

Minimal Intervention Shared Control with Guaranteed Safety under Non-Convex Constraints

Shivam Chaubey, Francesco Verdoja, Shankar Deka, Ville Kyrki

Abstract—Shared control combines human intention with autonomous decision-making. At the low level, the primary goal is to maintain safety regardless of the user’s input to the system. However, existing shared control methods—based on, e.g., Model Predictive Control, Control Barrier Functions, or learning-based control—often face challenges with feasibility, scalability, and mixed constraints.

To address these challenges, we propose a Constraint-Aware Assistive Controller that computes control actions online while ensuring recursive feasibility, strict constraint satisfaction, and minimal deviation from the user’s intent. It also accommodates a structured class of non-convex constraints common in real-world settings. We leverage Robust Controlled Invariant Sets for recursive feasibility and a Mixed-Integer Quadratic Programming formulation to handle non-convex constraints. We validate the approach through a large-scale user study with 66 participants—one of the most extensive in shared control research—using a simulated environment to assess task load, trust, and perceived control, in addition to performance. The results show consistent improvements across all these aspects without compromising safety and user intent. Additionally, a real-world experiment on a robotic manipulator demonstrates the framework’s applicability under bounded disturbances, ensuring safety and collision-free operation.

I. INTRODUCTION

Shared control is a framework where control authority is either continuously or intermittently shared between a human operator and an assistive system [1], [2], and is commonly used in teleoperation, semi-autonomous navigation, assistive robotics, and medical systems. It combines user intent with the enforcement of safety constraints, task-specific requirements, or optimal behavior. In practice, users often have implicit goals but limited knowledge of internal dynamics or system constraints (e.g., joint limits, actuator capabilities, or environmental factors), which may lead to unsafe situations. Thus, assistive systems are needed to ensure safety and constraint satisfaction without compromising user control intent [3]. Although user intent prediction and task-driven shared control can enhance autonomy, they risk overriding user control, diminish trust, and depend on accurate demonstrations for model learning [1], [4]. Preserving intent is thus crucial for user engagement and trust, especially in human-centered applications [5].

Equally important is strict constraint satisfaction, which arises from physical limitations (e.g., actuator bounds or robot configurations), environmental considerations (e.g., obstacle avoidance), and task-specific demands [6] (e.g., precise

The authors acknowledge the use of the MIDAS infrastructure of Aalto School of Electrical Engineering. S. Chaubey, F. Verdoja, S. Deka, and V. Kyrki are with the School of Electrical Engineering, Aalto University, Espoo, Finland. {firstname.lastname}@aalto.fi

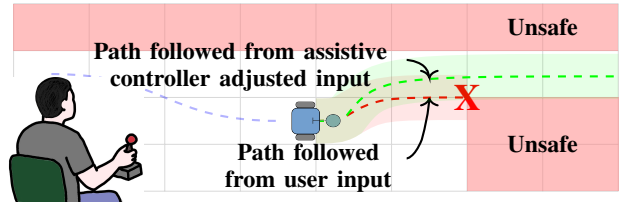


Fig. 1: Shared control framework: minimal adjustment of user input to avoid unsafe regions (red).

manipulation in activities like pouring [7]). When these constraints are not strictly met, safety, feasibility, task success, and user trust can be compromised. Methods such as Model Predictive Control (MPC) and Control Barrier Functions (CBFs) can enforce constraints, but in practice often require tuning of horizon length or slack penalties [5], [8]–[11] and barrier-function design (hand-crafted or learned) [12]–[14], and may be difficult to implement reliably under strict non-convex state constraints and bounded control inputs [15], [16]. Learning-based assistance can be context-aware [17], [18], but typically lacks hard guarantees beyond training data. These limitations motivate a unified framework that preserves user intent while providing strict constraint satisfaction and recursive feasibility without extensive tuning in such settings.

To address these limitations, we propose a Constraint-Aware Assistive Controller (CAAC) framework that ensures recursive feasibility and safety by leveraging Robust Controlled Invariant Sets (RCISs), enabling online, constraint-aware assistance via a single-step optimization. It also addresses a structured class of non-convex safe sets—specifically, those representable as the complement of a finite union of convex polytopic unsafe sets.

Our main contributions are as follows. First, we develop our framework by integrating offline-computed RCISs into a single-step Constrained Optimal Control Problem (COCP), reformulated as a Mixed-Integer Quadratic Programming (MIQP) to handle non-convex constraints. Second, we validate its effectiveness against our hypotheses through a large-scale user study with 66 participants—the first of its kind at this scale in shared control research. The study shows reduced workload, increased trust, preserved control authority, satisfied safety constraints, and improved performance. Finally, we validate the method on a robotic manipulator, where the controller maintains feasibility under bounded disturbances and ensures collision-free operation¹.

¹Supplementary material and code are available at: <https://version.aalto.fi/gitlab/irobotics/CAAC>

II. RELATED WORK

A wide range of shared control strategies have been proposed to assist users in task completion. These approaches differ in how they interpret user intent, ensure safety, and preserve the user’s sense of control.

Probabilistic methods, such as hindsight optimization [4] and policy blending [1], anticipate user intention and assist with task completion. Hindsight optimization learns the user’s goal through Partially Observable Markov Decision Processes (POMDPs). Policy blending, on the other hand, infers user intent and adjusts the controls accordingly. While both methods offer more proactive autonomy, they lack strict safety guarantees. Moreover, users often report a loss of control when the system enforces a specific inferred strategy, limiting their ability to adapt. This highlights a trade-off between autonomous assistance and user adaptability.

In the context of vehicles, shared control literature frequently focuses on predictive methods such as MPC to enforce constraints and optimize behavior. Without a suitable terminal/invariant set, ensuring recursive feasibility can require careful horizon selection and constraint handling, and obstacle anticipation often motivates longer horizons that increase computational cost [19]. To reduce this cost, some works adopt variable-step horizons [9], [20], and model obstacle avoidance via soft constraints [5], [9] or cost penalties [8], [10]. Other works emphasize user authority or reduce early intervention [5], [20]. Overall, challenges remain in achieving hard safety guarantees in practice with limited tuning and real-time computation, and in adequately evaluating user experience.

CBF-based methods provide theoretical safety via forward invariance safe sets, and are commonly implemented through real-time Quadratic Program (QP)/MPC safety filters or potential-field style controllers. In practice, barrier functions are either hand-crafted [21] or learned from data [13], while feasibility is often enforced through slack variables [3], [22], [23]. However, many CBF-based frameworks do not directly accommodate strict non-convex state and control constraints [15], [16], and extensive user studies remain limited.

Learning-based methods can adapt robot impedance and provide haptic guidance based on task dynamics, as discussed in [24]. Human behavior modeling using inverse differential games is explored in [25]. Operator workload reduction by learning manipulation tasks from demonstrations is proposed in [26]. While these approaches are promising for specific tasks, they often lack safety guarantees and restrict the user to learned tasks, thereby limiting flexibility.

