

Reactive Robot Navigation Using Quasi-Conformal Mappings and Control Barrier Functions

Gennaro Notomista¹, Member, IEEE, Gary P. T. Choi², and Matteo Saveriano³, Senior Member, IEEE

Abstract—This article presents a robot control algorithm suitable for safe reactive navigation tasks in cluttered environments. The proposed approach consists of transforming the robot workspace into the *ball world*, an artificial representation where all obstacle regions are closed balls. Starting from a polyhedral representation of obstacles in the environment, obtained using exteroceptive sensor readings, a computationally efficient mapping to ball-shaped obstacles is constructed using quasi-conformal (QC) mappings and Möbius transformations. The geometry of the ball world is amenable to provably safe navigation tasks achieved via control barrier functions (CBFs) employed to ensure collision-free robot motions with guarantees both on safety and on the absence of deadlocks. The performance of the proposed navigation algorithm is showcased and analyzed via extensive simulations and experiments performed using different types of robotic systems, including manipulators and mobile robots.

Index Terms—Constrained control, mobile robots, optimal control, robot control.

I. INTRODUCTION

SAFETY of dynamical systems (DSs) is receiving increasingly more attention, thanks to recently developed theoretical and computational tools that allow us to formulate several safety-critical controllers as convex optimization control policies. Applications include robot collision avoidance [1], energy-aware systems [2], and safe learning [3], [4]. Safety is intended as the forward invariance property of a subset of the state space of the DS describing the system. That is, a system is safe with respect to a set \mathcal{S} if the trajectory of the state, $x(t)$, satisfies $x(t_0) \in \mathcal{S} \implies x(t) \in \mathcal{S} \quad \forall t \geq t_0$. This notion of safety has been employed to formulate

Received 24 July 2024; revised 8 November 2024; accepted 9 December 2024. Date of publication 17 January 2025; date of current version 25 April 2025. The work of Gennaro Notomista was supported by the NSERC Discovery under Grant RGPIN-2023-03703. The work of Gary P. T. Choi was supported in part by the CUHK Vice-Chancellor Early Career Professorship Research Startup Fund under Project 4933624 and in part by the Croucher Foundation Start-Up Allowance. The work of Matteo Saveriano was supported by the European Union Project INVERSE GA under Grant 101136067. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Health and Digital Executive Agency (HADEA). Neither the European Union nor HADEA can be held responsible for them. Recommended by Associate Editor I. Prodan. (Corresponding author: Gennaro Notomista.)

Gennaro Notomista is with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada (e-mail: gennaro.notomista@uwaterloo.ca).

Gary P. T. Choi is with the Department of Mathematics, The Chinese University of Hong Kong, Hong Kong (e-mail: ptchoi@cuhk.edu.hk).

Matteo Saveriano is with the Department of Industrial Engineering, University of Trento, 38123 Trento, Italy (e-mail: matteo.saveriano@unitn.it).

Code is available online at <https://github.com/idra-lab/qcm-cbf>.

This article has supplementary material available at <https://youtu.be/phGiLKjysnQ>.

Digital Object Identifier 10.1109/TCST.2025.3526377

problems of safe motion planning for autonomous systems [5], navigation [6], [7], and autonomy [8], [9]. More recently, safe reinforcement learning problems have been formulated leveraging such a control-theoretic notion of safety [10], [11], [12], [13].

At the motion planning stage, safety constraints have been considered, which typically result in nonconvex problem formulations (e.g., [3], [14]). The recently developed control barrier function (CBF) [15] gained popularity also, thanks to the low computational complexity of convex optimization formulations, as well as their ability to encode a large variety of safety specifications. Nevertheless, the low computational complexity comes at the cost of a reactive controller synthesis approach, as opposed to planning-like strategies (see, e.g., [16] for a recent work trying to unify reactive and predictive techniques). When such reactive controllers are employed, it has been shown that the presence of competing objectives (more specifically, stability and safety) may generate undesired and asymptotically stable equilibrium points [17]. This phenomenon is particularly critical as undesirable asymptotically stable equilibria exist even in the presence of convex unsafe regions (e.g., obstacles).

The presence of undesirable equilibrium points arises in robot navigation and path planning problems, where robots have to navigate around obstacles to reach a desired goal. Over the past several decades, different approaches for robot navigation and path planning via mathematical transformations have been proposed [18], [19], [20], [21], [22] (see [23] for a survey). For instance, Sarkar et al. [24] developed a method for greedy routing in sensor networks using Ricci flow. Loizou [6] and Constantinou and Loizou [25] tackled the navigation problem by mapping a star world to a domain called the point world. Vlantis et al. [26] developed a method for robot navigation by transforming a real workspace into a punctured disk using harmonic maps. Methods for motion planning using Schwarz–Christoffel mappings were proposed in [27] and [28]. Fan et al. [29] proposed an iterative scheme for constructing conformal navigation transformations that map a complex workspace to a multiply connected circle domain.

In our recent work [30], we presented an approach to mitigate the problem of undesirable asymptotically stable equilibria in the case of multiple nonconvex unsafe regions. The control algorithm consists of mapping the system state space with possibly nonconvex unsafe regions, to a space where unsafe regions are either closed balls or the complement of open balls. A safety-preserving controller is computed in the ball world and mapped, via an appropriately defined

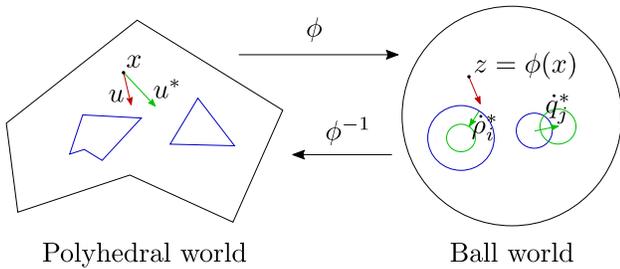


Fig. 1. Approach proposed in this article to ensure the safety of DSs in the presence of multiple nonconvex unsafe regions consists of mapping the state space—the *polyhedral world*—to a *ball world*, where obstacles are closed balls or the complement of open balls. In the ball world, obstacles are transformed by changing their centers and radii, by means of the inputs $\hat{\rho}_i^*$ and \hat{q}_j^* . This lets the mapped state z , and hence the real state x , to remain safe.

diffeomorphism, into the controller for the robot moving in the real world. The algorithm is formulated as a convex quadratic program (QP) regardless of the number and the convexity of unsafe regions. Fig. 1 pictorially illustrates this idea.

In this article, we design a computationally efficient reactive navigation policy for robotic systems which, thanks to the novel application of computational geometric tools to optimization-based robot control, significantly extends the range of applicability of the idea proposed in [30]. More concretely, the proposed control algorithm improves on [30] in the following ways.

- 1) We employ QC mappings to be able to represent obstacle regions in the real world—which will be referred to as the *polyhedral world*—of which we do not have an analytical expression.
- 2) We show how to compose Partial QC mappings in such a way as to get a diffeomorphism between the polyhedral world and the ball world, where controllers are synthesized.
- 3) We particularize the implementation of the developed algorithm for specific classes of systems modeling a large number of mobile robots employed in practical applications, namely systems near-identity diffeomorphic to single integrators, feedback linearizable systems, and differentially flat systems.
- 4) We report the results of extensive simulations and experiments with real robotic platforms to discuss the advantages and limitations of the proposed approach in terms of tradeoffs between robustness and computational complexity.

Compared to the work in [30], these contributions allow for the implementation of reactive controllers which, based on constrained-optimization-based formulations, guarantee a safe execution of robotic tasks, while preventing undesired equilibrium points from appearing and practically preventing the existence of deadlocks. Furthermore, we highlight how our proposed algorithm is amenable to be combined with existing tracking control techniques (based on, e.g., feedback linearization and differential flatness).

The remainder of the article is organized as follows. In Section II, we review the background and related work on robot navigation. In Section III, we describe our pro-

posed safety-preserving algorithm leveraging QC mappings and convex optimization problems. Numerical validation and comparisons are presented in Section IV. In Section V, we present both simulation and real experimental results for robotics applications. We conclude our work and discuss possible future directions in Section VI.

II. BACKGROUND AND RELATED WORK

A. Problem Definition

In this article, we consider robotic systems modeled by a control-affine DS

$$\dot{x} = f(x) + g(x)u \quad (1)$$

where $x \in \mathbb{R}^n$ is the robot state, $u \in \mathbb{R}^m$ is the robot control input, and $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ are locally Lipschitz continuous vector fields. The robots have to reach a desired goal in the environment, and this behavior is assumed to be achieved via a given control policy

$$\hat{u}: \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^m.$$

This way

$$\dot{x} = f(x) + g(x)\hat{u}(x, t) \quad (2)$$

where t denotes the time variable, makes the state x evolve to reach a desired goal state.

