

# A Bayesian Optimization Framework for the Automatic Tuning of MPC-based Shared Controllers

Anne van der Horst<sup>1\*</sup>, Bas Meere<sup>1\*</sup>, Dinesh Krishnamoorthy<sup>1</sup>,  
Saray Bakker<sup>1,2</sup>, Bram van de Vrande<sup>3</sup>, Henry Stoutjesdijk<sup>3</sup>, Marco Alonso<sup>3</sup> and Elena Torta<sup>1</sup>

**Abstract**—This paper presents a Bayesian optimization framework for the automatic tuning of shared controllers which are defined as a Model Predictive Control (MPC) problem. The proposed framework includes the design of performance metrics as well as the representation of user inputs for simulation-based optimization. The framework is applied to the optimization of a shared controller for an Image Guided Therapy robot. VR-based user experiments confirm the increase in performance of the automatically tuned MPC shared controller with respect to a hand-tuned baseline version as well as its generalization ability.

## I. INTRODUCTION

Nowadays, robots are deployed in increasingly complex environments, where the variability in tasks can be high. For such robotics applications, human judgement is preferred over full autonomous decision making due to humans' ability to adapt to task diversity [1]. To this end, shared control [2] is a favorable framework, as it combines the benefits of robotic automation with the flexibility and decision-making skills of humans. The performance of such a control algorithm often depends on the tuning of its parameters. The tuning process notoriously takes the form of a trial-and-error approach which is time-consuming, error-prone, and can become frustrating. Bayesian Optimization (BO) [3] is a data-based optimization technique that has surged in popularity as a viable way of automatically tuning controller parameters for robotic systems (e.g., [4]–[6]). Contrary to Reinforcement Learning, BO is a gradient-free method and typically reaches convergence with fewer iterations. Both aspects are important to incorporate direct user feedback in the tuning process of algorithms for human-robot interaction tasks.

While previous work has addressed the development of BO algorithms for autonomous motion tasks of manipulators [5], [6] and drones with emphasis on safety constraints [4], [7], little attention has been given to controller optimization for human-robot interaction tasks, such as shared control.

The contribution of this paper relates to the definition of a BO framework for the automatic tuning of shared controllers for robot navigation. In addition to the design choices of the algorithm itself, distinctive challenges relate to the definition of representative user inputs and the identification of the performance metrics for the optimization process.

\*Equal contribution.

<sup>1</sup> Eindhoven University of Technology, The Netherlands (e-mail: a.v.d.horst1@tue.nl)

<sup>2</sup> Delft University of Technology, The Netherlands

<sup>3</sup> Philips IGT Systems Mechatronics, The Netherlands




Fig. 1: The IGT robot with its prismatic joint  $q_1$ , two revolute joints  $q_2$  and  $q_3$  and the end-effector position. The 12 spheres for obstacle avoidance capture the relevant geometry of the robot.

The proposed BO framework is applied to the optimization of an MPC-based shared controller for an Image Guided Therapy (IGT) robot, shown in Fig. 1. These robots are used to create medical images via X-ray during minimally invasive procedures and are constituted by several links with the terminal link presenting the classical C-shape carrying the X-ray generator and detector. In the current state-of-practice, the (re)positioning of the robot in the operating room requires inputs from the physician through a joystick-like control module. The development of a shared controller is expected to facilitate positioning because the robot can automatically avoid obstacles while following users' commands as closely as possible. The motion performance of the automatically-tuned MPC shared controller is compared with a baseline, hand-tuned controller, by means of Virtual Reality (VR) experiments involving users interacting with the system through a joystick-like control module. The comparison is based on metrics that measure the safety and smoothness of the trajectories as well as the efficiency of the controller. Results show an overall performance improvement of the automatically tuned shared controller w.r.t. the baseline as well as a good generalization for user input sequences not used during optimization.

*Related Work:* BO has been successfully applied for the automatic tuning of controllers for several robotic tasks. One of the earlier examples is the work of Marco et al. [6] that optimized an LQR controller for an inverted pendulum actuated by a robotic arm. It showed the successful application of BO to a physical robotic system. Subsequent work has focused on safe BO [4], [7] which aims at enabling a safe exploration strategy of the algorithm's tuning parameters.

The latter is particularly relevant for the application of the optimization framework on physical systems or with direct user's involvement. Recent applications of BO to robotic systems also investigated its relation to the selection of rewards for reinforcement learning algorithms (e.g., [8], [9]) as well as the automatic tuning of MPC-based controllers (e.g., [10]–[12]).

Edwards et al. [10] devised `AutoMPC`; a software package to jointly identify the underlying system dynamics and automatically tune the parameters of data-driven MPC controllers for robotics applications. `AutoMPC` was validated in an OpenAI Gym environment. An example involving the application of BO to the automatic tuning of an MPC controller on a scaled model of a racing car is proposed in [11].