To address the need for strict safety guarantees, feasibility, and preservation of user control, we propose leveraging RCIS, which ensures that, from any state within the set, there exists an admissible control input that keeps the system safe. Recent works [27]–[29]—including implicit formulations, lifted-space methods, and efficient set approximations—have made RCIS practical for high-dimensional systems. While RCIS have been widely used in control theory, their appli-




Fig. 2: Illustration of (a) convex unsafe sets and (b) a non-convex unsafe region decomposed into convex components.

cation to shared control remains limited. Compared to approaches that often rely on learned or hand-designed barrier functions and feasibility tuning, RCIS provide a systematic way to certify safety and recursive feasibility for admissible control actions within the set.

Building on these insights, we first formulate the COCP problem for our CAAC framework, then embed RCISs to guarantee recursive feasibility, and finally handle non-convex constraints within a MIQP formulation.

III. PROBLEM FORMULATION

To formalize the shared control problem, we consider a Discrete-Time Linear System (DTLS) of the form:

$$x_{k+1} = Ax_k + Bu_k + Ew_k, \quad (1)$$

where $x_k \in \mathbb{R}^n$ denotes the state vector, $u_k \in \mathbb{R}^m$ the control input, and $w_k \in \mathbb{R}^d$ an unknown disturbance at time step k . The matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, and $E \in \mathbb{R}^{n \times d}$ define the system dynamics. The disturbance w_k is assumed to lie within a bounded set $\mathcal{W} = \{w \in \mathbb{R}^d \mid Gw \leq g_w\}$, where $G_w \in \mathbb{R}^{p \times d}$ and $g_w \in \mathbb{R}^p$ define the polyhedron.

a) *State Constraints*: We define the feasible region for states as a convex set \mathcal{P} and a non-convex set \mathcal{D} . The convex set is represented as a polytope:

$$\mathcal{P} = \{x \in \mathbb{R}^n \mid \mathbf{F}x \leq f\}, \quad (2)$$

where $\mathbf{F} \in \mathbb{R}^{m_F \times n}$ and $f \in \mathbb{R}^{m_F}$.

In this work, we define the non-convex safe set \mathcal{D} as the complement of an unsafe set: $\mathcal{D} = \mathbb{R}^n \setminus \mathcal{D}^c$. The unsafe set \mathcal{D}^c is modeled as a finite union of convex polytopes: $\mathcal{D}^c = \bigcup_{i=1}^{n_b} \mathcal{D}_i^c$, where each \mathcal{D}_i^c is either a convex region in itself (Fig. 2a) or a component obtained by decomposing a non-convex unsafe region into convex parts (Fig. 2b). The set $\mathcal{X} = \mathcal{P} \cap \mathcal{D}$ defines the state admissible region in which the state x must remain at each time step.

b) *Control constraints*: Control constraints ensure that the control input u lies within a feasible region represented as a convex polytope:

$$\mathcal{U} = \{u \in \mathbb{R}^m \mid \mathbf{G}u \leq g\}, \quad (3)$$

where $\mathbf{G} \in \mathbb{R}^{m_G \times m}$ and $g \in \mathbb{R}^{m_G}$ define the half-space representation (H-representation) of the polytope.

c) *Assistive Control Problem*: The objective of CAAC is to determine a control input u_k that minimally perturbs the user's input u_k^r when necessary to ensure that safety and task-specific constraints are satisfied for all future time steps. Although the optimization considers only the current input, constraints are enforced over an infinite horizon to ensure that the selected control does not lead the system into future states where constraint satisfaction becomes infeasible. This is particularly important in shared control, where user inputs may inadvertently drive the system toward the boundary of the infeasible region, requiring early intervention to avoid entering infeasible or unsafe states. Given $x_{k+h} \in \mathcal{P} \cap \mathcal{D}$, this problem can be formulated as:

$$J_u = \min_{u_k} \frac{1}{2} (u_k - u_k^r)^\top Q_u (u_k - u_k^r) \quad (4a)$$

$$\text{s.t. } \exists u_{k+h} \in \mathcal{U}, \forall h \geq 0, \text{ for which;} \quad (4b)$$

$$Ax_{k+h} + Bu_{k+h} + Ew_{k+h} \in \mathcal{P} \cap \mathcal{D}, \forall w_{k+h} \in \mathcal{W},$$

where Q_u is a positive definite weight matrix.

IV. METHODS

The problem (4) enforces infinite-horizon constraints, solving it directly is intractable. We instead constrain the system to a precomputed RCIS, guaranteeing recursive feasibility and reducing the problem to a single-step optimization solvable in real time. To address non-convex state constraints, the safe region—defined as the complement of a union of convex polytopic unsafe sets—is encoded using the Big-M method [30], which converts the disjunctive structure into an equivalent intersection of relaxed half-spaces. The resulting constraints are constructed in sparse form within the MIQP framework, ensuring compatibility with solver and efficient use of sparsity. Together with one-step optimization and the use of RCIS, this formulation enables real-time shared control. While MIQP problems have exponential worst-case complexity, the combination of one-step structure and sparsity greatly reduces solver overhead in practice.

Notation: Let $\Phi(A_1, A_2, \dots, A_n)$ denote a block-diagonal matrix with blocks $A_i \in \mathbb{R}^{a \times b}$. The vectors $\mathbf{1}_a$ and $\mathbf{0}_a$ are column vectors of ones and zeros of size $a \times 1$, and $\mathbf{0}_{a,b}$ is the $a \times b$ zero matrix. Time indices are omitted when clear from context: the state $\mathbf{x}_k = [x_k, x_{k+1}]$, control u_k , and binary variable $p_{i,j,k+1}$ are written simply as \mathbf{x} , \mathbf{u} , and $p_{i,j}$.

A. Robust Controlled Invariant Sets

Definition 1: Consider the discrete-time linear system (1), subject to state constraints $x_k \in \mathcal{X} \subseteq \mathbb{R}^n$ and control constraints $u_k \in \mathcal{U} \subseteq \mathbb{R}^m$, where $\mathcal{X} = \mathcal{P} \cap \mathcal{D}$. Let $\mathcal{S}_{x,u} = \mathcal{X} \times \mathcal{U}$ denote the set of admissible state-control pairs.

A set $\mathcal{S} \subseteq \mathcal{X}$ is called a RCIS if, for every $x_k \in \mathcal{S}$, there exists a control input $u_k \in \mathcal{U}$ such that the next state $x_{k+1} = Ax_k + Bu_k + Ew_k$ also belongs to \mathcal{S} for all $w_k \in \mathcal{W}$.