A computationally efficient safety-preserving control design for control affine DSs is based on the use of CBF [15], which results in a convex optimization control policy. The collision-free set—the *safe set*—is defined as the zero-superlevel set of a continuously differentiable function h , that is, the safe set S is defined as $S = \{x \in \mathbb{R}^n: h(x) \geq 0\}$. Then, if it exists, the following controller guarantees the collision-free motion of the robot:

$$\begin{aligned} u^*(x, t) &= \arg \min_u \|u - \hat{u}(x, t)\|^2 \\ \text{s.t. } L_f h(x) + L_g h(x)u &\geq -\alpha(h(x)) \end{aligned} \quad (3)$$

where $L_f h(x) = (\partial h / \partial x) f(x)$ and $L_g h(x) = (\partial h / \partial x) g(x)$ denote the Lie derivatives of h in the direction of the vector fields f and g , respectively, and $\alpha: \mathbb{R} \rightarrow \mathbb{R}$ is a class \mathcal{K} function, that is, a continuous, monotonically increasing function, with $\alpha(0) = 0$.

For robot navigation tasks, the safe set S is straightforwardly computable when the environment and the obstacles are balls. Therefore, the approach we take in this article to devise a reactive robot navigation policy consists of the following steps.

- 1) Reconstructing the polyhedral world from readings of exteroceptive sensors mounted on the robot.
- 2) Mapping the polyhedral world to a ball world, where the workspace and the obstacles are closed balls.
- 3) Solving for a robot control input that guarantees the collision-free motion of the robot in the ball world.
- 4) Mapping the robot control input from the ball world to the polyhedral world and sending it to the robot.

This is achieved by changing the position and size of the obstacles in the ball world, which, in turn, changes the diffeomorphism between the free space in the polyhedral world and the free space in the ball world. We then leverage this diffeomorphism and its Jacobian to compute the input for the robot

in the polyhedral (real) world that ensures its collision-free motion without introducing undesired equilibrium points.

To this end, let z denote the robot state in the ball world and $\phi: \mathbb{R}^n \rightarrow \mathbb{R}^n$ the mapping from the polyhedral world to the ball world, namely $z = \phi(x) \in \mathbb{R}^n$ (see Fig. 1). Assuming ϕ is continuously differentiable and invertible, we can write the robot dynamics in the ball world as follows:

$$\begin{aligned} \dot{z} &= \frac{\partial \phi}{\partial x} \dot{x} = \frac{\partial \phi}{\partial x} f(x) + \frac{\partial \phi}{\partial x} g(x)u \\ &= \underbrace{\frac{\partial \phi}{\partial x} (f \circ \phi^{-1})(z)}_{=: f_z} + \underbrace{\frac{\partial \phi}{\partial x} (g \circ \phi^{-1})(z)}_{=: g_z} u. \end{aligned} \quad (4)$$

As the algorithm is implemented on real systems at discrete time stamps, we can relax the differentiability assumption and approximate derivatives with finite differences, by always keeping the theoretical guarantees of the CBF. Having defined the CBF h whose zero-superlevel set is the collision-free set, we can solve for the robot controller u^* as follows:

$$\begin{aligned} u^*(x, z, t) &= \arg \min_u \|u - \hat{u}(x, t)\|^2 \\ \text{s.t. } L_{f_z} h(z) + L_{g_z} h(z)u &\geq -\alpha(h(z)). \end{aligned} \quad (5)$$

In Section II-B, we show how to build a diffeomorphism ϕ starting from a polyhedral description of the obstacles. The resulting function, together with its inverse and Jacobian, is locally Lipschitz continuous and bounded. As a result, the optimal solution to (5) is locally Lipschitz continuous [31].

B. Building Mappings Between Real and Ball Worlds

In the ball world, obstacles are mapped to convex sets (balls), which result in zero-measure sets of initial conditions from which deadlocks arise [17]. More importantly, for the approach we propose in this article, ball world obstacles can be parameterized by their position and radius. These two parameters will be controlled in our algorithm to ensure the safe motion of the robot in the polyhedral world. To build a computationally tractable diffeomorphism between the polyhedral and the ball world, we leverage QC mappings composed using the navigation functions in [19]. In the following, we will briefly recall the definition of the latter for the case of star-shaped obstacles—which results in an analytic expression of the function—while Section III is devoted to the definition and properties of the QC mappings, which allow us to obtain a mapping from the polyhedral world (where obstacles are polyhedrons) to the ball world.

Consider that the robot state is evolving in the n -dimensional manifold \mathcal{M}^n . To build an analytic diffeomorphism between the real world (containing star-shape obstacles) and the ball world, we will compose a number of real-valued analytic functions, denoted by $\beta_i: \mathcal{M}^n \rightarrow \mathbb{R}$, for which 0 is a regular value. The robot workspace will be denoted by $\mathcal{W} \subset \mathcal{M}^n$, a connected and compact n -dimensional submanifold of \mathcal{M}^n such that

$$\begin{aligned} \mathcal{W}^\circ &\subset \{x \in \mathcal{M}^n : \beta_0(x) > 0\} \\ \partial \mathcal{W} &\subset \{x \in \mathcal{M}^n : \beta_0(x) = 0\} \end{aligned} \quad (6)$$

are its interior and boundary, respectively. It is assumed that M static obstacles are present in the workspace \mathcal{W} . As will be explained in the following, the mapping between the polyhedral and the ball world is updated at each iteration. Therefore, while the motion of the obstacles is not considered explicitly, it is accounted for implicitly. These are denoted by \mathcal{O}_i (where $i = 1, 2, \dots, M$) and correspond to the interior of a connected and compact n -dimensional submanifold of \mathcal{M}^n such that

$$\begin{aligned} \bar{\mathcal{O}}_i &\subset \mathcal{W}^\circ \quad \forall i \\ \mathcal{W} \setminus \bar{\mathcal{O}}_i &\subset \{x \in \mathcal{M}^n : \beta_i(x) > 0\} \\ \partial \bar{\mathcal{O}}_i &\subset \{x \in \mathcal{M}^n : \beta_i(x) = 0\} \\ \bar{\mathcal{O}}_i \cap \bar{\mathcal{O}}_j &= \emptyset \quad \forall i \neq j. \end{aligned} \quad (7)$$

With these definitions, the *safe space* \mathcal{S} (or free space) is given by

$$\mathcal{S} = \mathcal{W} \setminus \bigcup_{i=1}^M \mathcal{O}_i. \quad (8)$$

In the ball world, both the robot workspace and the obstacles are balls. Then, an explicit representation of them can be as follows:

$$\begin{aligned} \hat{\mathcal{O}}_0 &= \left\{ q \in \mathbb{R}^m : \underbrace{\rho_0^2 - \|q - q_0\|^2}_{\hat{\beta}_0(q)} < 0 \right\} \\ \hat{\mathcal{O}}_i &= \left\{ q \in \mathbb{R}^m : \underbrace{\|q - q_i\|^2 - \rho_i^2}_{\hat{\beta}_i(q)} < 0 \right\} \end{aligned} \quad (9)$$

where $\hat{\mathcal{O}}_0$ is the representation of the robot workspace \mathcal{W}° , and q_i and ρ_i , $i = 1, \dots, M$, denote the center and the radius of the i th obstacle, respectively. Thus, the safe space in the ball world is

$$\hat{\mathcal{S}} = \{q \in \mathbb{R}^n : \hat{\beta}_0(q) \geq 0, \hat{\beta}_1(q) \geq 0, \dots, \hat{\beta}_M(q) \geq 0\}. \quad (10)$$

The functions β_i are not easy to obtain in the case of generic shapes of the obstacles. In the proposed approach, however, obstacles are conveniently mapped to balls, making the expression of such functions analytic and simple.

Following the procedure described in [19], we now describe how to obtain an analytic diffeomorphism between the real and the ball world in the case of star-shaped obstacles in the plane. The following is a parameterized example of such a diffeomorphism:

$$\begin{aligned} \phi_\lambda(x) &= \sum_{i=0}^M \sigma_{i,\lambda}(x) \underbrace{(\rho_i b_i(x) + q_i)}_{\phi_i(x) := \text{ith obstacle diffeomorphism}} + \sigma_{g,\lambda}(x) ((x - x_g) + q_g) \end{aligned} \quad (11)$$

where

$$b_i(x) = \frac{\|x - x_i\|}{r_i(\theta)} \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix} \quad (12)$$

with $\theta = \angle(x - x_i)$, x_i is the center of the obstacle, and r_i is the function describing how the radius changes as a function of the angle θ . The functions $\sigma_{i,\lambda}$ are defined as follows:

$$\begin{aligned} \sigma_{i,\lambda} &= \frac{\gamma_g \bar{\beta}_i}{\gamma_g \bar{\beta}_i + \lambda \beta_i}, \quad i = 0, \dots, M \\ \sigma_{g,\lambda} &= 1 - \sum_{i=0}^M \sigma_{i,\lambda} \end{aligned} \quad (13)$$

with $\gamma_g = \|x - x_g\|^2$ and $\bar{\beta}_i = \prod_{\substack{j=0 \\ j \neq i}}^M \beta_j$. Finally, x_g and q_g are goal points in the real and ball world, respectively, which in the context of the algorithm presented in this article can be set to an arbitrary point in the safe space in the real and ball world, respectively. Notice that, given the structure of the analytic diffeomorphism ϕ , the point x_g is mapped to q_g . While this approach works well for obstacles that can be accurately described by star-shaped sets, obtaining a mapping for obstacles described by generic polyhedrons is challenging. In Section III, we will introduce QC mappings and show how to employ them to construct a transformation from a polyhedral to a ball world.