Semi-autonomous navigation under shared control requires a seamless integration of inputs from a human operator and an autonomous controller. To this end, MPC is an advantageous technique as it can incorporate constraints on a robot's actuation limits as well as predictions of future system's states and user's commands. Several examples of MPC-based shared controllers for semi-autonomous robot navigation are proposed in the literature (e.g., [13]–[15]). In [13], the controller is employed to blend human input and autonomous behavior to enable the autonomous navigation of a mobile robot in a cluttered environment. Another example of an MPC-based shared controller is given in the experimental study of Zarei et al. [14], where a mobile robot could navigate semi-autonomously by minimizing the deviations from a human input while upholding obstacle avoidance constraints. Taking one step further, Chipalkatty et al. [15] design the control problem such that it includes the system dynamics to minimize deviations from a human input while safeguarding a linear state constraint. Since the current paper aims to apply the shared controller to a robot that resembles a mobile manipulator, we designed the controller based on an adaptation of the perceptive MPC architecture for mobile manipulators presented in [16].

## II. SHARED CONTROL FRAMEWORK

The shared controller for robot semi-autonomous navigation is designed to optimize the trade-off between following user commands as closely as possible while autonomously avoiding obstacles. The architecture proposed in [16] has been adapted for a navigation task under shared control by considering a target velocity provided by the user. The block diagram of the resulting architecture is illustrated in Fig 2. User input is gathered from a joystick-like control module in the form of a reference velocity vector, i.e.,  $\dot{\mathbf{x}}_d$ , which represents the velocity the robot is required to track. The reference velocity is applied to the robot's end-effector which, specifically for the IGT robot, is defined as the midpoint between the X-ray generator and detector (see Fig. 1).

A voxel map of the environment is created which is subsequently converted into a Euclidean Signed Distance Field (ESDF) [17] that reports, for every voxel, the distance to the closest obstacle. For computational efficiency, the entire geometry of the robot is approximated using a set of spheres




Fig. 2: The control architecture of the proposed MPC-based shared controller.

(as depicted in Fig. 1). The number of spheres for the specific IGT robot,  $n_{cs}$ , was 12, all with a radius  $r_{cs} = 0.4$  m. Relying on the gradient of the ESDF, it is possible to derive the preferred direction to move each sphere away from obstacles. The control effort, in the form of joints' velocities, is computed by solving a non-linear MPC problem whose expression is given by the following equations:

$$\min_{\mathbf{x}, \mathbf{u}} \alpha \sum_{k=0}^{N_c-1} l(\mathbf{x}(k), \mathbf{u}(k)) + (1 - \alpha) \sum_{k=1}^{N_p} \sum_{m=1}^{n_{cs}} B_m(\mathbf{x}(k)), \quad (1)$$

$$\text{s.t. } \mathbf{x}(k+1) = \mathbf{x}(k) + T_s f(\mathbf{x}(k), \mathbf{u}(k)), \quad (2)$$

$$\mathbf{x}(0) = \mathbf{x}_0, \quad (3)$$

$$\mathbf{h}(\mathbf{x}(k), \mathbf{u}(k)) \geq 0. \quad (4)$$

The state vector is defined as  $\mathbf{x}(k) = [\mathbf{x}_e(k) \quad \mathbf{q}(k)]^T$  with  $\mathbf{x}_e(k)$  representing the 2D pose of the end-effector and  $\mathbf{q}(k)$  representing the vector of the positions of the joints. For the IGT robot,  $\mathbf{q}(k)$  represents the positions of the prismatic joint  $q_1$  and the revolute joints  $q_2$  and  $q_3$  (see Fig. 1). The control effort,  $\mathbf{u}(k)$ , is chosen as the velocity of the joints. The stage cost in Eq. 1 includes a term to penalize the deviation of the end-effector's Cartesian velocity from the joystick input provided by the user:

$$l(\mathbf{x}(k), \mathbf{u}(k)) = \|\mathbf{J}_e(\mathbf{x}(k))\mathbf{u}(k) - \dot{\mathbf{x}}_d(k)\|_{\mathbf{Q}}^2, \quad (5)$$

where  $\mathbf{J}_e$  is the geometric Jacobian matrix of the robot and  $\mathbf{Q} \in \mathbb{R}^{3 \times 3}$  is a positive definite weight matrix used to weight the squared norm of the error,  $\|\mathbf{x}\|_{\mathbf{Q}}^2 = \mathbf{x}^T \mathbf{Q} \mathbf{x}$ . The reference input  $\dot{\mathbf{x}}_d(k)$  is assumed constant over the prediction horizon  $N_p$ . The other term of the stage cost penalizes joint configurations that bring the collision spheres closer to obstacles. Its expression is given by:

$$B_m(\mathbf{x}(k)) = \frac{c_1}{1 + e^{c_2(sd(\mathbf{x}(k)) - c_3)}}, \quad (6)$$

which is taken as a function of the state vector. Here,  $sd(\mathbf{x}(k))$  is the signed distance of the center of each collision sphere to the closest obstacle, and  $c_1, c_2$  and  $c_3$  are scaling factors. In this work,  $c_3 = r_{cs}$ . Individual contributions are summed in Eq. 1 over the entire prediction horizon  $N_p$ , where the matrix  $\mathbf{Q}$  from Eq. 5 and the scalar  $c_1$  from Eq. 6 normalize the contributions. The blending scalar, a typical element of shared controllers that specifies the priority between tracking user input and avoiding obstacles, is defined as  $\alpha$ , which is a scalar between 0 and 1. In this study,  $\alpha = 0.5$ , to assign equal value to the human input and obstacle avoidance term. In our formulation, the control horizon  $N_c$  can be different from the prediction horizon  $N_p$ .

Furthermore, Eq. 2 discretizes the forward kinematics of the system  $\forall k \in N_p$ , where  $T_s$  denotes the discrete sample time using the zero-order hold method [18]. The initial conditions are defined in Eq. 3, while Eq. 4 contains position, velocity, acceleration and jerk constraints on the joints of the robot as well as the end-effector. The constraints are summed  $\forall k \in N_p, N_c$ , depending on whether the constraint relies on states  $\mathbf{x}(k)$  or inputs  $\mathbf{u}(k)$ .

The choice of the MPC parameters for automatic tuning depends on several factors such as the robot's geometry and the task's specifications. Due to scalability issues of BO [3], [19], it is recommended to keep the amount of tuning parameters low. For the validation use-case, a sensitivity analysis with the elementary effect method [20] was performed which resulted in the selection of the following parameters; the prediction and control horizon  $N_p$  and  $N_c$ , the three elements of the  $\mathbf{Q}$  matrix in Eq. 5 that penalizes differences between reference end-effector velocities and computed velocities, and the two scaling variables of the obstacle avoidance term, i.e.,  $c_1$  and  $c_2$  of Eq. 6. To summarize, seven tuning parameters were selected, given by  $\boldsymbol{\xi} = [N_p, N_c, Q_x, Q_y, Q_\theta, c_1, c_2]$ . The choice of tuning parameters is in line with related works. For example, selecting the individual elements of the matrix  $\mathbf{Q}$  is quite common practice (e.g., [21], [22]). The selection of the control  $N_c$  and prediction horizon  $N_p$  is also reported in [23].

### III. BAYESIAN OPTIMIZATION FRAMEWORK

In this section, we describe the proposed setup of the BO framework for the shared controller. For background information on BO and its theoretical underpinning, the reader is referred to [3], [24].

We cast the BO problem as a multi-objective constraint optimization problem:

$$\begin{aligned} \min_{\boldsymbol{\xi}} J_{\boldsymbol{\xi}} &= n_{mov} \sum_{j=1}^{n_{mov}} \sum_{h=1}^{n_{obj}} w_h F_{hj}(\boldsymbol{\xi}), \\ \text{s.t.} \quad N_p &\geq N_c, \\ n_{succ}(\boldsymbol{\xi}) &= n_{mov}, \\ sd(\mathbf{x}) &\geq r_{cs}. \end{aligned} \quad (7)$$

The term  $\boldsymbol{\xi}$  represents the vector of tunable MPC parameters. We assume the objective function can be evaluated precisely only at the point  $\boldsymbol{\xi}_i$  selected for the  $i$ -th optimization iteration [3]. The objective function is thus approximated by a statistical surrogate model which commonly takes the form of a Gaussian Process [3], [24]. For this work, the kernel of the Gaussian Process is chosen as a Matérn 5/2 kernel [24]. The next sampling set,  $\boldsymbol{\xi}_i$ , is determined by maximizing the acquisition function of which an analytical expression is known. The acquisition function can take different expressions. For the validation use-case, we opted for the Expected Improvement (EI) [3]. The value of the objective function,  $J_{\boldsymbol{\xi}}$ , at the sampling point  $\boldsymbol{\xi}_i$ , can be evaluated by means of experiments in which the performance of the robot's motion is quantified with a set of  $n_{obj} = 6$  performance metrics (i.e.,  $F_{hj}(\boldsymbol{\xi})$ , where  $h$  denotes the metric index and  $j$  the

movement index). The proposed analytical formulation of  $F_{hj}(\boldsymbol{\xi})$  is described in detail in Section III-A. The relative importance of the metrics is expressed through the weights  $w_h$ . At every iteration,  $n_{mov}$  movements are performed with the MPC controller configured with the current parameter set  $\boldsymbol{\xi}_i$ . At every iteration of the optimization algorithm, the performance metrics are computed for each movement and summed over the total number of movements  $n_{mov}$ . For the validation setup reported in the paper, movements were performed in a simulation environment. Executing a movement in simulation requires the formulation of a user input which we elaborate on more extensively in Section III-B. The full procedure is reported in Algorithm 1. For