Equivalently, if the system starts at any $x_k \in \mathcal{S}$, there exists a control sequence $\{u_k\}_{k \geq 0} \in \mathcal{U}$ that keeps the system within $\mathcal{S} \subseteq \mathcal{X}$ under disturbance w_k , ensuring recursive feasibility and constraint satisfaction at all future time steps. In this work, \mathcal{S} denotes a polytopic representation of the RCIS for

the admissible set $\mathcal{S}_{x,u}$, computed offline using the explicit closed-form method proposed in [29].

B. COCP formulation

To ensure recursive feasibility while minimizing deviation from the user's control input, a safe COCP is formulated by explicitly incorporating the RCIS as a terminal constraint, $x_{k+1} \in \mathcal{S}$. Given $x_k \in \mathcal{S}$, the resulting optimization problem at k^{th} time-step is:

$$J_u = \min_u \frac{1}{2} (u_k - u_k^r)^\top Q_u (u_k - u_k^r) \quad (5)$$

$$\text{s.t. } x_{k+1} = Ax_k + Bu_k + Ew_k, u_k \in \mathcal{U},$$

$$x_{k+1} \in \mathcal{S} \subseteq \mathcal{X}, \forall w_k \in \mathcal{W}.$$

C. Quadratic Programming Formulation

Reformulating the problem in (5) as a standard MIQP allows us to encode non-convex constraints, while structuring it in sparse matrix form provides a solver-ready formulation that can exploit sparsity for online CAAC applications. The COCP problem (5) can thus be reformulated as:

$$J = \min_z \frac{1}{2} z^\top Q z + q^\top z \quad (6a)$$

$$\text{s.t. } \underline{z} \leq \mathcal{A} z \leq \bar{z}. \quad (6b)$$

Here, z is the decision vector including states, controls, and binary variables introduced to handle non-convex constraints. Q is the positive semi-definite Hessian matrix, and q is the corresponding linear cost vector. The matrix \mathcal{A} encodes all equality and inequality constraints, while \underline{z} and \bar{z} denote the element-wise lower and upper bounds, respectively.

1) *Equality State Constraints*: Since the optimization is performed over a single step, we define the stacked state and control vectors as $\mathbf{x} = [x_k^\top, x_{k+1}^\top]^\top$ and $\mathbf{u} = u_k$ to maintain a consistent vector notation. The combined nominal DTLS (1) and initial condition can be written compactly as:

$$[\mathcal{A}' \quad \mathcal{B}'] \begin{bmatrix} \mathbf{x} \\ \mathbf{u} \end{bmatrix} = -\mathbf{x}_0, \quad (7)$$

$\mathbf{x}_0 = [x_k^\top, \mathbf{0}_n^\top]^\top$ encodes the known initial state $x_k \in \mathcal{S}$ at time step k , and matrices $\mathcal{A}' \in \mathbb{R}^{2n \times 2n}$ and $\mathcal{B}' \in \mathbb{R}^{2n \times m}$ are defined as:

$$\mathcal{A}' = \begin{bmatrix} \mathbf{0}_{n,n} & \mathbf{0}_{n,n} \\ A & \mathbf{0}_{n,n} \end{bmatrix} - \mathbf{I}_{2n \times 2n}, \quad \mathcal{B}' = \begin{bmatrix} \mathbf{0}_{n,m} \\ B \end{bmatrix}.$$

2) *Convex State Inequality RCIS Constraints*: For the admissible set of state-input pairs (2)–(3) with $\mathcal{S}_{x,u} = \mathcal{P} \times \mathcal{U}$ and the DTLS (1), we compute a RCIS. The RCIS is represented by the polytope $\mathcal{S}_F = \{x \in \mathbb{R}^n \mid \mathbf{C}_F x \leq c_F\}$. To impose these constraints only on the successor state x_{k+1} , we use the modified representation:

$$\mathbf{C}_F \mathbf{x} \leq c_F, \quad (8)$$

where $\mathbf{C}_F = [\mathbf{0}_{m_{C_F} \times n} \quad \mathbf{C}_F]$, $c_F = c_F$, and m_{C_F} denotes the number of rows of \mathbf{C}_F .

3) *Non-Convex State Inequality RCIS Constraints*:

a) *Decomposition of Non-Convex Safe Sets:* The convex unsafe region \mathcal{D}_i^- is described as intersections of n_{c_i} half-spaces, with each half-space $P_D(i, j)$ defined by a linear inequality: $T_{i,j}^\top x > t_{i,j}$, with $T_{i,j}$ is the inward-pointing normal vector of the j^{th} half-space, and $t_{i,j}$ is the corresponding offset. Thus, the region \mathcal{D}_i^- is $\mathcal{D}_i^- = \bigcap_{j=1}^{n_{c_i}} \{x \mid T_{i,j}^\top x > t_{i,j}\}$. To define the safe region \mathcal{D}_i , we take the complement of the unsafe region \mathcal{D}_i^- , resulting in a union of half-spaces: $\mathcal{D}_i = \bigcup_{j=1}^{n_{c_i}} \{x \mid T_{i,j}^\top x \leq t_{i,j}\}$. Finally, the overall safe region \mathcal{D} is defined as the intersection of the individual regions \mathcal{D}_i :

$$\mathcal{D} = \left\{ x \in \mathbb{R}^n \mid \bigcap_{i=1}^{n_b} \left(\bigcup_{j=1}^{n_{c_i}} P_D(i, j) \right) \right\}, \quad (9)$$

$$P_D(i, j) = \left\{ x \in \mathbb{R}^n \mid T_{i,j}^\top x \leq t_{i,j} \right\}. \quad (9a)$$

b) *Conversion of Non-Convex to Convex sets:* To convert each non-convex set \mathcal{D}_i , which is the union of convex half-space $P_D(i, j)$ region, into intersection of half-spaces $P_D(i, j)$, we apply Big-M formulation using binary variables $p_{i,j} \in \{0, 1\}$. Each $p_{i,j}$ corresponds to a half-space $P_D(i, j)$ of the non-convex safe set \mathcal{D}_i . The binary variables determine whether the set $P_D(i, j)$ is active ($p_{i,j} = 0$) or relaxed ($p_{i,j} = 1$). A large constant $M \gg 0$ is used to relax the constraint when $p_{i,j} = 1$, effectively deactivating it. In this way, the union of convex regions is represented through the intersection of all constraints—active or relaxed—controlled by binary variables.

