

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

Contact-Aware Safety in Soft Robots Using High-Order Control Barrier and Lyapunov Functions

Kiwan Wong^{1,2}, Maximilian Stölzle^{1,2,3}, Wei Xiao^{4,2}, Cosimo Della Santina³, Daniela Rus^{*,2}, Gioele Zardini^{*,1}

Abstract—Robots operating alongside people, particularly in sensitive scenarios such as aiding the elderly with daily tasks or collaborating with workers in manufacturing, must guarantee safety and cultivate user trust. Continuum soft manipulators promise safety through material compliance, but as designs evolve for greater precision, payload capacity, and speed, and increasingly incorporate rigid elements, their injury risk resurfaces. In this letter, we introduce a comprehensive High-Order Control Barrier Function (HOCBF) + High-Order Control Lyapunov Function (HOCLF) framework that enforces strict contact force limits across the entire soft-robot body during environmental interactions. Our approach combines a differentiable Piecewise Cosserat-Segment (PCS) dynamics model with a convex-polygon distance approximation metric, named Differentiable Conservative Separating Axis Theorem (DCSAT), based on the soft robot geometry to enable real-time, whole-body collision detection, resolution, and enforcement of the safety constraints. By embedding HOCBFs into our optimization routine, we guarantee safety, allowing, for instance, safe navigation in operational space under HOCLF-driven motion objectives. Extensive planar simulations demonstrate that our method maintains safety-bounded contacts while achieving precise shape and task-space regulation. This work thus lays a foundation for the deployment of soft robots in human-centric environments with provable safety and performance.

Index Terms—Modeling and Control for Soft Robots, Robot Safety, Soft Robot Applications

I. INTRODUCTION

DEPLOYING robots in human-centered environments, such as assisting workers in manufacturing or supporting older adults in everyday activities [1], demands not only demonstrable safety but also user confidence in the robot’s behavior. Traditional rigid collaborative manipulators address this need through increasingly sophisticated algorithms for

Manuscript received: May 4, 2025; Revised August 6, 2025; Accepted September 24, 2025.

This paper was recommended for publication by Editor Cecilia Laschi upon evaluation of the Associate Editor and Reviewers’ comments.

*D. Rus and G. Zardini contributed equally as joint senior authors.

¹Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139, USA {kiwan588, mstolzle, gzardini}@mit.edu. ²Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, USA {weixy, rus}@mit.edu. ³Cognitive Robotics, Delft University of Technology, Delft, 2628 CD, Netherlands {M.W.Stolzle, C.DellaSantina}@tudelft.nl. ⁴Robotics Engineering Department, Worcester Polytechnic Institute, Worcester, MA 01609, USA wxiao3@wpi.edu.

The work by K. Wong was supported by The Hong Kong Jockey Club Scholarships; The work by M. Stölzle was supported under the European Union’s Horizon Europe Program from Project EMERGE - Grant Agreement No. 101070918, and by the Cultuurfonds Wetenschapsbeurzen 2024 and the Rudge (1948) and Nancy Allen Chair for his research visit to LIDS/Zardini Lab at MIT.

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

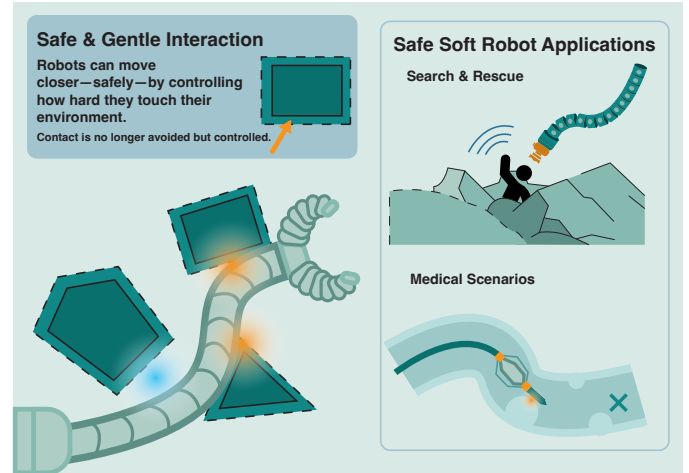


Fig. 1: Contact Safety-Aware Control of Soft Robots with High-Order CBFs (HOCBFs) and High-Order CLFs (HOCLFs). Illustration of compliant contact control with safety bounds guaranteed by HOCBFs. By respecting contact force limits, the robot can intentionally engage with its surroundings without sacrificing safety. Task goals are shaped through HOCLFs, while constraint satisfaction is upheld by HOCBFs. This approach enables the safe use of soft robots in demanding settings—from search-and-rescue missions to delicate medical procedures.

collision detection [2], impedance control [3], Model Predictive Control (MPC) [4], more recently, Control Barrier Functions (CBFs) [5], [6], and the integration of Lyapunov-based methods with reinforcement learning to ensure robotic safety [7]. Yet, perception errors or model inaccuracies can still expose users to hazardous impacts.