---

#### Algorithm 1 BO procedure for shared controller tuning.

---

- 1: Define a set of  $n_{mov}$  movements (see Section III-B)
  - 2:  $\boldsymbol{\xi} \leftarrow \boldsymbol{\xi}_0$  (initialize tuning parameters)
  - 3: **for**  $i = 0$  **to**  $n_{max}$  **do** (optimization iterations)
  - 4:   **for**  $j = 0$  **to**  $n_{mov}$  **do**
  - 5:     Perform a movement, evaluate metrics
  - 6:   Compute objective function for iteration  $i$  (Eq. 7)
  - 7:   Update the surrogate function
  - 8:   Update the acquisition function
  - 9:   Select a new value of  $\boldsymbol{\xi}$  by maximizing the acquisition function
  - 10: **return**  $\boldsymbol{\xi}$  with the smallest objective function
- 

the validation use-case, the procedure was implemented in Matlab using the Statistics and Machine Learning toolbox. Simulations were performed in Simulink Simscape.

#### A. Metrics definition

Based on prior work on the evaluation of mobile robot local planners [25] and pure-motion tasks of mobile robots [26], we identified  $n_{obj} = 6$  metrics by which to evaluate the performance of the MPC-based shared controller. Each metric relates to one of three fundamental aspects for human-robot interaction namely safety, smoothness and efficiency. *Safety* is arguably one of the most important aspects since shared control is often deployed in robots that operate in human-populated environments. We identified two metrics to evaluate safety, i.e., obstacle proximity and time spent close to obstacles.

The obstacle proximity metric, denoted as  $d_{ob}$ , serves as an indicator of safety as it represents the closest approach of the robot to a potential collision during its movement. The metric is calculated as the minimum distance between any link of the robot and obstacles throughout the entire motion of the robot:

$$F_1 : d_{ob} = \min_g \left\{ \min_m \{d_{g,m}\} \right\}, \quad 1 \leq g \leq N, \quad 1 \leq m \leq n_{cs}, \quad (8)$$

where  $g$  is the index of the sample,  $N$  is the total number of samples in a movement, and  $m$  is the index of the link. For the validation use-case, the links of the robots are approximated with the collection of spheres,  $n_{cs}$ , with the

distance measured from the center of each sphere.

The second safety metric,  $t_{ob}$ , quantifies the time that the robot spends within a certain distance from the obstacles during a movement. The metric is calculated as the fraction of time spent with any of the links of the robot closer than a threshold  $d_{safe}$  to obstacles. Its expression is given by:

$$F_2 : t_{ob} = \frac{\sum (t_b - t_a)}{t_F} \times 100\%, \quad (9)$$

where the subscripts  $a$  and  $b$  are the indices of timestamps satisfying  $d_s \leq d_{safe}$ ,  $a \leq k \leq b$ , and  $d_s$  representing the distance between any of the links and obstacles. For the validation use-case, we tailored  $d_{safe}$  to be equal to the radius of the collision spheres  $r_{cs}$ .

*Smoothness* metrics quantify the oscillations in a robot's trajectory. For shared controllers, oscillations often result from the trade-off between reference tracking and obstacle avoidance. We identified three smoothness metrics, i.e., end-effector path smoothness, curvature change and velocity smoothness. The end-effector path smoothness,  $f_{ps}$ , quantifies oscillations of the end-effector in Cartesian space and its expression is given by:

$$F_3 : f_{ps} = \frac{1}{S} \sum_{g=2}^{N-1} \|\Delta \mathbf{p}_{g+1} - \Delta \mathbf{p}_g\|^2, \quad (10)$$

where  $g$  denotes the sample,  $N$  the total number of samples,  $\mathbf{p}_g$  is the end-effector Cartesian position at sample  $g$  and  $S$  is the total travelled path computed as  $S = \sum_{g=2}^N \|\Delta \mathbf{p}_g\|$  which normalizes the path smoothness over the traveled distance. Curvature change is the second smoothness metric and it is complementary to path smoothness. Given the linear and angular Cartesian velocity,  $v(t)$  and  $\omega(t)$ , of the end-effector the curvature equation becomes:

$$\kappa(t) = \left| \frac{\omega(t)}{v(t)} \right|. \quad (11)$$

The curvature change metric as defined in [26] is given by:

$$F_4 : f_{cc} = \frac{1}{S} \int_0^{t_F} |\kappa'(t)| dt, \quad (12)$$

where the measure is divided by the path length to normalize for different movements.