Thus, the corresponding safe-set can be defined as:

$$\left\{ x \in \mathbb{R}^n \mid \bigcap_{i=1}^{n_b} \left(\bigcap_{j=1}^{n_{c_i}} \left(T_{i,j}^\top x - t_{i,j} \leq M p_{i,j} \right) \right) \right\}. \quad (10)$$

To ensure that the system state x lies within at least one convex region $T_{i,j}^\top x - t_{i,j} \leq 0$ of each non-convex safe set \mathcal{D}_i , we enforce:

$$0 \leq \sum_{j=1}^{n_{c_i}} p_{i,j} \leq n_{c_i} - 1, \quad \forall i \in \{1, \dots, n_b\}. \quad (11)$$

The above formulation of (10) and (11), applied for all $i \in \{1, n_b\}$ and $j \in \{1, n_{c_i}\}$, ensures that the user is not restricted to a specific region but can freely switch between regions that best satisfy the control objectives.

c) *Imposing RCIS Constraints:* For each convex region $P_D(i, j)$ that defines a half-space of a non-convex safe set \mathcal{D}_i , we define the corresponding set of admissible state-control pairs as: $\mathcal{S}_{x,u}(i, j) = \{(x, u) \in \mathbb{R}^n \times \mathbb{R}^m \mid \mathbf{F}x \leq f, T_{i,j}^\top x \leq t_{i,j}, \mathbf{G}u \leq g\}$.

A RCIS, $\mathcal{S}(i, j)$ is computed for each admissible set $\mathcal{S}_{x,u}(i, j)$. Each $\mathcal{S}(i, j)$ contains all states x_k for which there exists a control input u_k such that the successor state $x_{k+1} = Ax_k + Bu_k + Ew_k$ remains within the set for all $w_k \in \mathcal{W}$, satisfying all constraints. These sets are represented as polytopes in the state space of the form $\mathcal{S}(i, j) = \{x \in \mathbb{R}^n \mid \mathbf{C}_{i,j}x \leq c_{i,j}\}$, where $\mathbf{C}_{i,j} \in \mathbb{R}^{m_{c_{i,j}} \times n}$ and $c_{i,j} \in \mathbb{R}^{m_{c_{i,j}}}$ define the polytope.

Similar to (10), the enforcement of the RCIS constraint at time step $k+1$ for any admissible set $\mathcal{S}_{x,u}(i, j)$ follows the same approach. The resulting safe set is defined as:

$$\left\{ x_{k+1} \in \mathbb{R}^n \mid \bigcap_{i=1}^{n_b} \left(\bigcap_{j=1}^{n_{c_i}} \left(\mathbf{C}_{i,j}x_{k+1} - c_{i,j} \leq M \mathbf{1}_{m_{c_{i,j}}} p_{i,j,k+1} \right) \right) \right\}. \quad (12)$$

Following the conversion of non-convex to union of convex constraint formulation, we use equation (12) to impose the RCIS inequality constraints corresponding to the j^{th} subregion (half-space) of the non-convex safe set \mathcal{D}_i at next time step $k+1$: $\mathbf{C}_{i,j}x_{k+1} - M \mathbf{1}_{m_{c_{i,j}}} p_{i,j,k+1} \leq c_{i,j}$. For the full stacked state vector \mathbf{x}_k , this constraint is written as:

$$\underbrace{[\mathbf{0}_{m_{c_{i,j}} \times n}, \mathbf{C}_{i,j}]}_{\mathbf{C}_{i,j}} \mathbf{x}_k - M \mathbf{1}_{m_{c_{i,j}}} p_{i,j,k+1} \leq c_{i,j}.$$

For brevity, we omit the time indices and write \mathbf{x} in place of \mathbf{x}_k and $p_{i,j}$ in place of $p_{i,j,k+1}$. To include all $\mathcal{S}_{i,j}$ for all $j \in \{1, \dots, n_{c_i}\}$ corresponding to the i^{th} safe set \mathcal{D}_i , we write $\mathbf{C}_i \mathbf{x} - M \mathbf{L}_p \mathbf{p}_i \leq \mathbf{c}_i$, where $\mathbf{C}_i = [\mathbf{C}_{i,1}^\top, \dots, \mathbf{C}_{i,n_{c_i}}^\top]^\top$, $\mathbf{L}_p \mathbf{p}_i = \mathbf{\Phi}(\mathbf{1}_{m_{c_{i,1}}}, \dots, \mathbf{1}_{m_{c_{i,n_{c_i}}}})$, binary vector $\mathbf{p}_i = [p_{i,1}, \dots, p_{i,n_{c_i}}]^\top$, and $\mathbf{c}_i = [c_{i,1}, \dots, c_{i,n_{c_i}}]^\top$.

Similarly, to impose the intersection over all sets \mathcal{D}_i for $i \in \{1, \dots, n_b\}$, we obtain:

$$\mathcal{C}_T \mathbf{x} - M \mathbf{L}_p \mathbf{p} \leq \mathbf{c}_T, \quad (13)$$

where $\mathcal{C}_T = [\mathbf{C}_1^\top, \dots, \mathbf{C}_{n_b}^\top]^\top$, $\mathbf{L}_p = \mathbf{\Phi}(\mathbf{L}_{p_1}, \dots, \mathbf{L}_{p_{n_b}})$ is a block-diagonal matrix of dimension $\left(\sum_{i=1}^{n_b} \sum_{j=1}^{n_{c_i}} m_{c_{i,j}} \right) \times n_p$, $\mathbf{p} = [\mathbf{p}_1^\top, \dots, \mathbf{p}_{n_b}^\top]^\top$, and $\mathbf{c}_T = [\mathbf{c}_1^\top, \dots, \mathbf{c}_{n_b}^\top]^\top$.

d) *Enforcing constraints:* From equation (11), at least one region $P_D(i, j)$ belong to \mathcal{D}_i must be active in order to ensure avoidance of \mathcal{D}_i^- unsafe convex set for time step $k+1$. This can be compactly written as:

$$\mathbf{p}_i \leq \mathbf{1}_{n_{c_i}}^\top \underbrace{[p_{i,1} \quad p_{i,2} \quad \dots \quad p_{i,n_{c_i}}]^\top}_{\mathbf{p}_i} \leq \bar{\mathbf{p}}_i,$$

where $\mathbf{p}_i = \mathbf{0}$, $\bar{\mathbf{p}}_i = (n_{c_i} - 1)$. To enforce this constraint for all safe regions ($\forall i \in \{1, \dots, n_b\}$), we define:

$$\mathbf{p} \leq \mathbf{E} \mathbf{p} \leq \bar{\mathbf{p}}, \quad (14)$$

where $\mathbf{E} = \mathbf{\Phi}(\mathbf{1}_{n_{c_1}}^\top, \dots, \mathbf{1}_{n_{c_{n_b}}}^\top)$, $\mathbf{p} = \mathbf{0}_{n_p}$, and $\bar{\mathbf{p}} = [(n_{c_1} - 1) \dots n_{c_{n_b}} - 1]^\top$.