III. SAFE ROBOT NAVIGATION ALGORITHM

A. Full QC Mapping

In this section, we develop a fast method for mapping a polyhedral world to a ball world using conformal mappings and QC mappings. Conformal mappings are angle-preserving mappings that have been commonly used for shape transformations. More specifically, let $\phi : \mathbb{C} \rightarrow \mathbb{C}$ be a map on the complex plane with $\phi(z) = u_\phi(x, y) + i v_\phi(x, y)$, where $z = x + iy$ and u_ϕ, v_ϕ are real-valued functions and i is the imaginary number with $i^2 = -1$. ϕ is conformal if its derivative $\phi'(z)$ is nonzero everywhere and it satisfies the Cauchy–Riemann equations $(\partial u_\phi / \partial x) = (\partial v_\phi / \partial y)$ and $(\partial u_\phi / \partial y) = -(\partial v_\phi / \partial x)$. While conformal mappings can preserve the local geometry under transformations, most of the existing conformal mapping algorithms require a considerable amount of computation and hence are not suitable for real-time robot navigation. Here, we consider a generalization of conformal mappings called the QC mappings, which are planar homeomorphisms that map small circles to small ellipses with bounded eccentricity [32]. Mathematically, a QC mapping $\phi : \Omega_1 \rightarrow \Omega_2$ from a planar domain $\Omega_1 \subset \mathbb{C}$ to another planar domain $\Omega_2 \subset \mathbb{C}$ is an orientation-preserving homeomorphism that satisfies the Beltrami equation $(\partial \phi / \partial \bar{z}) = \mu(z)(\partial \phi / \partial z)$ for some complex-valued function μ with $\|\mu\|_\infty < 1$. We propose a fast QC mapping method based on a recent conformal parametrization algorithm [33] with modifications. Specifically, since the safety-preserving algorithm does not require the mapping in this step to be perfectly conformal, we can relax the conformality requirement of the algorithm in [33] by simplifying certain procedures, thereby producing a QC mapping more efficiently.

Let $\mathcal{S} \subset \mathbb{C}$ be a planar polyhedral domain with k polygonal holes, and denote the boundary as $\partial \mathcal{S} = \Gamma_0 - \Gamma_1 - \dots - \Gamma_k$, where Γ_0 is the outer boundary and $\Gamma_1, \dots, \Gamma_k$ are the inner boundaries. We first fill all k holes and obtain a simply

connected domain $\tilde{\mathcal{S}}$. We can then compute a disk harmonic map $\varphi : \tilde{\mathcal{S}} \rightarrow \mathbb{D}$ by solving the Laplace equation

$$\Delta \varphi = 0 \quad (14)$$

subject to a circular boundary constraint $\varphi(\Gamma_0) = \partial \mathbb{D}$, where Δ is the Laplace–Beltrami operator defined on $\tilde{\mathcal{S}}$ discretized using the cotangent formulation [34] (see [35] for more details). Next, we remove the filled regions and transform the k holes in $\varphi(\mathcal{S})$ into k circles such that the center of each circle is the centroid of the corresponding hole and the area of the circle equals that of the hole (here we remark that alternatively, one may also map the k holes to k circles based on some prescribed centers and radii). Finally, we obtain a QC map $\tilde{\varphi} : \mathcal{S} \rightarrow \mathbb{D}$ using the linear Beltrami solver (LBS) method [36], [37] subject to the updated circular boundary constraints. More specifically, $\tilde{\varphi}$ is obtained by solving the generalized Laplace equation

$$\nabla \cdot (A \cdot \nabla \tilde{\varphi}) = 0 \quad (15)$$

subject to the updated boundary constraints for all Γ_j , where

$$A = \begin{pmatrix} \frac{(\operatorname{Re}(\mu_\varphi) - 1)^2 + (\operatorname{Im}(\mu_\varphi))^2}{1 - |\mu_\varphi|^2} & -\frac{2\operatorname{Im}(\mu_\varphi)}{1 - |\mu_\varphi|^2} \\ -\frac{2\operatorname{Im}(\mu_\varphi)}{1 - |\mu_\varphi|^2} & \frac{(\operatorname{Re}(\mu_\varphi) + 1)^2 + (\operatorname{Im}(\mu_\varphi))^2}{1 - |\mu_\varphi|^2} \end{pmatrix} \quad (16)$$

and $\mu_\varphi = (\varphi_z / \varphi_{\bar{z}})$ is the Beltrami coefficient of the disk harmonic map φ . This completes the proposed fast method for mapping the polyhedral world to the ball world. In the discrete case, note that both (14) and (15) are $n_v \times n_v$ sparse linear systems, where n_v is the number of vertices in the polyhedral domain \mathcal{S} . In other words, the Full QC method only requires solving two linear systems without any iterations and hence is highly efficient. Also, the resulting QC map ϕ has bounded distortion as guaranteed by QC theory.

B. Partial Conformal Mapping

In the proposed method above, the entire polyhedral world is mapped to a ball world by solving linear systems. However, in some cases, it may be desirable to have an even more efficient method for mapping only part of the polyhedral world to a partial ball world, where the computation of the mapping should be as simple as possible for facilitating real-time navigation. Here, we develop an algorithm for conformally mapping a polygonal hole in the polyhedral world to the unit disk and everything outside the hole to the exterior of the unit disk. To achieve this, we adopt a similar strategy as in [38] and utilize the geodesic algorithm [39], which consists of a series of analytic maps on the complex plane.

Denote the boundary vertices of the polygonal hole as p_1, p_2, \dots, p_n (in anticlockwise order). We first consider the following mapping:

$$\varphi_0(z) = \sqrt{\frac{z - p_2}{z - p_1}} \quad (17)$$

with the branching $(-1)^{1/2} = i$, which maps p_1 to ∞ , p_2 to 0, and the interior of the polygon to the right half-plane. We then construct a sequence of mappings $\varphi_1, \dots, \varphi_{n-2}$ such that the

remaining p_i 's are mapped to the imaginary axis one by one via their compositions

$$\varphi_j(z) = \sqrt{\left(\frac{\frac{\operatorname{Re} \xi_j}{|\xi_j|^2} z}{1 + \frac{\operatorname{Im} \xi_j}{|\xi_j|^2} z i}\right)^2 - 1} \quad (18)$$

where $\xi_j = (\varphi_{j-1} \circ \varphi_{j-2} \circ \dots \circ \varphi_0)(p_j)$, with the branching $(-1)^{1/2} = -i$. Next, we apply the following analytic map:

$$\varphi_{n-1}(z) = \left(\frac{z}{1 - \frac{z}{(\varphi_{n-2} \circ \varphi_{n-3} \circ \dots \circ \varphi_0)(p_1)}}\right)^2 \quad (19)$$

which enforces the interior of the polygon to be in the upper half-plane while keeping p_1 to be mapped to ∞ . The Cayley transform

$$\varphi_n(z) = \frac{z - i}{z + i} \quad (20)$$

can then be applied to map the real axis to the unit circle, the upper half-plane to the interior of the unit disk, and the lower half-plane to the exterior of the disk. Next, we normalize the transformation and ensure that ∞ remains fixed by applying a reflection with respect to the unit circle $\eta(z) = (1/\bar{z})$, followed by a Möbius transformation that shifts ∞ to 0:

$$\tau(z) = \frac{z - (\eta \circ \varphi_n \circ \dots \circ \varphi_0)(\infty)}{1 - (\eta \circ \varphi_n \circ \dots \circ \varphi_0)(\infty)z}. \quad (21)$$

Finally, we apply an inverse reflection $\eta^{-1}(z) = (1/\bar{z}) = \eta(z)$.

Altogether, the composition of the above-mentioned mappings $\Psi = \eta^{-1} \circ \tau \circ \eta \circ \varphi_n \circ \dots \circ \varphi_0$ maps the polygonal hole to the unit disk and everything outside the hole to the exterior of the unit disk. Note that Ψ is fully analytic and does not require any equation solving.

C. Convex Optimization Formulation

In this section, we develop the main optimization-based controller, which is designed to move the positions of the obstacles as well as their radii so that the robot in the ball world moves in a collision-free fashion. Leveraging the diffeomorphism built using the techniques recalled in the previous section, the motion of the robot in the ball world is mapped into the motion of the robot in the polyhedral (real) world (see also the Appendix for more implementation details). It is important to remark that the motion of the obstacles in the ball world is artificial and does not correspond to the motion of the obstacles in the polyhedral world. While the approach proposed in this article is capable of accounting for the latter, for the sake of clarity we do not consider it in this work.

As the ball world is a mathematical representation of the real robot workspace, it can be freely deformed and modified to ensure the robot never collides with the obstacles present in the environment and with the boundary of the latter. This approach can be interpreted as a *robot-avoidance paradigm*, as opposed to the more traditional obstacle-avoidance one. In these settings, obstacles in the ball world are displaced from their positions and shrunk with respect to their initial radii so that the following conditions are satisfied in the ball world:

(C1) The robot does not collide with the obstacles.

(C2) The obstacles do not collide with each other.

(C3) The obstacles do not exit the environment.

(C4) The obstacles do not overlap with the goal point q_g used to build the mapping between real and ball worlds.