The third smoothness metric measures velocity smoothness and captures oscillations at the joint level. This metric is relevant when the controller can influence multiple actuated joints. The metric,  $f_{vs}$ , is calculated as the zero crossing rate of the acceleration for all joints composing the robot:

$$a_g = \frac{v_{g+1} - v_g}{t_{g+1} - t_g}, \quad (13)$$

$$F_5 : f_{vs} = \frac{1}{N-1} \sum_{g=1}^{N-1} |\text{sign}(a_g) - \text{sign}(a_{g-1})|, \quad (14)$$

where  $v_g$  denotes the joint linear velocity at sample  $g$ . The joint subscript is omitted.

*Efficiency* relates to the computation performance of the MPC controller and measures how fast, on average, the

controller can compute a sample. The average computation time  $t_C$  is defined as:

$$F_6 : t_C = \frac{1}{N} \sum_{g=1}^N t_{g+1} - t_g, \quad (15)$$

where  $t$  is the time required to compute the sample,  $g$  is the sample index and  $N$  is the total number of samples. To summarize, Eq. (8), (9), (10), (12), (14), (15) are the six performance metrics used in this work.

## B. User input representation

The evaluation of the performance metrics requires gathering data about the performance of the robot during a movement under shared control. As we rely on simulation for data collection, our challenge lies in accurately representing user inputs to encompass a wide spectrum of scenarios, including both typical and corner cases. For every movement, we propose to represent user inputs as a concatenation of  $n$  velocity vectors which are applied each for a time  $T/n$  where  $T$  is the total time allowed to complete a movement. The velocity vector can have different magnitudes and directions which is what is expected from typical inputs of joystick control. The concatenation of velocity vectors with different magnitudes and directions allows the simulation of different acceleration and deceleration profiles. The initial robot joints' configuration is assumed to be randomized for every movement which means that for some movements the application of the velocity input might result in a request to move towards obstacles. This is beneficial during optimization since it replicates edge cases that exemplify the trade-off between following user commands and avoiding obstacles. For the validation use-case, we concatenated two velocity vectors per movement and considered a maximum execution time  $T$  of 20s. Three magnitudes were defined:  $\frac{k}{3} v_{max}$  for  $k = 1, 2, 3$ , where  $v_{max}$  represents the maximum velocity magnitude the user can request via the joystick. A simulation environment was created to emulate rooms where IGT robots usually operate. The environment comprises a room with two walls and a patient table, similar in size and shape to real-world environments, and was implemented in Simscape Multibody. The lateral view of the simulated room is visible in Fig 3. In the same figure, an image of an actual room is also reported. A total of  $n_{mov} = 40$  different concatenation options were generated. A top-down view of the room with superimposed all generated movements is displayed in Fig. 4.




Fig. 3: The environment created for the simulations (left). An image of a real operating room (right).




Fig. 4: All concatenated velocity vectors for simulation. The first vector is in red, the second is in orange, and obstacles are in black.

#### IV. EXPERIMENTAL VALIDATION

We applied the BO framework of Alg. 1 to the automatic tuning of the MPC-based shared controller in Eq. 1 for the IGT robot in Fig.1. The performance of the optimized controller is compared to the performance of a baseline, hand-tuned, controller. Furthermore, we performed a VR-based user study to validate the generalization of the optimized controller to user input not provided during the automatic tuning procedure.

A weighted sum of the metrics is a crucial term in the definition of the objective function  $J_{\xi}$  in Eq. 7. The choice of the weights reflects the relative prioritization of the metrics and it is a task-dependent choice. For the validation use-case, their prioritization is reported in Table I. As mentioned in Section II, seven parameters were selected for the optimization procedure, i.e.,  $\xi = [N_p, N_c, Q_x, Q_y, Q_{\theta}, c_1, c_2]$ . Their values after the application of the BO procedure are reported in Tab. II. For comparison purposes, the values of those parameters for the baseline, hand-tuned controller, are reported as well. The hand-tuned controller parameters were obtained by a trial and error approach during the course of several weeks. A full optimization procedure could be executed on a computer with an Intel Core i7 10th generation processor running at 2.7 GHz in less than 120 minutes. The performance metrics, defined in Section III-A, were evaluated for a set of 50 movements executed with the optimized MPC parameter set as well as with the hand-tuned set. The results are reported in Table III. For every metric, the collected performance values were subject to a one-tailed t-test to assess if the observed difference between the optimized and the nominal parameter set is statistically significant. The optimized set returned an overall objective function improvement of 14% with respect to the hand-

TABLE I: Overview of the metrics and their weights.

