The mean raw scores for each controller are as follows: 30.7

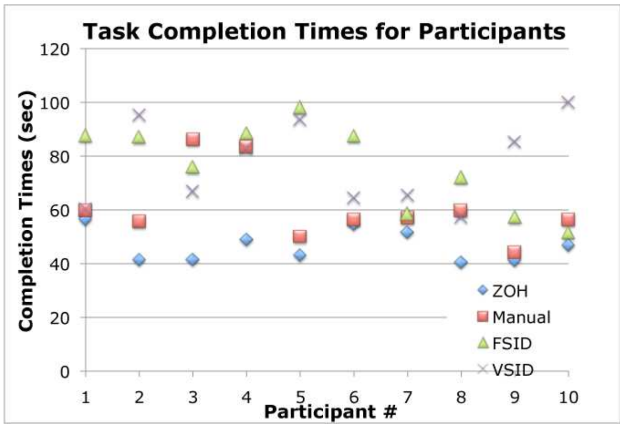

Fig. 4. Task completion times for all participants

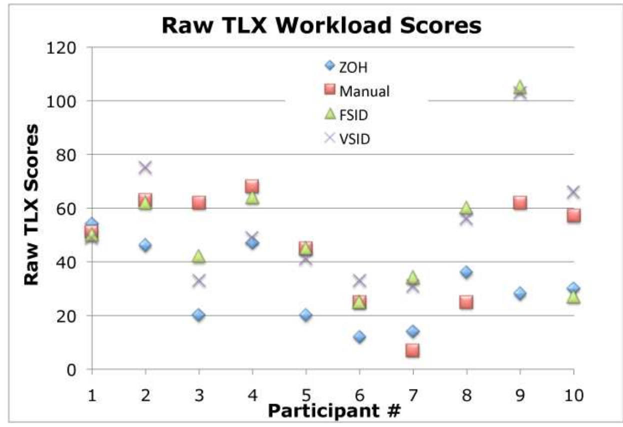

Fig. 5. Raw Total NASA TLX Workload survey scores for all participants.

for ZOH, 51.4 for Manual, 51.4 for FSID, and 53.6 for VSID. A repeated measures pairwise t-test was performed for the ZOH controller against the other three controllers to test the hypothesis in this paper. The resulting t-value and p-values for each pairwise test are ZOH versus Manual ($t(10)= 3.898$, $p=0.0036$), ZOH versus FSID ($t(10)= 2.938$, $p=0.0165$), and ZOH versus VSID ($t(10)= 3.318$, $p=0.0089$). Hence, the ZOH controller was shown to statistically significantly have a lower operator workload than manual control and the other prediction methods. The results of these human studies have shown that not only will the mixed-initiative control scheme guarantee low-level task completion, but we have shown that human operators have the freedom to accomplish high-level tasks with benefits to performance and operator workload in this particular experiment.

The reader may note that the TLX scores for Participant 1 are very close and a look at that participant's completion time will show that except for the FSID controller, those times are close as well. One interpretation of these results is that this operator is more skilled at the task and hence the assistance from the autonomous controller has less effect. In addition, the NASA TLX scores are highly dependent on the participant's individual viewpoint of the amount of workload, so scores may vary greatly between participants as seen in the lower overall scores of Participants 5-8. This dispersion in TLX scores has been accounted for by the repeated measures experiments conducted in this experiment as the aforementioned statistical analysis focuses on the differences in performance of each controller for a particular subject, rather than the absolute scores across all participants and controllers.

Finally, the final survey asked the operators to choose which of the controllers they would prefer to use again, which was the most frustrating to use, and which controller did they trust the most. The results of the survey are shown in Figure 6. The ZOH controller was both the most preferred and trusted controller while the VSID controller was the most frustrating. ZOH, along with FSID, was the least frustrating to use. It can also be seen that no one preferred or trusted Manual. In summary, less is indeed more in this particular context.