4) *QP Constraint Matrices:* The decision variable is defined as $z = [\mathbf{x}^\top \quad \mathbf{u}^\top \quad \mathbf{p}^\top]^\top$. Combining the control input constraints (3), equality constraint (7), convex (8) and non-convex (13) RCIS state inequality constraints, and binary variable constraints (14). The complete constraint set in (6b) can be expressed as $\underline{z} \leq \mathcal{A}z \leq \bar{z}$, where

$$\underline{z} = \begin{bmatrix} -\mathbf{x}_0 \\ -\infty \\ -\infty \\ -\infty \\ -\infty \\ \mathbf{p} \end{bmatrix}, \quad \mathcal{A} = \begin{bmatrix} \mathcal{A}' & \mathcal{B}' & \mathbf{0} \\ \mathbf{C}_F & \mathbf{0} & \mathbf{0} \\ \mathcal{C}_T & \mathbf{0} & -M \mathbf{L}_p \\ \mathbf{0} & \mathbf{G} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{E} \end{bmatrix}, \quad \bar{z} = \begin{bmatrix} -\mathbf{x}_0 \\ c_F \\ \mathbf{c}_T \\ g \\ \bar{\mathbf{p}} \end{bmatrix}. \quad (15)$$

5) *QP Objective Matrices*: The total objective function is expressed as $J = J_x + J_u + J_p$, corresponding to state, control, and binary variable costs, respectively. Since we only track the user-provided control input, $J_x = 0$ and $J_p = 0$. The objective reduces to the control cost: $J_u = \frac{1}{2} \mathbf{u}^\top \mathcal{Q}_u \mathbf{u} + \mathbf{q}_u^\top \mathbf{u}$.

The full objective function (6a) is then:

$$J = \min \frac{1}{2} \mathbf{z}^\top \underbrace{\Phi(\mathcal{Q}_x, \mathcal{Q}_u, \mathcal{Q}_p)}_{\mathcal{Q}} \mathbf{z} + \underbrace{[\mathbf{q}_x^\top \ \mathbf{q}_u^\top \ \mathbf{q}_p^\top]}_{\mathbf{q}^\top} \mathbf{z}, \quad (16)$$

where $\mathcal{Q}_x = 0_{2n \times 2n}$, $\mathcal{Q}_u = \mathbf{I}_{m \times m}$, $\mathcal{Q}_p = 0_{n_p \times n_p}$, $\mathbf{q}_x = 0_{n-2}$, $\mathbf{q}_u^\top = -\mathbf{u}^\top \mathcal{Q}_u$, $\mathbf{q}_p = 0_{n_p}$, and $n_p = \sum_{i=1}^{n_b} n_{c_i}$ represents the total number of binary variables. In this formulation, \mathcal{Q} is positive semi-definite, as only \mathcal{Q}_u contributes a strictly convex term. The worst-case computational complexity for this MIQP is $\mathcal{O}(2^{n_p} \cdot (2n + m)^3)$.

V. USER STUDY

This section presents a user study designed to evaluate the effectiveness of the proposed Constraint-Aware Assistive Controller (CAAC). Our goal is to investigate how its use impacts task load, trust, perceived control, and task completion (performance). The study involves two modes—one with CAAC and one with User-Only Control (UOC), where the user task is to navigate a maze-like environment by guiding a robot toward goals while avoiding obstacles. Given concerns about user and machine safety arising from letting novice users interact without assistance on a real system, we decided to conduct the user study in a simulation setting. We first describe our hypotheses, simulation setup, and user study setup. We then explain the evaluation methodology and statistical testing approach used to assess our hypotheses.

A. Hypotheses

The following hypotheses were formulated to evaluate the effectiveness of CAAC compared to UOC:

- H1**: CAAC decreases task load from the user.
- H2**: CAAC increases user's trust of being able to complete the task safely and successfully.
- H3**: CAAC does not decrease the user's perception of being in control of the system.
- H4**: CAAC decreases task completion time.

B. Experimental Setting

To evaluate the hypotheses, we designed a simulation (shown in Fig. 3) in which the objective is to guide a robot, represented by a small blue circle (1), to sequentially reach green goal circles (2), with progress tracked by a goal-reached counter (3), while avoiding black obstacles (5). Navigation is restricted to safe gray regions (4).

1) *Robot Behavior*: The robot state is defined as $x_k = [X_k, Y_k, v_k^x, v_k^y]^\top$ and the control input as $u_k = [a_k^x, a_k^y]^\top$, where (X_k, Y_k) denote position, (v_k^x, v_k^y) velocity, and (a_k^x, a_k^y) acceleration in x and y directions. All quantities are expressed in consistent, unitless values without mapping to

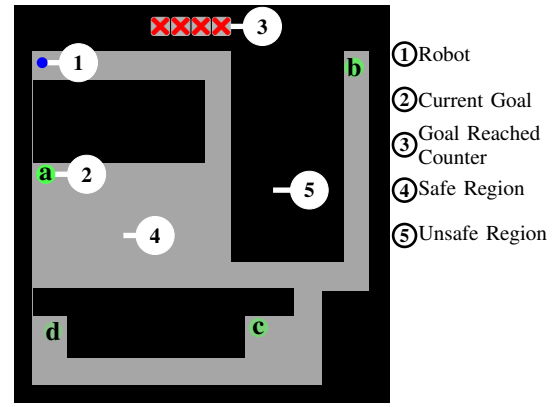


Fig. 3: Simulation layout with goal sequence (a \rightarrow d) shown in green for illustration; only one goal appears at a time.

a physical scale. The robot follows linear damped double-integrator dynamics $x_{k+1} = Ax_k + Bu_k$, where

$$A = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 - \gamma \Delta t & 0 \\ 0 & 0 & 0 & 1 - \gamma \Delta t \end{bmatrix}, \quad B = \begin{bmatrix} 0.5 \Delta t^2 & 0 \\ 0 & 0.5 \Delta t^2 \\ \Delta t & 0 \\ 0 & \Delta t \end{bmatrix}. \quad (17)$$

The damping coefficient is set to $\gamma = 0.1$, causing velocity to build gradually with input and decay without input. Velocity and acceleration are bounded along both axes. The user provides the reference acceleration u_k^r through the joystick's right analog stick, where direction sets the acceleration vector and deflection magnitude controls its strength.

The environment in Fig. 3 defines the robot's workspace bounded by outer walls and five axis-aligned rectangular obstacles (unsafe region). The constraints are expanded to account for the robot's shape, ensuring collision avoidance considers its geometry, not just its center point. The robot must stay within the environment bounds and avoid each obstacle by satisfying a union of axis-aligned half-space constraints.

Based on the system dynamics in (17) and the defined constraints, we constructed the admissible set $S_{x,u}(i,j)$ for each j^{th} half-space belonging to the i^{th} obstacle. We then computed a RCIS for every $S_{x,u}(i,j)$, constructed the constraint matrices (15) and objective (16), and solved the resulting MIQP using Gurobi solver [31] on a 12th Gen Intel Core i7-12700H with 32 GB RAM; the average solve time of 1.81 ms per step (solver-limited upper bound ≈ 552 Hz). The simulation ran at 30 Hz and the CAAC node at 50 Hz.

2) *Sequence*: Each mode begins with a fixed 120-second training session where the user can familiarize with the controls and freely move around. After training, participants begin a performance session where they are tasked to sequentially reach goals (green circles) in sequence as quickly as possible (shown in Fig. 3), with only one goal appearing at a time.