Condition (C1) is directly related to safety: If the robot is kept in the safe set in the ball world, so is the case in the polyhedral world. Therefore, having the obstacles move away from the robot in the ball world will result in the robot avoiding the obstacles in the polyhedral world. Conditions (C2)–(C4) are introduced to always obtain a valid mapping between real and ball worlds by satisfying the assumptions introduced in Section III-B. Given the simple geometry of the ball world, where every object is a ball, constraint satisfaction for the obstacles can be efficiently enforced through the use of CBF.

In this article, we are going to enforce constraints on the obstacles by leveraging CBF. To do so, we need to define the dynamics of the obstacle motion. We choose to control both the position of the center of each obstacle, q_j , via its velocity, u_{q_j} , as well as the radius of each obstacle, ρ_j , via the radius rate of change u_{ρ_j} . Therefore, we can define the following single integrator obstacle dynamical model:

$$\begin{cases} \dot{q}_j = u_{q_j} \\ \dot{\rho}_j = u_{\rho_j} \end{cases} \quad (22)$$

We are now ready to present the CBF employed to enforce Conditions (C1)–(C4) introduced above.

a) *CBF for Condition (C1)*: The objective of keeping the robot at position q away from obstacle j can be enforced by defining formally the CBF $h_j := \hat{\beta}_j$, that is,

$$h_j(q_j, \rho_j) = \|q_j - q\|^2 - \rho_j^2. \quad (23)$$

The following inequality constraint on u_{q_j} and u_{ρ_j} can be defined to enforce the positivity of h_j [15]:

$$\underbrace{\begin{bmatrix} -2(q_j - q)^\top & 2\rho_j \end{bmatrix}}_{=:A_{C1,j}} \begin{bmatrix} u_{q_j} \\ u_{\rho_j} \end{bmatrix} \leq \underbrace{-2(q_j - q)^\top \dot{q} + \alpha(h_j(q_j, \rho_j))}_{=:b_{C1,j}}. \quad (24)$$

This constraint on the control inputs u_{q_j} and u_{ρ_j} will be enforced in an optimization program formulated to synthesize the controller to move and deform the obstacles in the ball world.

b) *CBF for Condition (C2)*: To make the obstacles not collide with each other, we can proceed to define the following CBF:

$$h_{jk}(q_j, q_k, \rho_j, \rho_k) = \|q_j - q_k\|^2 - (\rho_j + \rho_k)^2 \quad (25)$$

and corresponding inequality constraint for the inputs to the obstacles

$$\underbrace{\begin{bmatrix} -2(q_j - q_k)^\top & 2(q_j - q_k)^\top & 2(\rho_j + \rho_k) & 2(\rho_j + \rho_k) \end{bmatrix}}_{=:A_{C2,jk}} \begin{bmatrix} u_{q_j} \\ u_{q_k} \\ u_{\rho_j} \\ u_{\rho_k} \end{bmatrix} \leq \underbrace{\alpha(h_{jk}(q_j, q_k, \rho_j, \rho_k))}_{=:b_{C2,jk}}. \quad (26)$$

c) *CBF for Condition (C3)*: To prevent the obstacle from exiting the environment, we can proceed similarly to the previous two cases and define the following CBF for each obstacle j :

$$h_{j0}(q_j, \rho_j) = (\rho_0 - \rho_j)^2 - \|q_j - q_0\|^2 \quad (27)$$

resulting in the following constraint on u_{q_j} and u_{ρ_j} :

$$\underbrace{\begin{bmatrix} 2(q_j - q_0)^\top & 2(\rho_0 - \rho_j) \end{bmatrix}}_{=:A_{C3,j}} \begin{bmatrix} u_{q_j} \\ u_{\rho_j} \end{bmatrix} \leq \underbrace{\alpha(h_{j0}(q_j, \rho_j))}_{=:b_{C3,j}}. \quad (28)$$

d) *CBF for Condition (C4)*: Finally, the nonoverlapping condition between any of the obstacles and the goal of the analytic diffeomorphism can be obtained with the following CBF:

$$h_{jg}(q_j, \rho_j) = \|q_j - q_g\|^2 - \rho_j^2 \quad (29)$$

from which the following constraint on u_{q_j} and u_{ρ_j} follows:

$$\underbrace{\begin{bmatrix} -2(q_j - q_g)^\top & 2\rho_j \end{bmatrix}}_{=:A_{C4,j}} \begin{bmatrix} u_{q_j} \\ u_{\rho_j} \end{bmatrix} \leq \underbrace{\alpha(h_{jg}(q_j, \rho_j))}_{=:b_{C4,j}}. \quad (30)$$

At this point, to select the control input for the obstacles, we formulate the following QP:

$$\begin{aligned} & \underset{u_{q_j}, u_{\rho_j}}{\text{minimize}} \quad \|u_{q_j} - \hat{u}_{q_j}\|^2 + \kappa \|u_{\rho_j} - \hat{u}_{\rho_j}\|^2 \\ & \text{subject to} \quad \begin{bmatrix} A_{C1,j} \\ A_{C3,j} \\ A_{C4,j} \end{bmatrix} \begin{bmatrix} u_{q_j} \\ u_{\rho_j} \end{bmatrix} \leq \begin{bmatrix} b_{C1,j} \\ b_{C3,j} \\ b_{C4,j} \end{bmatrix} \quad \forall j = 1, \dots, M \\ & \quad \quad \quad \begin{bmatrix} A_{C2,jk} \\ A_{C2,jk} \\ A_{C2,jk} \end{bmatrix} \begin{bmatrix} u_{q_j} \\ u_{q_k} \\ u_{\rho_j} \end{bmatrix} \leq b_{C2,jk} \quad \forall j, k = 1, \dots, M, j > k \end{aligned} \quad (31)$$

where u_{q_j} and u_{ρ_j} are the stacked obstacle velocity and radius change rate, respectively,

$$\begin{aligned} u_{q_j} &= [u_{q_1}^\top \dots u_{q_M}^\top]^\top \\ u_{\rho_j} &= [u_{\rho_1} \dots u_{\rho_M}]^\top \end{aligned} \quad (32)$$

and \hat{u}_{q_j} and \hat{u}_{ρ_j} are the stacked nominal control inputs for the velocity and the radius change rate of the obstacles defined as follows to keep the original positions $q_j(0)$ and radius $\rho_j(0)$ if possible:

$$\begin{aligned} \hat{u}_{q_j} &= K_p(q_j(0) - q_j(t)) \\ \hat{u}_{\rho_j} &= K_p(\rho_j(0) - \rho_j(t)) \end{aligned} \quad (33)$$

where $K_p > 0$ is a controller gain. The parameter κ determines the relative weight between the change of position and the change of radius of the obstacles and can be chosen to prioritize one over the other. In [30], it is shown that the QP in (31) is always feasible, regardless of the number of obstacles, their relative position to the robot, and the boundary of the environment.

Algorithm 1 summarizes the required steps to execute the safe robot navigation algorithm. The symbol (\star) denotes a

Algorithm 1 Safety With Multiple Obstacles

Require: ϕ , Δt , α , κ , K_p , $q_i(t_0)$ and $\rho_i(t_0)$, $i = 0, \dots, M$

- 1: $k = 0$
- 2: **while** true **do**
- 3: $k \leftarrow k + 1$
- 4: $\dot{q}^{(k)} = L_f \phi^{(k)}(x^{(k)}) + L_g \phi^{(k)}(x^{(k)}) u^{(k)}$
- 5: Compute \hat{u}_q and \hat{u}_ρ ▷ (33)
- 6: Compute u_q^* and u_ρ^* ▷ (31)
- 7: $q_i^{(k+1)} \leftarrow q_i^{(k)} + u_{q_i}^* \Delta t$, $i = 0, \dots, M$
- 8: $\rho_i^{(k+1)} \leftarrow \rho_i^{(k)} + u_{\rho_i}^* \Delta t$, $i = 0, \dots, M$
- 9: Update $\phi^{(k+1)}$
- 10: $\dot{x}^{(k)} = \frac{\partial \phi^{(k+1)^{-1}}}{\partial q} \dot{q}^{(k)}$
- 11: Compute $u^{(k+1)}$ to track $\dot{x}^{(k)}$ ▷ (\star)
- 12: Apply $u^{(k+1)}$
- 13: **end while**

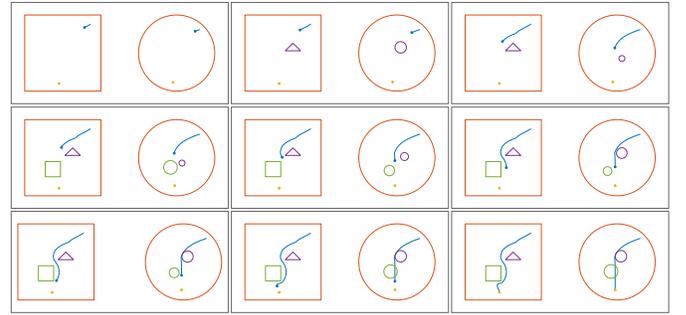


Fig. 2. Full QC mapping (using the inverse mapping to transform the states). In each panel, the left box is the polyhedral (real) world, and the right circle is the ball world. Notice the motion of the obstacles in the ball world to prevent the state of the robot mapped into the ball world from colliding with them at each point in time.

step which, depending on the properties of the system to control, can be executed in different ways. In the Appendix, we illustrate how to implement step (\star) for three classes of DSs which encompass a large variety of, yet not all, robotic systems.