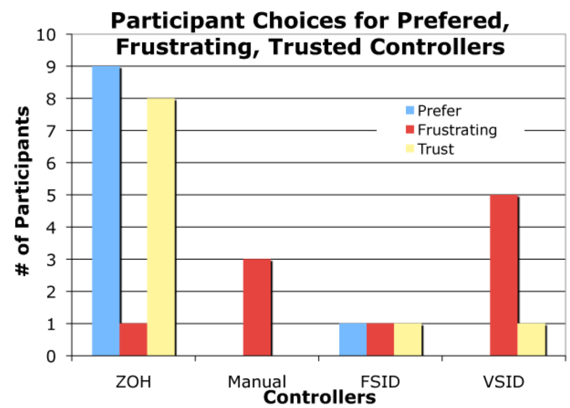

Fig. 6. Final Survey results showing number of participants indicating the controllers they preferred, thought were the most frustrating to use, and thought they could trust the most.

V. Conclusions

In this paper, an MPC controller is presented that combines human input signals with an automatic control signal to produce a controller with naturally sliding levels of autonomy. While theoretical results give low-level task completion guarantees, experimental results with human operators show that the control scheme allows for high-level task completion with specific benefits to performance and operator workload for a specific search-and-rescue-motivated mobile robot navigation task. These benefits were shown to be statistically significant with the Zero-Order Hold with Fixed Horizon version of the controller.

Acknowledgment

This work was supported by the US National Science Foundation through Grant # 0820004. The authors wish to thank Tiffany Chen and Dr. Brian Levinthal for helpful discussions.

References

[1] M. A. Goodrich and A. C. Schultz, "Human robot interaction: A survey," *Foundations and Trends in HumanComputer Interactio*, vol. Vol. 1: No 3, pp. pp 203–275, 2007.

[2] J. Desai, J. Ostrowski, and V. Kumar, "Controlling formations of multiple mobile robots," in *Robotics and Automation, 1998. Proceedings. 1998 IEEE International Conference on*, vol. 4. IEEE, 2002, pp. 2864–2869.

[3] G. Wasson, J. Gunderson, S. Graves, and R. Felder, "Effective shared control in cooperative mobility aids," in *In Proceedings of the Fourteenth international Florida Artificial intelligence Research Society Conference (May 21 - 23*. AAAI Press, 2001, pp. 509–513.

[4] P. Griffiths and R. B. Gillespie, "Shared control between human and machine: haptic display of automation during manual control of vehicle heading," in *Proc. 12th International Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems HAPTICS '04*, 27–28 March 2004, pp. 358–366.

[5] D.-J. Kim and A. Behal, "Human-in-the-loop control of an assistive robotic arm in unstructured environments for spinal cord injured users," in *HRI '09: Proceedings of the 4th ACM/IEEE international conference on Human robot interaction*. New York, NY, USA: ACM, 2009, pp. 285–286.

[6] S. G. Tzafestas and E. S. Tzafestas, "Human–machine interaction in intelligent robotic systems: A unifying consideration with implementation examples," *J. Intell. Robotics Syst.*, vol. 32, no. 2, pp. 119–141, 2001.

[7] C. Wang, H. Ma, and D. J. Cannon, "Human-machine collaboration in robotics: Integrating virtual tools with a collision avoidance concept using conglomerates of spheres," *J. Intell. Robotics Syst.*, vol. 18, no. 4, pp. 367–397, 1997.

[8] R. Wegner and J. Anderson, "Agent-based support for balancing teleoperation and autonomy in urban search and rescue," *Int. J. Robot. Autom.*, vol. 21, no. 2, pp. 120–128, 2006.

[9] D. J. Bruemmer, R. L. Boring, D. A. Few, J. L. Marble, and M. C. Walton, "i call shotgun!: An evaluation of mixed-initiative control for novice users of a search and rescue robot," in *In Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, 2004, pp. 2847–2852.

[10] M. Burstein, G. Ferguson, and J. Allen, "Integrating agent-based mixed-initiative control with an existing multi-agent planning system," in *Proc. Fourth International Conference on MultiAgent Systems*, 2000, pp. 389–390.