3) *Collisions and Goals*: A collision is triggered when the robot enters a black region, causing it to respawn at the initial or last goal position. A goal is reached only when the robot fully enters the green circle at under 0.9 units/s.

4) *Termination*: A session ends either when all goals are reached (performance session) or when the time limit is reached (training session). For each performance session, we recorded Completion Duration (**CD**) and collision count. After the session, the user completed the feedback survey.

C. User Study Setup

Participants were divided into two groups: Group A experienced their first sequence (training + performance session) with **UOC**, followed by a second sequence with the **CAAC**, while Group B followed the reverse order. Each sequence began with a dedicated training session before the corresponding performance session.

1) *Participants*: The study included 66 participants, compensated with a free lunch. All surveys were optional. Most were aged 25–34 (54.5%), followed by 18–24 (36.4%) and 35–44 (9.1%). The sample was predominantly male (71.2%), with female participants comprising 28.8%. Professionally, the majority were students (41.5%), followed by engineers (30.8%), researchers (16.9%), and others (10.8%). Regarding prior experience, 43.9% of participants reported rarely playing video games (less than once every two months), while only 10.6% played daily. Joystick use was even less common: 77.3% of participants used a joystick rarely, and 3% daily. The study was approved by our institution’s research ethics committee, and informed consent was obtained from all participants.

2) *Survey*: Participants completed a survey after each simulation session, which included RAW-TLX Questions (RTQs) listed in NASA Raw Task Load Index (Raw-TLX) [32], [33], Additional Questions (AQs), and two Open-ended Questions (OEQs) were asked to capture qualitative feedback. All items except the OEQs were rated on a five-point Likert scale. For RTQs ranging from “Failure” to “Perfect” for Successful (SC) and “Very Low” to “Very High” for the rest. For AQs, from “Strongly Disagree” to “Strongly Agree”. The exact questions asked are:

Mental Demand (MD): *How mentally demanding was the task?*
Physical Demand (PD): *How physically demanding was the task?*
Temporal Demand (TD): *How rushed was the task pace?*
Successful (SC): *How successful were you?*
Effort (EF): *How hard did you work?*
Frustration (FR): *How stressed or annoyed were you?*

AQs **Safe (SF)**: *I felt safe avoiding obstacles.*
Control (CT): *The agent accurately followed my actions.*

OEQs *The assistive controller affected my performance by [text].*
My experience with and without assistive controller was [text].

D. Evaluation and Testing Criteria

To evaluate our hypotheses, we used survey responses (RTQs, AQs) for **H1–H3** and the metric Completion Duration (**CD**) for **H4**, supplemented by open-ended questions (OEQs) for qualitative insights. We applied the Wilcoxon Signed-Rank (WSR) test for paired samples and the Mann-Whitney U (MWU) test for unpaired samples. Normality was assessed using the Shapiro-Wilk (SW) test.

Based on our predefined hypotheses, we conducted directional statistical tests to assess whether the median values

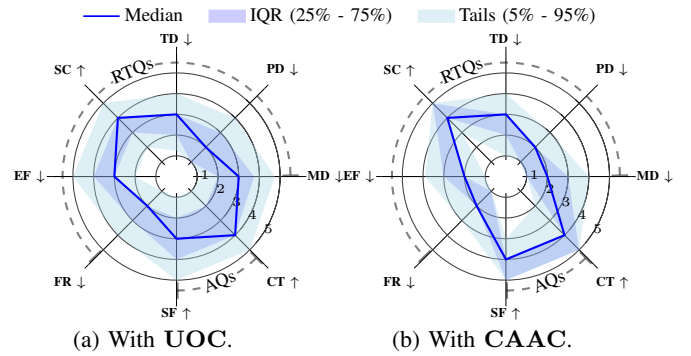


Fig. 4: RTQs and AQs responses: **UOC** vs. **CAAC**. Arrows show the desirable direction: \uparrow = higher is better, \downarrow = lower is better.

under the **CAAC** were significantly higher or lower than those under the **UOC**. Tab. I outlines the null hypotheses (H_0^{\geq} or H_0^{\leq}) associated with each question and the corresponding hypothesis (**H1–H4**) it supports. A two-sided test (H_0^{\neq}) was applied only in cases where assessing distributional symmetry was relevant. The significance level for each statistical test is categorized as follows: $\alpha \leq 0.05^*$ (significant), $\alpha \leq 0.01^{**}$ (highly significant), and $\alpha \leq 0.001^{***}$ (very highly significant).

VI. USER STUDY RESULTS

We evaluated the **CAAC**’s effectiveness through qualitative analysis for **H1–H3** and quantitative analysis for **H4** to assess overall impact.

a) *Qualitative Analysis*: To assess subjective experience and workload, we aggregated responses across all participants and analyzed the RTQs, AQs, and OEQs surveys, targeting hypotheses **H1–H3**. As shown in Fig. 4, **CAAC** consistently outperformed **UOC** across all questions supporting our hypothesis **H1–H3**.

To further validate these results, a paired one-tailed WSR test confirmed statistically significant improvements for **CAAC** across all metrics (Tab. I), with very highly significant difference ($\alpha < 0.001^{***}$), except for Temporal Demand (**TD**), which only showed significant difference ($\alpha < 0.05^*$). In addition, most of the OEQs responses from the users anecdotally reinforced that **CAAC** reduced the workload, made them feel much safer, and kept them feeling in control. One user reported:

“It is interesting with assistive system when the agent can avoid the obstacles by itself to make the player feel safe and secured despite how careless the player is. The agent was driven faster to reach the goal.”

b) *Quantitative Analysis*: To evaluate hypothesis **H4**, we used OEQs and recorded experimental data to calculate task Completion Duration (**CD**) metric. For a fair comparison, **CD** under the **UOC** was computed only from collision-free trajectories, excluding any collision or respawn delays. No collisions occurred with **CAAC**, supporting its safety guarantees.

This analysis revealed that **CAAC** significantly reduced **CD**, supporting hypothesis **H4**. As visualized in Fig. 5,

| | MD | PD | TD | SC | EF | FR | SF | CT | CD |
|------------------------------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|-----------|
| Null Hypothesis (H_0) | $H_0^>$ | $H_0^>$ | $H_0^>$ | $H_0^<$ | $H_0^>$ | $H_0^>$ | $H_0^<$ | $H_0^<$ | $H_0^>$ |
| Alternate Hypothesis (H) | H1 | H1 | H1 | H1 | H1 | H1 | H2 | H3 | H4 |
| Significance (α) | 0.001*** | 0.001*** | 0.05* | 0.001*** | 0.001*** | 0.001*** | 0.001*** | 0.001*** | 0.001*** |
| Sample Size (N) | 66 | 65 | 64 | 60 | 60 | 61 | 62 | 66 | 66 |

TABLE I: Wilcoxon Signed-Rank (WSR) test results comparing UOC vs. CAAC.