IV. VALIDATION AND COMPARISONS

This section presents an evaluation and a comparison between different variants of the proposed approach. We start by comparing the Partial Conformal and the Full QC mappings, then highlighting the tradeoff between computational complexity and robustness of the approach. Then, we will illustrate the difference between mapping the state of the system between polyhedral and ball world, rather than the control input. The discussion in this section will serve as a guide to the choice of algorithm employed in Section IV.

A. Partial Conformal Versus Full QC Mapping

We provide a simulation comparison between the Partial Conformal (see Section III-B) and the Full QC (see Section III-A) mappings. To this end, we consider the dynamics

$$\dot{x} = -Ax + u = \begin{bmatrix} 60 & 0 \\ 0 & 10 \end{bmatrix} (x_g - x) \quad (34)$$

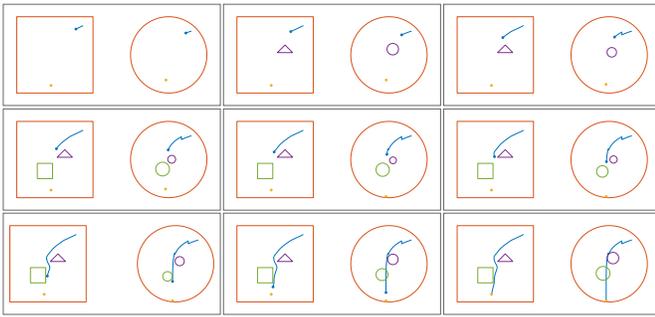


Fig. 3. Partial conformal mapping ($\lambda = 10000$). In each panel, the left box is the polyhedral (real) world, and the right circle is the ball world. Notice the motion of the obstacles in the ball world to prevent the state of the robot mapped into the ball world from colliding with them at each point in time.

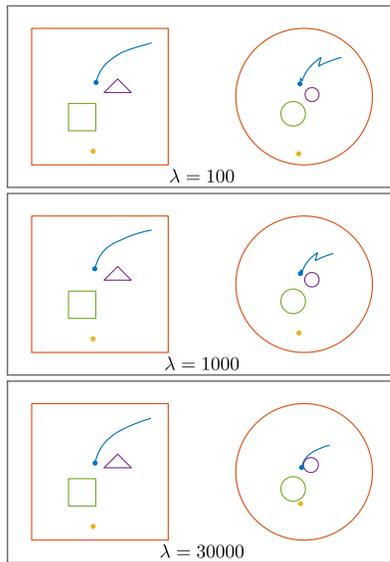
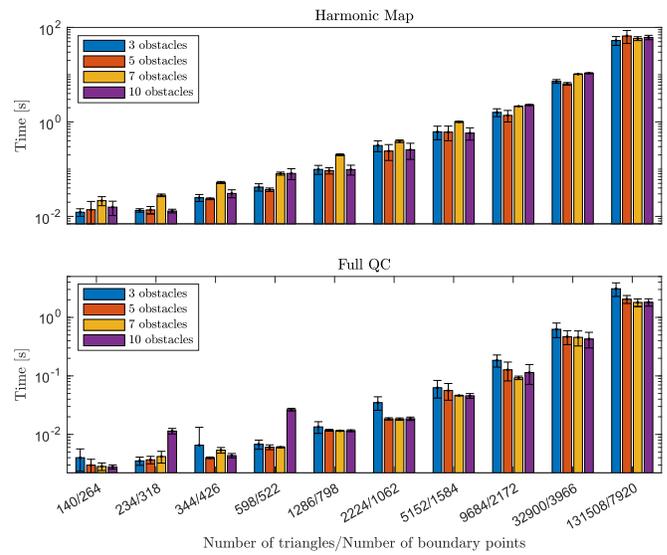
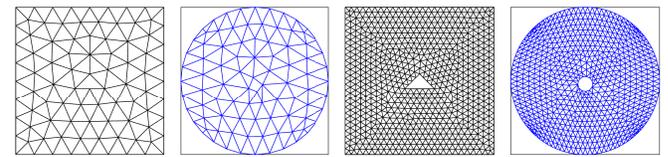


Fig. 4. Partial conformal mapping for different values of λ . In each panel, the left box is the polyhedral (real) world, and the right circle is the ball world.

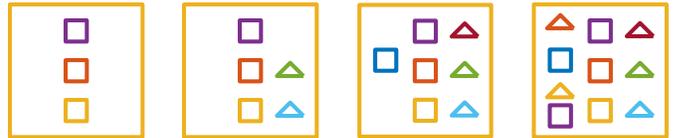
where the 2-D state vector is defined as $x = [x_1, x_2]^T$, the goal state is $x_g = -[0.25, 2]^T$, and $u = Ax_g$. We consider a navigation scenario where the state of the dynamics (34) has to reach x_g starting from $x(0) = [2, 2]^T$ while avoiding obstacles that dynamically appear in the state space (the triangle and the square in Fig. 3). The navigation task is successfully completed using either the Partial (see Fig. 3) or the Full (see Fig. 2) mapping. The Partial is computationally more efficient than the Full mapping, with computational time increasing only when adding new obstacles. However, the Partial mapping is a proper QC mapping only for certain values of λ . Indeed, as shown in Fig. 3, the state in the ball world exhibits a discontinuous behavior when new obstacles are added. The jump can be reduced by reducing the effect of the obstacle at larger distances, that is, by increasing λ as shown in Fig. 4, but this requires manual tuning of λ at run-time which makes the approach application specific. On the other hand, updating the Full mapping at each iteration is computationally more expensive (see Section IV-B), but it always generates a proper QC mapping, and, therefore, smooth state trajectories (see Fig. 2).



(a)



(b)



(c)

Fig. 5. (a) Computational complexity of Harmonic Map [26] and the proposed Full QC for different domain approximations. Note the logarithmic scale on the time axis. (b) Mesh domain approximation with 144 (left) and 10070 (right) triangles. (c) Polyhedral worlds of increasing complexity considered in this comparison.

B. Computation Versus Robustness Tradeoff

In this experiment, we evaluate the computational time of the proposed Full QC mapping, and compare it against the Harmonic Map approach presented in [26]. We consider the same box domain (polyhedral world) of Section IV-A containing two obstacles (a box and a triangle, as shown in Fig. 5). Given this domain, we create a triangular mesh using the built-in MATLAB¹ function `generateMesh`. The function accepts as input the maximum element size h_{\max} that we sample from the vector $[0.05, 0.1, 0.15, 0.2, 0.3, 0.4, 0.5, 0.6]$. As shown in Fig. 5, the value of h_{\max} affects the number of triangles in the mesh, for example, for $h_{\max} = 0.6$, the mesh contains 144 triangles. The Harmonic Map approach approximates the domain and the boundary of each obstacle with a set of points (boundary points). To make the two approaches comparable, we compute the number of boundary points for each value of h_{\max} and use it to compute the Harmonic Map. For each value of h_{\max} (or number of boundary

¹Registered trademark.

points), we compute the Full QC mapping ten times and report statistics in Fig. 5. Results are obtained using MATLAB 2018b on a laptop equipped with an Intel i7 7th generation CPU and 16 GB of DDR4 RAM.

Fig. 5(a) shows the computation time obtained with Harmonic Map and Full QC approaches in the polyhedral domains shown in Fig. 5(c). We observe that the computation time of both Harmonic Map and Full QC is almost independent of the number of obstacles, while it increases with the number of boundary points/triangles used to approximate the domain. For all the considered resolutions, our approach is about one order of magnitude more efficient than the Harmonic Map, which makes it possible to use the Full QC in dynamic scenarios as the one considered in Section IV-A. Fig. 5(b) shows the tradeoff between computational complexity (number of triangles) and robustness margin (size of triangles). Fast computation requires fewer triangles. However, having a too-coarse mesh may result in a significant loss of representation details. As shown in the left panels of Fig. 5(b), with 144 mesh elements, the triangular obstacle is poorly approximated both in the real and the ball worlds. A practical solution could be to set the number of triangles (or, equivalently, h_{\max}) based on the controller loop time and then set the desired error bounds based on the size of the triangles. This will introduce more errors, which can be accounted for in the safety constraints by increasing the safety margin accordingly. Alternatively, one can create a coarse mesh and refine it locally around the obstacles.

C. State Mapping Versus Input Mapping

The safe robot navigation algorithm requires mapping back and forth the state and the control input from the real to the ball world (see Appendix). These mappings can be realized in two ways: 1) using the Jacobian of the transformation and its inverse to map control inputs; or 2) using the transformation (QC mapping) to map states.

To compare the two approaches, we consider the same setup used in Section IV-A. Mapping the state of the robot from the ball world to the polyhedral world (see Fig. 2) is effective only for dynamics with specific properties like differential flatness or feedback linearizability (see Appendix). Instead, mapping control inputs using the Jacobian and its inverse is a more general approach that does not require further assumptions on the dynamics beyond smoothness. Results obtained with this approach are shown in Fig. 6. In the considered scenario, both approaches can successfully complete the navigation task.