Type	Metric	$F_{h,j}$	$w_h$
Safety	Obstacle proximity	$d_{ob}$ [m]	0.15
	Time spent near obstacles	$t_{ob}$ [%]	0.30
Smoothness	Path smoothness	$f_{ps}$ [m]	0.15
	Curvature change	$f_{cc}$ [rad/m]	0.25
	Velocity smoothness	$f_{vs}$ [m/s <sup>2</sup> ]	0.10
Efficiency	Computation time	$t_C$ [ms]	0.05

TABLE II: The baseline and optimized MPC parameters.

	$N_p$	$N_c$	$Q_x$	$Q_y$	$Q_{\theta}$	$c_1$	$c_2$
base	25	13	1	1	1	5	20
opt.	21	5	3.047	0.62117	5.9981	7.3219	20.138

TABLE III: Metric values for 50 simulated movements with the confidence interval in brackets. The \* symbol indicates a statistically significant lower mean according to the Shapiro-Wilkinson test, with  $p < 0.05$ .

	$d_{ob}$ [m]	$t_{ob}$ [%]	$f_{ps}$ [m]	$f_{cc}^*$ [rad/m]	$f_{vs}$ [m/s <sup>2</sup> ]	$t_C$ [ms]
base	0.68 (0.05)	0 (0)	$1.2 \times 10^{-5}$ ( $1.6 \times 10^{-5}$ )	38.38 (8.91)	0.29 (0.05)	43.60 (1.82)
opt.	0.63 (0.06)	0 (0)	$1.1 \times 10^{-5}$ ( $1.4 \times 10^{-5}$ )	19.96 (4.63)	0.25 (0.07)	50.63 (2.26)

tuned controller (from 0.42 to 0.36) without introducing any infeasible solution. Overall, the optimized set significantly improved the curvature change metric. No other metric was found to be statistically different. This result is in line with the weights associated to the curvature change metric (see Table II) which is the second largest after time spent near obstacles.

#### A. Comparison in Virtual Reality

To validate the generalization of the optimized controller for user inputs that were not provided during the optimization procedure, we tested the controller performance in a VR setup comparing the optimized controller with respect to the hand-tuned one. A VR simulation of the IGT robot was created which was connected to a joystick-like control module. The control module was used to gather user input in the form of desired velocity of the robot end-effector. Commands were communicated to the MPC controller implemented in Matlab. The VR simulation was realized in Unity. A two way Matlab/Unity connection was established via a ROS bridge. The environment used for the simulations during optimization (see Fig. 3) was replicated exactly in Unity. During the experiment, users were requested to move the end-effector of the robot from an initial position on one side of the patient table to the opposite side, mimicking typical usage. A visualization of the VR experimental setup is displayed in Fig. 5. The experiment had a within-subject design. Two variables were manipulated during the experiment, the pair initial and final position of the end-effector (goal) and the controller configuration (hand-tuned or optimized). This led to a 3 (goals) x 2 (MPC configurations) experimental design which resulted in 6 trials per participant. A total of



Fig. 5: The setup of the VR user experiment.

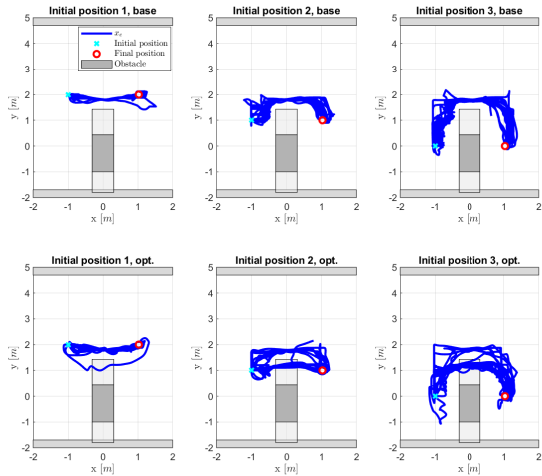


Fig. 6: The end-effector path for all users with the initial position (light blue cross) and the desired final position (red circle). Top hand-tuned controller. Bottom optimized controller.

25 participants, with a mean age of 38 years, took part in the user study. However, data from 3 participants were excluded due to setup errors. Approximately 73% of the participants had prior experience with the IGT robot used in the study, while 45% had prior experience with VR. The 6 trials were presented in random order.