Fig. 5: Aggregated Survey. Fig. 6: Group-wise Survey.

participants completed the task faster under CAAC. To validate these findings, we used paired data across all participants. The two-sided Shapiro-Wilk (SW) test yielded very highly significant difference ($\alpha < 0.001^{***}$), rejecting the null hypothesis (H_0^-) of normality, prompting the use of non-parametric one-tailed WSR test, which confirmed a very highly significant reduction in CD when using CAAC ($\alpha < 0.001^{***}$). Feedback from the OEQs further supports this trend. Many users felt safer with CAAC, which reduced perceived risk and made them less conservative, resulting in lower CD. One user noted:

“With no assistive system, I tend to choose more conservative route-move slower, always ready to stop and take detour if it can avoid collision. With assistive system, I tend to choose more aggressive (and shorter) route.”

c) *Learning effect:* To investigate how prior experience in one condition influences performance in the other, we conducted additional group-wise analyses using unpaired data within each group. Qualitative group-wise results were generally consistent with the paired data findings but did not reach significance for TD, likely due to reduced statistical power in the smaller subgroups. In contrast, quantitative group-wise results statistically mirrored the paired analysis, confirming that CAAC significantly reduced task CD across both groups, as shown in Fig. 6. Additionally, the analysis revealed a learning effect under UOC: participants who first used the CAAC took longer time when later operating without it, suggesting a potential reliance. No such effect was observed under CAAC—prior experience did not affect CD—indicating that CAAC supports consistent behavior regardless of potential earlier experience without it.

VII. REAL-WORLD EXPERIMENT

Shared control robotic applications such as surgery, industrial assembly, and hazardous material handling often require strict constraint satisfaction under disturbances. To demonstrate relevance in such settings, we evaluate our method on a Franka Panda manipulator (Fig. 7a). Measurement noise, delay, and unmodeled dynamics are treated as bounded disturbances, with bounds found experimentally.

The user was tasked to navigate the end-effector, fitted with a 19mm peg, through a maze of three rectangular blocks with 32mm passages, leaving 6.5mm clearance per

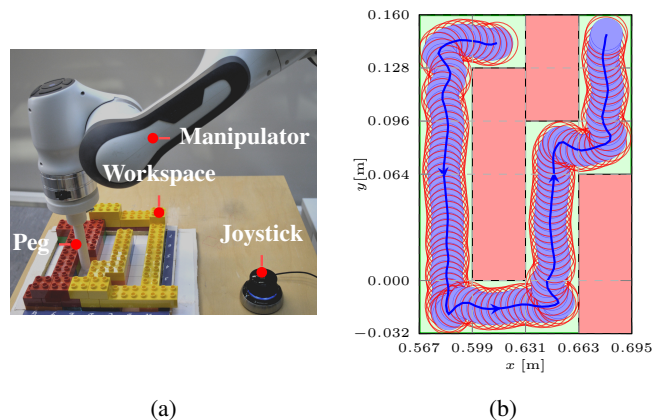


Fig. 7: Real experiment setup. (a) shows the maze workspace, and (b) the trajectory with disturbance bound (red circle).

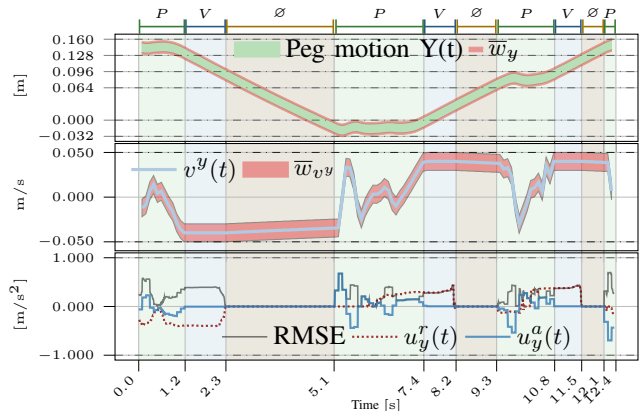


Fig. 8: Y-axis position, velocity, and control inputs (top-bottom), with RMSE of user vs. controller input.

side, reduced to about 3.5mm under ± 3 mm disturbance bounds. Five trials were conducted, and the trajectory from one trial is shown in Fig. 7b. Velocity is constrained to $[-0.05, 0.05]$ m/s (disturbance ± 0.01 m/s) and acceleration to $[-1, 1]$ m/s². A Cartesian controller drives the end-effector in the X-Y plane using assistive inputs u_k^a generated by the proposed CAAC framework, with states (X, Y, v^x, v^y) governed by (17). Acceleration-level control is used instead of position-level control, making the task more challenging. To analyze the controller’s behavior, we focus on the y-axis, where the manipulator covers a larger distance. As the x- and y-dynamics are decoupled, it is sufficient to examine this axis. Fig. 8 shows the evolution of position, velocity, and control inputs(top-bottom).

The controller’s behavior can be categorized into three regions. In region (P), the controller intervenes to avoid wall collisions in narrow passages. In region (V), the controller

enforces velocity limits whenever user inputs u_y^r would exceed the limits. In region (\emptyset), the controller followed the user inputs exactly (RMSE ≈ 0), as the end-effector was well aligned with the passage and required no correction to satisfy the limits.

The experiment confirms that the controller enforces safety and constraint satisfaction with minimal intervention under disturbance. Importantly, the system remains safe under arbitrary user inputs, as the controller intervenes only when necessary. This property is also relevant in teleoperation with communication delays, where deploying the controller locally ensures safety.

VIII. CONCLUSIONS

We proposed a constraint-aware CAAC framework for real-time shared control with formal safety guarantees under structured non-convex constraints. The user study confirmed the effectiveness of the CAAC. Qualitative results showed improved safety, control, and reduced workload, while quantitative analysis demonstrated faster task completion with zero collisions. Overall, the CAAC enabled more direct routes without loss of perceived control, reflecting greater trust and confidence.

The real-world manipulator experiment further confirmed the framework's practicality by showing that it can maintain safety under disturbances and operate reliably with structured non-convex constraints. This highlights its applicability to assistive tasks and teleoperation scenarios where safety assurance is paramount.

While the approach assumes known linear time-invariant dynamics, it establishes a principled foundation for constraint-aware shared control. Thus, extending the work to dynamic and uncertain settings, nonlinear models, and learning-based safety sets is an appealing avenue for the future, towards transparent, reliable, and minimally intrusive real-time assistance.