Continuum soft manipulators offer a fundamentally different path to safety: instead of relying solely on software, they seek to embed safety directly through material and structural compliance [8]. However, material softness is not a panacea [9], [10]. As the field advances toward greater precision and functionality, emerging designs are expected to incorporate increased stiffness, exert larger forces and velocities [11], and adopt hybrid rigid-soft architectures [12]. Such developments reintroduce risks traditionally associated with rigid systems. Thus, mechanical compliance must be augmented with algorithmic guarantees that ensure real-time safety and foster user trust.

Our approach embraces, rather than avoids, physical contact¹ [10], [13], [14], deviating from the predominant paradigm in rigid robotics [2], [15]. Instead of treating contact as a failure, we exploit the soft robot’s embodied intelligence [16]—its intrinsic physical coupling with the environment—to enhance robustness, stabilize deformation and motion, and adapt to

¹Please note that we use “contact” and “collision” interchangeably.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

external constraints, all while ensuring that every interaction respects safety standards such as the injury severity thresholds set out in ISO/TS 15066:2016 [15].

To date, no method explicitly enforces upper bounds on the contact force or pressure applied across the entire surface of a soft robot, while accounting for the system's inertia. Classical impedance and force control schemes cannot impose strict bounds [17], and CBF-based methods [12] address only self-contact avoidance. A recently published paper [10] adapts CBFs as a safety filter to constrain end-effector forces, but it neglects distributed body interactions and relies on a simplified template model that approximates the continuum bending behavior with an articulated chain of masses connected by prismatic joints [17] while neglecting important strains such as shear or elongation. Instead, Xu *et al.* [14] measures contact forces between the soft robot body and obstacles in the environment, but defines a kinematics-aware CBF instead of a dynamics-aware HOCBF; therefore, it neglects the dynamics of the system, which can, in turn, cause safety issues.

To fill this research gap, we present an integrated control framework that imposes contact-force limits along the entire body of a soft manipulator based on differentiable dynamic strain models. Building on the well-established CBF+Control Lyapunov Function (CLF) framework [5]—specifically its high-order extension [18]—our method optimizes for a control objective encoded in an HOCLF while keeping the system's trajectory inside a certified safe set as specified by HOCBF constraints by solving a constrained Quadratic Program (QP) online. The proposed HOCBF+HOCLF control scheme rests on two pillars: (i) a fully differentiable implementation of the Piecewise Constant Strain (PCS) [19], [20], and (ii) a fast, differentiable collision detection & resolution routine that represents the soft manipulator with convex-polygon approximations.

In support of pillar (ii), we propose a new convex-polygon distance measure—Differentiable Conservative SAT (DC-SAT)—that serves as a conservative, differentiable proxy for the standard Separating Axis Theorem (SAT) metric. Compared with recent differentiable SAT surrogates such as Smooth SAT (SSAT) [21], our approach (1) systematically underestimates the separation distance, yielding the conservative buffer required for formal safety guarantees, and (2) achieves an increase of roughly 1.5–3× in computational efficiency, enabling real-time, full-body collision checks. We then validate the proposed framework through extensive simulations in a planar setting.

In summary, this letter (i) develops a principled HOCBF-based method to enforce global contact force constraints on soft robots, (ii) adapts the HOCBF+HOCLF framework for operational space regulation supporting navigation tasks, and (iii) presents DCSAT, a new, fast, and conservative differentiable collision-detection method for arbitrary convex geometries, improving the feasibility of real-time use for full-body safety guarantees. A video attachment is available on YouTube² and the code is open-sourced on GitHub³.

²<https://youtu.be/ahUhVXiRDPE>

³<https://edu.nl/yekue>

II. BACKGROUND

This section introduces the background necessary for introducing the methodology and, later, the baseline methods. For clarity and simplicity, we focus on the planar case throughout this letter, though the framework naturally generalizes to 3D scenarios.