The resulting paths of the end-effector for all users for every experimental condition are displayed in Fig. 6. For every condition that configured the controller with the hand-tuned set, none of the users was able to move the end-effector on the other side of the table by letting the C-shaped link traverse the tabletop (Fig 6). It was achievable, though, with the optimized set (Fig 6). Verbal user feedback provided during the experiment confirmed that the hand-tuned shared controller was too conservative w.r.t. repulsive actions in the vicinity of obstacles which was not the case for the optimized set. By observing the paths obtained from the optimized set (Fig 6) one can see that not every user achieved the goal position by traversing the tabletop. This is likely due to the design of the experiment itself, users that experienced the hand-tuned condition first were biased thinking that the robot could not traverse the tabletop and they did not try to do so in the follow-up trials. This was also confirmed by verbal feedback from users during the experiment.

The values of the metrics computed on the data set from the VR experiments are reported in Table IV. Surprisingly, a one-tail t-test shows a statistically significant improvement in all metrics apart from time spent close to obstacles, as the optimized parameters are less conservative in penalizing proximity to obstacles. We can argue that the difference in performance between simulation and VR is likely due to differences in the type of user input. Nevertheless, the optimized controller appears to have generalized to user-requested movements that were not provided during the optimization procedure. A problem that emerged during VR validation is that the optimized controller encountered an

TABLE IV: Metric values for the VR user experiments with the confidence interval in brackets. The \* symbol indicates a statistically significant lower mean according to the Shapiro-Wilkinson test, with  $p < 0.05$ .

	$d_{ob}^*$ [m]	$t_{ob}$ [%]	$f_{ps}^*$ [m]	$f_{cc}^*$ [rad/m]	$f_{vs}^*$ [m/s <sup>2</sup> ]	$t_C^*$ [ms]
base	0.55 (0.01)	0 (0)	$2.4 \times 10^{-5}$ ( $3 \times 10^{-6}$ )	142.30 (88.5)	0.13 (0.02)	112.90 (0.64)
opt.	0.52 (0.02)	0.08 (0.16)	$2.0 \times 10^{-5}$ ( $3.1 \times 10^{-6}$ )	49.11 (30.05)	0.22 (0.02)	105.30 (1.46)

infeasible solution on 9% of its iterations, while the baseline controller only encountered them on 0.2% of its iterations. Recall that no infeasible solutions were encountered for the simulated movements. During VR experiments, when infeasibility occurred, the robot kept executing the feasible effort computed in the previous sample and users autonomously adapted their input until a feasible solution was found again. To reduce the number of infeasible solutions, a possible improvement is the addition of real user inputs alongside simulated ones during optimization, at the likely cost of longer optimization time. It is also important to notice that in one trial the distance between the robot and obstacles was less than the radius of the collision spheres. This is an unwanted situation and further underlines the need to improve user input representation during optimization. The computation time  $t_C$  of the optimized controller was reduced with respect to the baseline controller, likely due to the shorter control horizon  $N_c$ . For the smoothness metrics, the larger improvement is seen for the curvature change, as expected since that metric has the largest weight in its category.

## V. CONCLUSION

This paper presented a framework for the tailoring of Bayesian optimization to the automatic tuning of MPC-based shared controllers, including a proposal for performance metrics and representation of user inputs. The performance of the tuned controller has been validated in simulation and with a VR-based user study. Results showed a significant performance improvement of the optimized controller w.r.t. a hand-tuned version. While tuning in simulation allowed to rapidly reach an optimized controller configuration, the user study outlined further challenges for the generalization of the results. Future work will look at enhancing the representation of the user input during the optimization process, possibly complementing it with data from experiments with users. While the proposed methods effectively optimized the performance metrics, it can be argued that other factors might be important in shared control such as user-perceived comfort or ergonomics. Extending the work considering an enlarged set of metrics constitutes another interesting research direction.

## ACKNOWLEDGEMENT

We kindly acknowledge the effort to enable the VR setup by the Philips IGT department and by Carolina Vissers in the context of the AMBER project.