When it comes to numerical implementation, the two approaches exhibit different behaviors. In particular, the approach proposed to compute the QC mapping ensures a certain level of smoothness. However, the mapping, especially close to the obstacle boundary, can be significantly deformed. Since the Jacobian is numerically computed using finite differences, this may result in a close to singular Jacobian matrix, and, as a consequence, in an almost unbounded inverse mapping. To prevent these issues, we used a varying-step numerical approach to compute derivatives, but this comes at an extra computational cost. On the other hand, mapping the

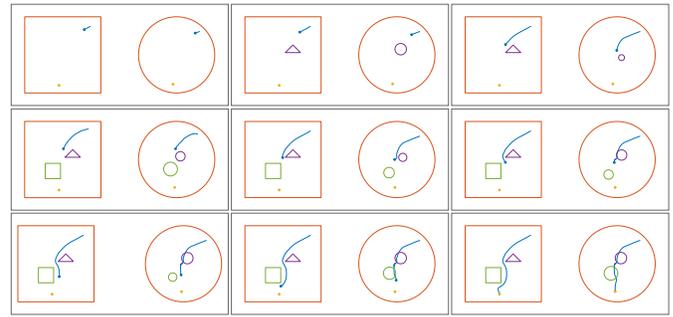


Fig. 6. Full QC input mapping (using the mapping Jacobian to transform inputs). In each panel, the left box is the polyhedral (real) world, and the right circle is the ball world. As for the previous simulations, notice the motion of the obstacles in the ball world to prevent the state of the robot mapped into the ball world from colliding with them at each point in time.

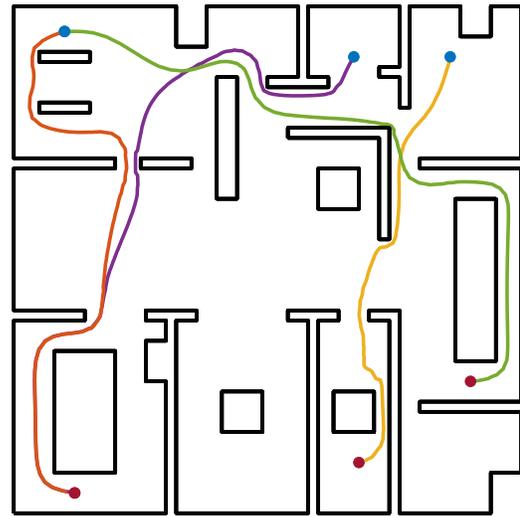


Fig. 7. Results obtained with the Full QC mapping in a cluttered office workspace. Blue/red bullets indicate start/goal positions.

state only requires the mapping and its inverse that are smooth everywhere in the state space.

D. Cluttered Office Workspace

In this experiment, we demonstrate that the Full QC approach can be used to navigate cluttered domains with a complex boundary. To this end, we construct a cluttered office workspace (see Fig. 7) similar to the one used in [26], [29]. This domain is designed to reproduce a realistic office workspace, and it is complex to navigate as it contains ten objects of different shapes and sizes and has complex boundaries. As shown in Fig. 7, the Full QC approach can navigate the office from different initial to different final positions.

V. ROBOTICS APPLICATIONS

This section presents experimental results in typical robotics applications, both in simulation and on real devices.

A. Avoiding Low-Manipulability Areas

In this simulation, we consider a planar two-link manipulator that has to move its end-effector between two points

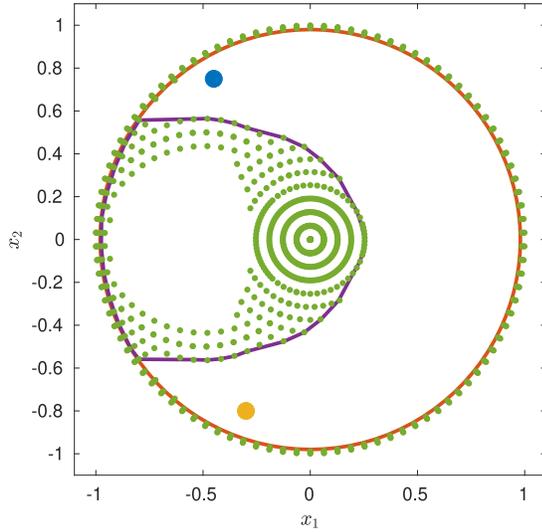


Fig. 8. Workspace of an RR planar manipulator. The blue dot is the starting point, the yellow dot is the goal, and the green dots are regions of low manipulability ($\mu < 0.1$).

(blue and yellow dots in Fig. 8) while avoiding a region of low manipulability (the region within the purple boundary). The manipulator has two links of equal length, that is, $l_1 = l_2 = 0.5$ m. The joint ranges are limited to $q_1, q_2 \in [q, \bar{q}] = [-\pi, \pi]$ rad. Given a joint configuration $q = [q_1, q_2]^T$, and the corresponding manipulator Jacobian $J(q)$, the manipulability is computed as

$$\mu = k_\mu \sqrt{\det(J(q)J^T(q))} \quad (35)$$

where the gain k_μ is used to reduce the manipulability while approaching the joint limits, and it is computed as

$$k_\mu = 1 - \exp\left(-100 \prod_{i=1}^2 \frac{(q_i - \underline{q})(\bar{q} - q_i)}{(\bar{q} - \underline{q})^2}\right). \quad (36)$$

The region of low manipulability (green dots in Fig. 8) is computed by uniformly sampling (50×50 grid) the configuration space of the robot, computing the manipulability μ at each point using (35), and marking the point as low manipulability if $\mu < 0.1$. Then, we create the boundary of the polyhedral world (orange circle in Fig. 8) as a circle of radius 0.98 m. This is enough to exclude low-manipulability configurations corresponding to the manipulator fully stretched. Other low-manipulability areas² (purple boundary in Fig. 8) are considered as an obstacle to be avoided. As shown in Fig. 9, using the Full QC mapping and only reducing the obstacle's radius in the ball world, the safe reactive robot navigation algorithm successfully executes the task (reach a goal joint configuration) while preserving safety (avoid low-manipulability regions).

B. Avoiding Forbidden Areas With a Pan-Tilt Camera

We consider a surveillance task where a pan-tilt camera has to scan the space between two configurations (blue and green

²This boundary is computed using the built-in MATLAB[®] function `boundary`.

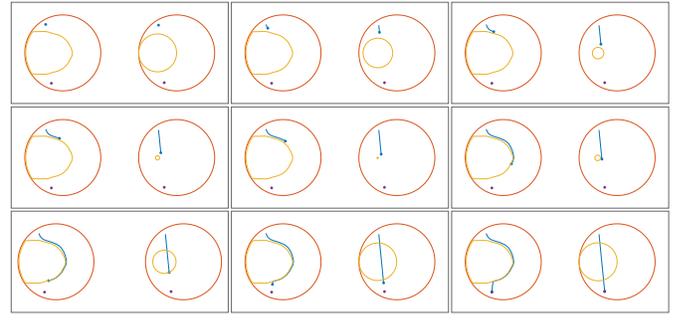


Fig. 9. Avoiding low-manipulability areas with the Full QC mapping. In each panel, the left circle is the polyhedral (real) world, representing the robot's workspace, and the right circle is the ball world. Notice the motion of the obstacles in the ball world to prevent the state of the robot mapped into the ball world from colliding with them at each point in time.

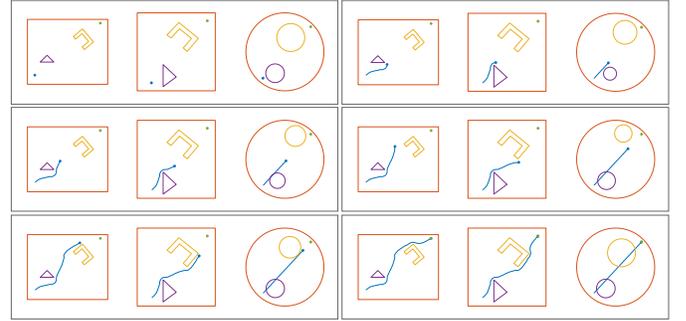


Fig. 10. Avoiding forbidden areas with a pan-tilt camera performing a surveillance task. In each panel, the left rectangle is the xy plane, the middle square is the pan-tilt angles (the polyhedral world), and the circle is the ball world. Also, in this figure, the motion of the obstacles in the ball world is what prevents the state of the robot mapped into the ball world from colliding with them at each point in time.

dots in Fig. 10) while avoiding two forbidden areas (purple triangle and yellow polygon in Fig. 10). In practice, inspecting inside the forbidden areas may correspond to privacy-violating areas, for example, if the area contains the windows of a private building.