[11] S. Anderson, S. Peters, T. Pilutti, and K. Iagnemma, "An optimal-control-based framework for trajectory planning, threat assessment, and semi-autonomous control of passenger vehicles in hazard avoidance scenarios," *International Journal of Vehicle Autonomous Systems*, vol. 8, no. 2, pp. 190–216, 2010.

[12] D. Bruemmer, D. Few, R. Boring, J. Marble, M. Walton, and C. Nielsen, "Shared understanding for collaborative control," *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, vol. 35, no. 4, pp. 494–504, 2005.

[13] J. Kofman, X. Wu, T. Luu, and S. Verma, "Teleoperation of a robot manipulator using a vision-based human-robot interface," *Industrial Electronics, IEEE Transactions on*, vol. 52, no. 5, pp. 1206–1219, 2005.

[14] Y. Horiguchi and T. Sawaragi, "Effects of probing to adapt machine autonomy in shared control systems," in *Proc. International Conference on Systems, Man and Cybernetics*, vol. 1, 2005, pp. 317–323.

[15] S. Parikh, V. Grassi Jr, V. Kumar, and J. Okamoto Jr, "Usability study of a control framework for an intelligent wheelchair," in *Robotics and Automation, 2005. ICRA 2005. Proceedings of the 2005 IEEE International Conference on*. IEEE, 2005, pp. 4745–4750.

[16] S. McLachlan, J. Arblaster, O. Liu, J. Miro, and L. Chenoweth, "A multi-stage shared control method for an intelligent mobility assistant," in *Rehabilitation Robotics, 2005. ICORR 2005. 9th International Conference on*. IEEE, 2005, pp. 426–429.

[17] C. Urdiales, J. Peula, M. Fdez-Carmona, C. Barrué, E. Pérez, I. Sánchez-Tato, J. del Toro, F. Galluppi, U. Cortés, R. Annichiaricco *et al.*, "A new multi-criteria optimization strategy for shared control in wheelchair assisted navigation," *Autonomous Robots*, pp. 1–19, 2011.

[18] S. G. Loizou and V. Kumar, "Mixed initiative control of autonomous vehicles," in *Proc. IEEE International Conference on Robotics and Automation*, Apr. 10–14, 2007, pp. 1431–1436.

[19] R. Chipalkatty and M. Egerstedt, "Human-in-the-Loop: Terminal constraint receding horizon control with human inputs," in *Robotics and Automation (ICRA), 2010 IEEE International Conference on*. IEEE, 2010, pp. 2712–2717.

[20] R. Chipalkatty, M. Egerstedt, and S. Azuma, "Multi-pendulum synchronization using constrained agreement protocols," in *International Conference on Robot Communication and Coordination*, 2009.

[21] G. Goodwin, M. Seron, and J. D. Dona, *Constrained Control and Estimation*, E. Sontag, Ed. Springer, 2005.

[22] P. O. M. Scokaert, D. Q. Mayne, and J. B. Rawlings, "Suboptimal model predictive control (feasibility implies stability)," *IEEE Transactions on Automatic Control*, vol. 44, no. 3, pp. 648–654, 1999.

[23] R. Findeisen and F. Allgower, "An introduction to nonlinear model predictive control," in *In 21st Benelux Meeting on Systems and Control*, 2002.

[24] D. Mayne, J. Rawlings, C. Rao, and P. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, p. 789814, 2000.

[25] L. Ljung and E. Ljung, *System identification: theory for the user*. Prentice-Hall Englewood Cliffs, NJ, 1987.

[26] G. Droge and M. Egerstedt, "Adaptive time horizon optimization in model predictive control," in *American Control Conference (ACC), 2011*. IEEE, 2011, pp. 1843–1848.

[27] Y. Zhou, M. Egerstedt, and C. Martin, "Hilbert space methods for control theoretic splines: A unified treatment," *Communications in Information and Systems*, vol. 6, No. 1, pp. 55–62, 2006.

[28] M. Quigley, B. Gerkey, K. Conley, J. Faust, T. Foote, J. Leibs, E. Berger, R. Wheeler, and A. Ng, "Ros: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.