## REFERENCES

- [1] A. Bauer, D. Wollherr, and M. Buss, "Human–robot collaboration: a survey," *International Journal of Humanoid Robotics*, vol. 5, no. 01, pp. 47–66, 2008.
- [2] D. A. Abbink, T. Carlson, M. Mulder, J. C. De Winter, F. Amiravan, T. L. Gibo, and E. R. Boer, "A topology of shared control systems—finding common ground in diversity," *IEEE Transactions on Human-Machine Systems*, vol. 48, no. 5, pp. 509–525, 2018.
- [3] P. I. Frazier, "Bayesian optimization," in *Recent advances in optimization and modeling of contemporary problems*. Informa, 2018, pp. 255–278.
- [4] F. Berkenkamp, A. Krause, and A. P. Schoellig, "Bayesian optimization with safety constraints: safe and automatic parameter tuning in robotics," *Machine Learning*, pp. 1–35, 2021.
- [5] A. Marco, F. Berkenkamp, P. Hennig, A. P. Schoellig, A. Krause, S. Schaal, and S. Trimpe, "Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with bayesian optimization," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 1557–1563.
- [6] A. Marco, P. Hennig, J. Bohg, S. Schaal, and S. Trimpe, "Automatic LQR tuning based on gaussian process global optimization," in *2016 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2016, pp. 270–277.
- [7] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.
- [8] N. Jaquier, L. Rozo, S. Calinon, and M. Bürger, "Bayesian optimization meets riemannian manifolds in robot learning," in *Conference on Robot Learning*. PMLR, 2020, pp. 233–246.
- [9] M. Turchetta, A. Krause, and S. Trimpe, "Robust model-free reinforcement learning with multi-objective bayesian optimization," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 10 702–10 708.
- [10] W. Edwards, G. Tang, G. Mamakoukas, T. Murphey, and K. Hauser, "Automatic tuning for data-driven model predictive control," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 7379–7385.
- [11] L. P. Fröhlich, C. Küttel, E. Arcari, L. Hewing, M. N. Zeilinger, and A. Carron, "Contextual tuning of model predictive control for autonomous racing," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 10 555–10 562.
- [12] M. H. Yeganegi, M. Khadiv, A. Del Prete, S. A. A. Moosavian, and L. Righetti, "Robust walking based on MPC with viability guarantees," *IEEE Transactions on Robotics*, vol. 38, no. 4, pp. 2389–2404, 2021.
- [13] J. G. Storms and D. M. Tilbury, "Blending of human and obstacle avoidance control for a high speed mobile robot," in *2014 American Control Conference*. IEEE, 2014, pp. 3488–3493.
- [14] M. Zarei, N. Kashi, A. Kalhor, and M. Tale Masouleh, "Experimental study on shared-control of a mobile robot via a haptic device with an optimal velocity obstacle based receding horizon control approach," *Journal of Intelligent & Robotic Systems*, vol. 97, no. 2, pp. 357–372, 2020.
- [15] R. Chipalkatty and M. Egerstedt, "Human-in-the-loop: Terminal constraint receding horizon control with human inputs," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, 2010, pp. 2712–2717.
- [16] J. Pankert and M. Hutter, "Perceptive model predictive control for continuous mobile manipulation," *IEEE Robot. Autom. Lett.*, vol. 5, no. 4, pp. 6177–6184, 2020.
- [17] C. R. Maurer, R. Qi, and V. Raghavan, "A linear time algorithm for computing exact euclidean distance transforms of binary images in arbitrary dimensions," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 25, no. 2, pp. 265–270, 2003.
- [18] I. D. Landau and G. Zito, *Digital control systems: design, identification and implementation*. Springer, 2006, vol. 130.
- [19] B. Letham, R. Calandra, A. Rai, and E. Bakshy, "Re-examining linear embeddings for high-dimensional bayesian optimization," *Advances in neural information processing systems*, vol. 33, pp. 1546–1558, 2020.
- [20] M. Binois and N. Wycoff, "A survey on high-dimensional gaussian process modeling with application to bayesian optimization," *ACM Transactions on Evolutionary Learning and Optimization*, vol. 2, no. 2, pp. 1–26, 2022.
- [21] M. Mehndiratta, E. Camci, and E. Kayacan, "Automated tuning of nonlinear model predictive controller by reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 3016–3021.
- [22] A. Gharib, D. Stenger, R. Ritschel, and R. Voßwinkel, "Multi-objective optimization of a path-following MPC for vehicle guidance: A bayesian optimization approach," in *2021 European Control Conference (ECC)*. IEEE, 2021, pp. 2197–2204.
- [23] A. Kapnopoulos and A. Alexandridis, "A cooperative particle swarm optimization approach for tuning an MPC-based quadrotor trajectory tracking scheme," *Aerospace Science and Technology*, vol. 127, p. 107725, 2022.
- [24] B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. De Freitas, "Taking the human out of the loop: A review of bayesian optimization," *Proceedings of the IEEE*, vol. 104, no. 1, pp. 148–175, 2015.
- [25] J. Wen, X. Zhang, Q. Bi, Z. Pan, Y. Feng, J. Yuan, and Y. Fang, "Mrpb 1.0: A unified benchmark for the evaluation of mobile robot local planning approaches," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 8238–8244.
- [26] D. Calisi and D. Nardi, "Performance evaluation of pure-motion tasks for mobile robots with respect to world models," *Autonomous Robots*, vol. 27, no. 4, pp. 465–481, 2009.