The considered pan-tilt camera has three links, namely $l_0 = 0$ m and $l_1 = l_2 = 1$ m. Also, in this case, the joint ranges are limited to $q_1, q_2 \in [q, \bar{q}] = [-\pi, \pi]$ rad. Given the pan-tilt angles $q = [q_1, q_2]^T$, the 3-D position of the camera (end-effector) is computed as

$$p = \begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} l_1 + l_2 \sin(q_2) \\ l_2 \sin(q_1) \cos(q_2) \\ l_0 - l_2 \cos(q_1) \cos(q_2) \end{bmatrix}. \quad (37)$$

The inverse mapping from position p to joint angles q can be computed as

$$q = \begin{bmatrix} q_1 \\ q_2 \end{bmatrix} = \begin{bmatrix} \text{atan2}(y, l_0 - z) \\ \text{atan2}\left(\frac{x - l_1}{l_2}, \sqrt{\frac{(y^2 + (l_0 - z)^2)}{l_2^2}}\right) \end{bmatrix}. \quad (38)$$

The forbidden areas in Fig. 10 are defined in task space. To obtain the polyhedral world in Fig. 10, we map the forbidden areas to the configuration space using (38). The polyhedral world is then mapped into the ball world using the Full QC mapping. Results in Fig. 10 show that the proposed

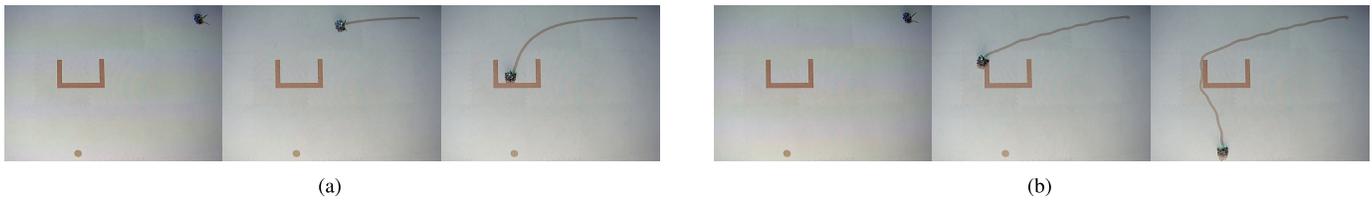


Fig. 11. Snapshots of the Robotarium experiments. (a) Results obtained without mapping the free space in the polyhedral world to the free space in a ball world. The controller preserves only the safety. (b) Results obtained using the Partial Conformal mapping. The controller preserves both safety and stability.

safe robot navigation algorithm successfully executes the task (reach a goal Cartesian position) while preserving safety (avoid forbidden areas). The experiment also shows that the proposed approach can be combined with other smooth mappings [in this case, the inverse kinematics in (38)] to solve specific problems.

C. Navigation in the Robotarium

In this experiment, we consider a mobile robot that has to navigate between two points while avoiding a concave obstacle (see Fig. 11). Experiments are performed inside the Robotarium [40]. Using the traditional CBF-QP formulation directly in the polyhedral world generates an undesired equilibrium point [see Fig. 11(a)], which is eliminated using the proposed approach [see Fig. 11(b)].

VI. CONCLUSION

In this article, we presented a computationally efficient approach for safe reactive robot navigation that exploits QC mappings and CBFs. QC mappings are used to compute a smooth mapping between the state space of the robot, containing possibly multiple and nonconvex unsafe regions, into a space where the unsafe regions are either closed balls or the complement of open balls. QC mappings allow us to have mild assumptions on the shape of the unsafe regions and, importantly, do not require an analytical representation of the unsafe regions. We presented and analyzed two versions of the proposed mapping, a Partial and a Full one. The Full mapping updates the entire workspace at each iteration of the control loop, while the Partial mapping computes the transformation for each unsafe region independently and then combines them. The Full mapping is computationally more intensive, but the resulting mapping is guaranteed to be a diffeomorphism. On the contrary, Partial mapping is more efficient—the mapping can be locally updated as novel unsafe regions appear or disappear, but it requires careful tuning of a parameter to ensure a proper diffeomorphic mapping. Given the mapping, we design a controller in a ball world—a transformed robot workspace where all obstacles are closed balls—that shifts and shrinks unsafe ball-shaped regions to ensure safe navigation, relying on the CBF formalism. The controller and the two mapping strategies have been validated through simulated and real experiments, showing promising results.

Our future research will focus on improving the efficiency of the Full QC mapping. A possibility is to consider meshes of

different granularity, that is, coarse in the free space and fine close to unsafe regions. We also plan to test our algorithm in the presence of partial occlusions, moving obstacles, and dynamic models of robotic systems. Finally, although our proposed algorithm applies to a large variety of DS models, we are interested in extending it to design a computationally efficient safety layer for existing tracking controllers.

APPENDIX

This section shows how the reactive navigation strategy developed in this article can be applied to a large variety of robotic systems.

A. Dynamics (Near-Identity) Diffeomorphic to a Single Integrator

For systems that are diffeomorphic (resp. near-identity diffeomorphic) to a single integrator [41], one can apply a (nonlinear) change of coordinates to the state such the transformed state (resp. part of the state) has linear dynamics. As an example of a system near-identity diffeomorphic to a single integrator, consider the following unicycle dynamics:

$$\begin{cases} \dot{p}_x = v \cos \theta \\ \dot{p}_y = v \sin \theta \\ \dot{\theta} = \omega \end{cases} \quad (39)$$

where p_x and p_y are the positions of the unicycle, θ its orientation, and v and ω are the longitudinal and angular velocity inputs, respectively. Using the change of coordinates

$$z = \begin{bmatrix} p_x \\ p_y \end{bmatrix} + l \begin{bmatrix} \cos \theta \\ \sin \theta \end{bmatrix}$$

leads to

$$\dot{z} = \begin{bmatrix} \cos \theta & -l \sin \theta \\ \sin \theta & l \cos \theta \end{bmatrix} \begin{bmatrix} v \\ \omega \end{bmatrix}. \quad (40)$$

The dynamics of z may be controlled using single-integrator dynamics to achieve the desired $\dot{x}^{(k)}$ in Step 10 of Algorithm 1, from which v and ω can be computed using (40) to control the unicycle dynamics. This approach was taken in Section V-C to implement the safe robot navigation algorithm using a differential drive robot.

B. Feedback Linearizable Dynamics

For feedback linearizable system dynamics [42], one can apply a (nonlinear) change of coordinates $y = T(x)$ such that there exists a (nonlinear) feedback control law that cancels the nonlinearities of the system. The feedback linearized dynamics are represented by a chain of integrators. As far as the application of the approach presented in this article is concerned, this class of systems represents a generalization of systems (near-identity) diffeomorphic to a single integrator. Step 11 in Algorithm 1 can be executed by designing a stabilizing state-feedback controller with a feedforward term equal to the desired $\dot{x}^{(k)}$, given in Step 10 of Algorithm 1.

C. Differentially Flat Dynamics

For differentially flat systems, there exists an output thereof—the flat output—such that their state and input can be determined from the output without integration [43]. This way, we can proceed similarly to the case of feedback linearizable dynamics and design a stabilizing flat-output-feedback controller with a feedforward term equal to the desired $\dot{x}^{(k)}$, given in Step 10 of Algorithm 1.