[29] M. Aicardi, G. Casalino, A. Bicchi, and A. Balestrino, "Closed loop steering of unicycle like vehicles via lyapunov techniques," *Robotics & Automation Magazine, IEEE*, vol. 2, no. 1, pp. 27–35, 1995.

APPENDIX

A. Proof of Theorem 2.1

The proof is based on showing that the value function goes to zero as k goes to infinity by proving that the value function is bounded above by a value function that converges to zero, for both the expanding, constant, and contracting horizon cases.

We will use $\mathcal{U}^{opt}_{k,N_k}$ as shorthand for the optimal control sequence of length N_k at time k, and note that this sequence really is a function of the human inputs $\mathcal{V}_{k,N_k}$ over the same horizon as well as the state of the system x_k at time k. Similarly, we will let $\mathcal{U}^{fea}_{k,N_k}$ denote a feasible (not necessarily optimal) control sequence.

Now, by Assumption A4, there exists a horizon, M, such that $M \leq N_k$ for all k, i.e. all horizons are bounded from below by M. At time k with human input $\mathcal{V}_{k,M} = \{v_k, v_{k+1}, \ldots, v_{k+M-1}\}$, let the following be an optimal control and state sequence using the horizon, M,

$$\begin{aligned} \mathcal{U}^{opt}_{k,M} &= \{u^{opt}_{k,M}, \ldots, u^{opt}_{k+M-1,M}\} \\ \mathcal{X}^{opt}_{k,M} &= \{x^{opt}_{k,M}, \ldots, x^{opt}_{k+M-1,M}, x^{opt}_{k+M,M}\}, \end{aligned}$$

with $x^{opt}_{k+M,M} \in \mathbb{X}_f$.

Then, at time $k+1$, a feasible control and state sequence for the reference horizon M, given human input $\mathcal{V}_{k+1,M} = \{v_{k+1}, v_{k+2}, \ldots, v_{k+M}\}$ is

$$\begin{aligned} \mathcal{U}^{fea}_{k+1,M} &= \{u^{opt}_{k+1,M}, \ldots, u^{opt}_{k+M-1,M}, v_{k+M}\} \\ \mathcal{X}^{fea}_{k+1,M} &= \{x^{opt}_{k+1,M}, \ldots, x^{opt}_{k+M-1}, x^{opt}_{k+M,M}, x_{k+M+1}\}, \end{aligned}$$

with (by Assumption A3) $x^{opt}_{k+M,M}, x_{k+M+1,M} \in \mathbb{X}_f$, and where $x_{k+M+1} = f(x^{opt}_{k+M,M}, v_{k+M})$. For both the expanding and contracting horizon cases (of which the constant horizon is a special case), we will show that the optimal costs are bounded above by this feasible cost over the horizon M, and that this feasible cost converges to zero as $k \to \infty$. We do this by showing that a feasible cost over horizon N_k can be constructed whose value is equal to the feasible cost over horizon M. And, by necessity, the optimal cost is bounded by this vanishing, feasible cost.

Case 1: $N_k \leq N_{k+1}$ (Expanding Horizon)

A feasible control and state sequence for the human input $\mathcal{V}_{k,N_k}$ at time k is given by a concatenation of two sequences:

$$\begin{aligned} \mathcal{U}^{fea}_{k,N_k} &= \mathcal{U}^{opt}_{k,M} \cdot \{v_{k+M}, \ldots, v_{k+N_k-1}\} \\ \mathcal{X}^{fea}_{k,N_k} &= \mathcal{X}^{opt}_{k,M} \cdot \{x_{k+M+1}, \ldots, x_{k+N_k}\}, \end{aligned}$$

where the second sequence in $\mathcal{X}^{fea}_{k,N_k}$ is entirely contained in $\mathbb{X}_f$, as per Assumption A3.