REFERENCES

- [1] A. D. Dragan and S. S. Srinivasa, "A policy-blending formalism for shared control," *The Int. J. of Robot. Res.*, vol. 32, no. 7, pp. 790–805, 2013.
- [2] D. P. Losey, C. G. McDonald, E. Battaglia, and M. K. O'Malley, "A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction," *Appl. Mechanics Reviews*, vol. 70, no. 1, p. 010804, 02 2018.
- [3] R. Periotta, M. Ferizbegovic, F. S. Barbosa, and R. C. Sundin, "Mpc-cbf with adaptive safety margins for safety-critical teleoperation over imperfect network connections," in *Proc. 2024 Eur. Control Conf. (ECC)*, 2024, pp. 1609–1615.
- [4] S. Javdani, S. Srinivasa, and J. A. Bagnell, "Shared autonomy via hindsight optimization," in *Proc. Robotics: Sci. and Syst. (RSS)*, 2015.
- [5] Y. Lu, L. Bi, and H. Li, "Model predictive-based shared control for brain-controlled driving," *IEEE Trans. on Intell. Transp. Syst.*, vol. 21, no. 2, pp. 630–640, 2020.
- [6] S. Chaubey, F. Verdoja, and V. Kyrki, "Jointly learning cost and constraints from demonstrations for safe trajectory generation," in *Proc. 2024 IEEE/RSJ Int. Conf. on Intell. Robots and Syst. (IROS)*, 2024, pp. 3635–3642.
- [7] K. Kronhardt, M. Pascher, and J. Gerken, "Understanding shared control for assistive robotic arms," 2023, arXiv:2303.01993.
- [8] C. Huang, F. Naghdy, H. Du, and H. Huang, "Shared control of highly automated vehicles using steer-by-wire systems," *IEEE/CAA J. of Automatica Sinica*, vol. 6, no. 2, pp. 410–423, 2019.

- [9] S. M. Ertien, S. Fujita, and J. C. Gerdes, "Shared steering control using safe envelopes for obstacle avoidance and vehicle stability," *IEEE Trans. on Intell. Transp. Syst.*, vol. 17, no. 2, pp. 441–451, 2016.
- [10] J. Talbot, M. Brown, and J. C. Gerdes, "Shared control up to the limits of vehicle handling," *IEEE Trans. on Intell. Vehicles*, vol. 9, no. 1, pp. 2977–2987, 2024.
- [11] A. Broad, T. Murphey, and B. Argall, "Highly parallelized data-driven mpc for minimal intervention shared control," 2019, arXiv:1906.02318.
- [12] H. Dai and F. Permenter, "Convex synthesis and verification of control-lyapunov and barrier functions with input constraints," in *Proc. 2023 Amer. Control Conf. (ACC)*, 2023, pp. 4116–4123.
- [13] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, "Learning control barrier functions from expert demonstrations," 2020, arXiv:2004.03315.
- [14] C. Dawson, Z. Qin, S. Gao, and C. Fan, "Safe nonlinear control using robust neural lyapunov-barrier functions," 2021, arXiv:2109.06697.
- [15] G. Notomista and M. Saveriano, "Safety of dynamical systems with multiple non-convex unsafe sets using control barrier functions," *IEEE Control Syst. Lett.*, vol. 6, pp. 1136–1141, 2022.
- [16] M. Srinivasan, A. Dabholkar, S. Coogan, and P. Vela, "Synthesis of control barrier functions using a supervised machine learning approach," 2020, arXiv:2003.04950.
- [17] D. P. Losey, H. J. Jeon, M. Li, K. Srinivasan, A. Mandlekar, A. Garg, J. Bohg, and D. Sadigh, "Learning latent actions to control assistive robots," 2021, arXiv:2107.02907.
- [18] M. Rubagotti, B. Sangiovanni, A. Nurbayeva, G. P. Incremona, A. Ferrara, and A. Shintemirov, "Shared control of robot manipulators with obstacle avoidance: A deep reinforcement learning approach," *IEEE Control Syst. Mag.*, vol. 43, no. 1, pp. 44–63, 2023.
- [19] Z. Chen, J. Lai, P. Li, O. I. Awad, and Y. Zhu, "Prediction horizon-varying model predictive control (mpc) for autonomous vehicle control," *Electronics*, vol. 13, no. 8, 2024.
- [20] W. Schwarting, J. Alonso-Mora, and D. Rus, "Parallel autonomy in automated vehicles: Safe motion generation with minimal intervention," *2017 IEEE Intell. Vehicles Symp. (IV)*, 2017.
- [21] A. Hebri, S. Acharya, M. Theofanidis, and F. Makedon, "A teleoperation framework for robots utilizing control barrier functions in virtual reality," in *Proc. 16th Int. Conf. on Pervasive Technologies Related to Assistive Environments*, ser. PETRA '23. New York, NY, USA: Association for Computing Machinery, 2023, p. 408–412.
- [22] J. Dallas, J. Talbot, M. Suminaka, M. Thompson, T. Lew, G. Orosz, and J. Subosits, "Control barrier functions for shared control and vehicle safety," 2025, arXiv:2503.19994.
- [23] X. Xu, "Control sharing barrier functions with application to constrained control," in *Proc. 2016 IEEE 55th Conf. on Decis. and Control (CDC)*, 2016, pp. 4880–4885.
- [24] Y. Michel, Z. Li, and D. Lee, "A learning-based shared control approach for contact tasks," *IEEE Robot. and Automat. Lett.*, vol. 8, no. 12, pp. 8002–8009, 2023.
- [25] H.-N. Wu and M. Wang, "Learning human behavior in shared control: Adaptive inverse differential game approach," *IEEE Trans. on Cybern.*, vol. 54, no. 6, pp. 3705–3715, 2024.
- [26] B. Xi, S. Wang, X. Ye, Y. Cai, T. Lu, and R. Wang, "A robotic shared control teleoperation method based on learning from demonstrations," *Int. J. of Adv. Robotic Syst.*, vol. 16, no. 4, 2019.
- [27] M. Fiacchini and M. Alami, "Computing control invariant sets is easy," 2017, arXiv:1708.04797.
- [28] T. Anevlavis and P. Tabuada, "A simple hierarchy for computing controlled invariant sets," in *Proc. 23rd Int. Conf. on Hybrid Systems: Computation and Control*, ser. HSCC '20, 2020.
- [29] T. Anevlavis, Z. Liu, N. Ozay, and P. Tabuada, "Controlled invariant sets: Implicit closed-form representations and applications," *IEEE Trans. on Autom. Control*, vol. 69, no. 7, pp. 4506–4521, 2024.
- [30] T. Schouwenaars, B. De Moor, E. Feron, and J. How, "Mixed integer programming for multi-vehicle path planning," in *Proc. 2001 Eur. control Conf. (ECC)*. IEEE, 2001, pp. 2603–2608.
- [31] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2024.
- [32] S. Hart, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," *Human mental workload/Elsevier*, 1988.
- [33] S. G. Hart, "Nasa-task load index (nasa-tlx); 20 years later," in *Proc. human factors and Ergonom. Soc. Annu. meeting*, vol. 50, no. 9. Sage publications Sage CA: Los Angeles, CA, 2006, pp. 904–908.