REFERENCES

- [1] L. Wang, A. D. Ames, and M. Egerstedt, "Safety barrier certificates for collisions-free multirobot systems," *IEEE Trans. Robot.*, vol. 33, no. 3, pp. 661–674, Jun. 2017.
- [2] G. Notomista and M. Egerstedt, "Persistification of robotic tasks," *IEEE Trans. Control Syst. Technol.*, vol. 29, no. 2, pp. 756–767, Mar. 2021.
- [3] A. Aswani, H. Gonzalez, S. S. Sastry, and C. Tomlin, "Provably safe and robust learning-based model predictive control," *Automatica*, vol. 49, no. 5, pp. 1216–1226, 2013.
- [4] M. Saveriano and D. Lee, "Learning barrier functions for constrained motion planning with dynamical systems," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Jul. 2019, pp. 112–119.
- [5] S. Petti and T. Fraichard, "Safe motion planning in dynamic environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Aug. 2005, pp. 2210–2215.
- [6] S. G. Loizou, "The navigation transformation," *IEEE Trans. Robot.*, vol. 33, no. 6, pp. 1516–1523, Dec. 2017.
- [7] E. H. Thyri, E. A. Basso, M. Breivik, K. Y. Pettersen, R. Skjetne, and A. M. Lekkass, "Reactive collision avoidance for ASVs based on control barrier functions," in *Proc. Conf. Control Technol. Appl. (CCTA)*, 2020, pp. 380–387.
- [8] S. Jha, V. Raman, D. Sadigh, and S. A. Seshia, "Safe autonomy under perception uncertainty using chance-constrained temporal logic," *J. Automated Reasoning*, vol. 60, no. 1, pp. 43–62, Jan. 2018.
- [9] G. Notomista, "Long-duration robot autonomy: From control algorithms to robot design," Ph.D. dissertation, Georgia Inst. Technol., Atlanta, GA, USA, 2020.
- [10] R. Cheng, G. Orosz, R. M. Murray, and J. W. Burdick, "End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks," in *Proc. AAAI Conf. Artif. Intell.*, 2019, vol. 33, no. 1, pp. 3387–3395.
- [11] Y. Emam, G. Notomista, P. Glotfelter, Z. Kira, and M. Egerstedt, "Safe reinforcement learning using robust control barrier functions," *IEEE Robot. Autom. Lett.*, early access, Oct. 25, 2022, doi: 10.1109/LRA.2022.3216996.
- [12] M. Ohnishi, L. Wang, G. Notomista, and M. Egerstedt, "Barrier-certified adaptive reinforcement learning with applications to brushbot navigation," *IEEE Trans. Robot.*, vol. 35, no. 5, pp. 1186–1205, Oct. 2019.
- [13] G. Notomista, "A constrained-optimization approach to the execution of prioritized stacks of learned multi-robot tasks," 2023, *arXiv:2301.05346*.
- [14] K.-C. Hsu, H. Hu, and J. F. Fisac, "The safety filter: A unified view of safety-critical control in autonomous systems," *Annu. Rev. Control, Robot., Auto. Syst.*, vol. 7, no. 1, pp. 47–72, Jul. 2024.
- [15] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proc. IEEE 18th Eur. Control Conf. (ECC)*, Oct. 2019, pp. 3420–3431.
- [16] J. Breeden and D. Panagou, "Predictive control barrier functions for online safety critical control," in *Proc. IEEE 61st Conf. Decis. Control (CDC)*, Dec. 2022, pp. 924–931.
- [17] M. F. Reis, A. P. Aguiar, and P. Tabuada, "Control barrier function-based quadratic programs introduce undesirable asymptotically stable equilibria," *IEEE Control Syst. Lett.*, vol. 5, no. 2, pp. 731–736, Apr. 2021.
- [18] E. Romon and D. E. Koditschek, "Exact robot navigation in geometrically complicated but topologically simple spaces," in *Proc. IEEE Int. Conf. Robot. Autom.*, Jun. 1990, pp. 1937–1942.
- [19] E. Rimon and D. E. Koditschek, "The construction of analytic diffeomorphisms for exact robot navigation on star worlds," *Trans. Amer. Math. Soc.*, vol. 327, no. 1, pp. 71–116, 1991.
- [20] E. Rimon and D. E. Koditschek, "Exact robot navigation using artificial potential functions," *IEEE Trans. Robot. Autom.*, vol. 8, no. 5, pp. 501–518, Oct. 1992.
- [21] C. Belta, V. Isler, and G. J. Pappas, "Discrete abstractions for robot motion planning and control in polygonal environments," *IEEE Trans. Robot.*, vol. 21, no. 5, pp. 864–874, Oct. 2005.
- [22] P. Rousseas, C. Bechlioulis, and K. J. Kyriakopoulos, "Harmonic-based optimal motion planning in constrained workspaces using reinforcement learning," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 2005–2011, Apr. 2021.
- [23] C. Goerzen, Z. Kong, and B. Mettler, "A survey of motion planning algorithms from the perspective of autonomous UAV guidance," *J. Intell. Robot. Syst.*, vol. 57, nos. 1–4, pp. 65–100, Jan. 2010.
- [24] R. Sarkar, X. Yin, J. Gao, F. Luo, and X. D. Gu, "Greedy routing with guaranteed delivery using ricci flows," in *Proc. Int. Conf. Inf. Process. Sensor Netw.*, Apr. 2009, pp. 121–132.
- [25] N. Constantinou and S. G. Loizou, "Robot navigation on star worlds using a single-step navigation transformation," in *Proc. 59th IEEE Conf. Decis. Control (CDC)*, Dec. 2020, pp. 1537–1542.
- [26] P. Vlantis, C. Vrohidis, C. P. Bechlioulis, and K. J. Kyriakopoulos, "Robot navigation in complex workspaces using harmonic maps," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 1726–1731.
- [27] S. Gao and N. Bezzo, "A conformal mapping-based framework for robot-to-robot and sim-to-real transfer learning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Sep. 2021, pp. 1289–1295.
- [28] G. Notomista and M. Egerstedt, "Coverage control for wire-traversing robots," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 5042–5047.
- [29] L. Fan, J. Liu, W. Zhang, and P. Xu, "Robot navigation in complex workspaces using conformal navigation transformations," *IEEE Robot. Autom. Lett.*, vol. 8, no. 1, pp. 192–199, Jan. 2023.
- [30] 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.
- [31] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Trans. Autom. Control*, vol. 62, no. 8, pp. 3861–3876, Aug. 2017.
- [32] O. Lehto, *Quasiconformal Mappings in the Plane*, vol. 126. Cham, Switzerland: Springer, 1973.
- [33] G. P. T. Choi, "Efficient conformal parameterization of multiply-connected surfaces using quasi-conformal theory," *J. Scientific Comput.*, vol. 87, no. 3, pp. 1–19, Jun. 2021.
- [34] U. Pinkall and K. Polthier, "Computing discrete minimal surfaces and their conjugates," *Experim. Math.*, vol. 2, no. 1, pp. 15–36, Jan. 1993.
- [35] P. T. Choi and L. M. Lui, "Fast disk conformal parameterization of simply-connected open surfaces," *J. Sci. Comput.*, vol. 65, no. 3, pp. 1065–1090, Dec. 2015.
- [36] L. M. Lui, K. C. Lam, T. W. Wong, and X. Gu, "Texture map and video compression using Beltrami representation," *SIAM J. Imag. Sci.*, vol. 6, no. 4, pp. 1880–1902, Jan. 2013.
- [37] P. T. Choi, K. C. Lam, and L. M. Lui, "FLASH: Fast landmark aligned spherical harmonic parameterization for genus-0 closed brain surfaces," *SIAM J. Imag. Sci.*, vol. 8, no. 1, pp. 67–94, Jan. 2015.
- [38] G. P. T. Choi, Y. Leung-Liu, X. Gu, and L. M. Lui, "Parallelizable global conformal parameterization of simply-connected surfaces via partial welding," *SIAM J. Imag. Sci.*, vol. 13, no. 3, pp. 1049–1083, Jan. 2020.
- [39] D. E. Marshall and S. Rohde, "Convergence of a variant of the zipper algorithm for conformal mapping," *SIAM J. Numer. Anal.*, vol. 45, no. 6, pp. 2577–2609, Jan. 2007.

- [40] S. Wilson et al., “The Robotarium: Globally impactful opportunities, challenges, and lessons learned in remote-access, distributed control of multirobot systems,” *IEEE Control Syst. Mag.*, vol. 40, no. 1, pp. 26–44, Feb. 2020.
- [41] R. Olfati-Saber, “Near-identity diffeomorphisms and exponential ϵ -tracking and ϵ -stabilization of first-order nonholonomic SE(2) vehicles,” in *Proc. Amer. Control Conf.*, Jun. 2002, pp. 4690–4695.
- [42] H. K. Khalil, *Nonlinear Control*. New York, NY, USA: Pearson, 2015.
- [43] M. Fliess, J. Lévine, P. Martin, and P. Rouchon, “Flatness and defect of non-linear systems: Introductory theory and examples,” *Int. J. Control*, vol. 61, no. 6, pp. 1327–1361, 1995.



Gennaro Notomista (Member, IEEE) received the B.S. degree in mechanical engineering from the Università degli Studi di Napoli Federico II, Naples, Italy, in 2012, the M.Eng. degree in automotive engineering from the Technische Hochschule Ingolstadt, Ingolstadt, Germany, in 2015, the M.S. degree in mechanical engineering from the Università degli Studi di Napoli Federico II, in 2016, and the M.S. degree in mathematics and the Ph.D. degree in robotics from Georgia Institute of Technology, Atlanta, GA, USA, in 2019 and 2020, respectively.

He is currently an Assistant Professor and the Varma Family Professor in robotics with the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, ON, Canada. Before joining the University of Waterloo, he was a Post-Doctoral Researcher at the CNRS/Inria/IRISA, Rennes, France. His main research interests include the intersection of design and control of robotic systems for long-duration autonomy with applications to environmental monitoring.

Dr. Notomista is a Fulbright Scholar and was a recipient of the Alumni Small Grant in 2020 and the IEEE ARSO Best Paper Award in 2022. He is an Associate Editor of *IEEE ROBOTICS AND AUTOMATION LETTERS* and *The International Journal of Robotics Research*.



Gary P. T. Choi received the Ph.D. degree in applied mathematics from Harvard University, Cambridge, MA, USA, in 2020.

He is a Vice-Chancellor Assistant Professor with the Department of Mathematics, The Chinese University of Hong Kong, Hong Kong. After his Ph.D., he worked as an NSF Post-Doctoral Fellow and Instructor in applied mathematics at the Massachusetts Institute of Technology, Cambridge, from 2020 to 2023. His research interests include applied and computational geometry, interdisciplinary mathematical modeling, metamaterial design, quantitative biology, medical imaging, geometry processing, and scientific computing.



Matteo Saveriano (Senior Member, IEEE) received the B.Sc. and M.Sc. degrees from the University of Naples “Federico II,” Naples, Italy, in 2008 and 2011, respectively, and the Ph.D. degree from the Technical University of Munich, Munich, Germany, in 2017.

He is an Associate Professor of control engineering with the Department of Industrial Engineering, University of Trento, Trento, Italy. After his Ph.D., he was a Post-Doctoral Researcher at the German Aerospace Center (DLR) and a Tenure-Track Assistant Professor at the Department of Computer Science and at the Digital Science Center, University of Innsbruck, Innsbruck, Austria. He is the coordinator of the HE EU Project INVERSE (GA 101136067). He has (co-)authored more than 70 scientific papers in international journals and conferences. His research interests include the intersection between learning and control and attempts to integrate cognitive robots into smart factories and social Environments through the embodiment of AI solutions, inspired by human behavior into robotic devices.

Dr. Saveriano serves as an Associate Editor for the main robotics conferences (ICRA, IROS, and HUMANOIDS), for *IEEE ROBOTICS AND AUTOMATION LETTERS*, and *The International Journal of Robotics Research*.