If we now advance time one step to $k+1$, the corresponding feasible sequences $\mathcal{U}^{fea}_{k+1,N_{k+1}}$ and $\mathcal{X}^{fea}_{k+1,N_{k+1}}$ are obtained by removing the first element from the feasible sequences at time k and then concatenating these sequences with $\{v_{k+N_k}, \ldots, v_{k+N_{k+1}-1}\}$ and $\{x_{k+N_k+1}, \ldots x_{k+N_{k+1}}\}$ respectively. Since the new states stay in the target set and the last control signals are equal to the human inputs, no extra cost is incurred, as per Assumptions A1-A3. Therefore

$$V_M(\mathcal{V}_{k+1,M}, \mathcal{U}^{fea}_{k+1,M}) = V_{N_{k+1}}(\mathcal{V}_{k+1,N_{k+1}}, \mathcal{U}^{fea}_{k+1,N_k+1}),$$

which establishes the bound in the expanding horizon case.

Case 2: $N_k \geq N_{k+1}$ (Contracting Horizon)

This part of the proof is similar and it also involves the construction of the same suitable, feasible sequences at time k as before $\mathcal{U}^{fea}_{k,N_k}$ and $\mathcal{X}^{fea}_{k,N_k}$. As the horizon is contracting, the difference now is with the construction of the subsequent sequences, where instead of the sequence length increasing, it is now decreasing due to the contracting horizon, and $\mathcal{U}^{fea}_{k+1,N_{k+1}}$ and $\mathcal{X}^{fea}_{k+1,N_{k+1}}$ are again obtained by removing the first element from the feasible sequences at time k and then, just as before, concatenating these sequences with $\{v_{k+N_k}, \ldots, v_{k+N_{k+1}-1}\}$ and $\{x_{k+N_k+1}, \ldots x_{k+N_{k+1}}\}$ respectively. Again, due to Assumptions A1-A4,

$$V_M(\mathcal{V}_{k+1,M}, \mathcal{U}^{fea}_{k+1,M}) = V_{N_{k+1}}(\mathcal{V}_{k+1,N_{k+1}}, \mathcal{U}^{fea}_{k+1,N_k+1}).$$

What remains to show is that the upper bound given by the fixed horizon feasible sequence (over horizon M) does indeed have a cost that converges to zero. Since

$$V_{N_k}(\mathcal{V}_{k,N_k}, \mathcal{U}^{opt}_{k,N_k}) \leq V_{N_k}(\mathcal{V}_{k,N_k}, \mathcal{U}^{fea}_{k,N_k}) = V_M(\mathcal{V}_{k,M}, \mathcal{U}^{fea}_{k,M}) \tag{12}$$

convergence would have been established if we can show that $V_M(\mathcal{V}_{k,M}, \mathcal{U}^{fea}_{k,M}) \to 0$ as $k \to \infty$. But, once the target set has been reached, this feasible control sequence lets the control inputs be given by the human inputs. And, per Assumption A3, the corresponding cost is zero, from which the convergence result is established, following the argument in [21] and [22]. ■

B. User Study Procedure

Ten operators were recruited for the experiment from the Georgia Institute of Technology community. None of the operators have had previous experience with mobile robot control. The participants were between the ages 20-30 with 3 female participants and 7 male participants. The participants read standard written instructions on the task and then physically shown the task environment with verbal instructions on how to complete the navigation task. The participants were first allowed to practice the task using only manual control for a maximum of three times. Then, the participants were given a training session with each of the controllers before performing the task with recorded data. Each training session consisted of a maximum of three attempts at the given task. Each recorded run was followed by a NASA TLX workload survey. The order of the four controllers were counter-balanced to account for any ordering effects. After the four trials, the participants were given an exit survey comparing the four controllers.

For each trial, the participant was asked to drive the robot to Goal 1, then Goal 2, then to pass through Point 3, on the way to Goal 4 as seen in Figure 2. The ordering of the goals was set before the trials by the study administrator and were the same for every participant. The controller was programmed with Goal 1, Goal 2, and Goal 4 given a priori however, Point 3 was not part of the low-level task. Moreover, the controller was not given the order in which the goals must be visited. In this way, human control naturally shifts from a supervisory type of control towards that of a more manual control to visit Point 3.