A. Soft Robotic Kinematics and Dynamics

We model the soft robot kinematics using the PCS formulation [19], which approximates the continuous backbone by discretizing it into N segments, where each segment exhibits the spatially constant strain $\xi_i \in \mathbb{R}^3$. The robot configuration \mathbf{q} is then defined as the deviation of these strains from their equilibrium values, yielding a generalized coordinate vector $\mathbf{q} \in \mathbb{R}^{3N}$. Based on this kinematic model, the forward kinematics map $\chi = \text{FK}(\mathbf{q}, s) : \mathbb{R}^{3N} \times (0, L] \rightarrow SE(2)$ returns the Cartesian position $\chi = [\theta \ p_x \ p_y]^\top$ at a given arc-length position s along the backbone, where $\theta \in [-\pi, \pi)$ is the planar orientation, p_x, p_y are the x- and y-positions, and $L \in \mathbb{R}$ is the total arc length of the robot's centerline. Differentiating $\text{FK}(\mathbf{q}, s)$ with respect to the configuration results in the Jacobian $\mathbf{J}(\mathbf{q}, s) = \frac{\partial \text{FK}(\mathbf{q}, s)}{\partial \mathbf{q}} \in \mathbb{R}^{3 \times 3N}$.

The corresponding dynamics can be derived leveraging established multibody modeling procedures [22], resulting in the following equations of motion:

$$\underbrace{M(\mathbf{q})\ddot{\mathbf{q}} + C(\mathbf{q}, \dot{\mathbf{q}})\dot{\mathbf{q}} + G(\mathbf{q})}_{\text{Multibody dynamics}} + \underbrace{K\mathbf{q} + D\dot{\mathbf{q}}}_{\text{Elast. and Diss. Forces}} = \underbrace{A(\mathbf{q})\mathbf{u}}_{\text{Act. Model}} + \underbrace{\tau_c}_{\text{Contact}}, \quad (1)$$

where $M(\mathbf{q}), C(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{3N \times 3N}$ are the mass and Coriolis matrices, respectively, $G(\mathbf{q})$ captures gravitational effects, and $K, D \in \mathbb{R}^{3N \times 3N}$ are the linear stiffness and damping matrices, respectively. The actuation matrix $A(\mathbf{q}) \in \mathbb{R}^{3N \times m}$ maps the actuation/control input $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^m$ into generalized torques and $\tau_c \in \mathbb{R}^{3N}$ accounts for forces and torques generated by contact with the environment.

Furthermore, $\mathbf{J}_{o,M}^+(\mathbf{q}) \in \mathbb{R}^{3N \times o}$ is the dynamically consistent pseudo-inverse of the Jacobian \mathbf{J}_o that maps into the operational space $\mathbf{o} \in \mathbb{R}^o$ [3], [17], [23], [24].

B. Model-Based Operational Space Controller

A model-based setpoint regulator $\mathbf{u}_o(\mathbf{q}, \dot{\mathbf{q}}) : \mathbb{R}^{6N} \rightarrow \mathbb{R}^{3N}$ that drives the system towards the desired operational space reference \mathbf{o}^d in exponential time can be designed based on the operational space dynamics [3], [17], [23], [24]

$$\mathbf{f} = \underbrace{K_p e_o + K_i \int e_o dt + K_d \dot{e}_o}_{\text{Operational Space PID}} + \underbrace{J_{o,M}^+(K_q + D\dot{\mathbf{q}})}_{\text{Cancel. Elastic \& Diss. Forces}}, \quad (2)$$

where $\mathbf{f} \in \mathbb{R}^o$ is a Cartesian force that serves as the control input in operational-space, $e_o(t) = \mathbf{o}^d(t) - \mathbf{o}(t)$ is the regulation error in operational space, and $K_p, K_i, K_d \in \mathbb{R}^{o \times o}$ are the corresponding PID gain matrices.

C. High Order Control Barrier Functions and High Order Control Lyapunov Functions

The soft robot dynamics can be expressed in control-affine form as an ODE:

$$\dot{\mathbf{x}} = \underbrace{\begin{bmatrix} \dot{\mathbf{q}} \\ M^{-1}(\tau_c - C\dot{\mathbf{q}} - G + K\mathbf{q} - D\dot{\mathbf{q}}) \end{bmatrix}}_{f(\mathbf{x})} + \underbrace{\begin{bmatrix} \mathbf{0}_{3N \times m} \\ M^{-1}A(\mathbf{q}) \end{bmatrix}}_{g(\mathbf{x})} \mathbf{u}, \quad (3)$$

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

with the locally Lipschitz continuous dynamics functions $f(\mathbf{x}) : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$, where $\mathbf{x} = [\mathbf{q}^\top \ \dot{\mathbf{q}}^\top]^\top \in \mathbb{R}^n$ with $n = 6N$ is soft robot state with the corresponding time derivative $\dot{\mathbf{x}} \in \mathbb{R}^n$. In many soft robotic control tasks, constraints specified in operational space do not yield an explicit dependence on the input \mathbf{u} after a single time derivative. This motivates the use of high-order extensions of CLFs and CBFs, where constraints are enforced on higher-order derivatives that expose the control input explicitly [18]. Given a function b with relative degree r , one can recursively build an HOCBF and HOCLF as follows [18].

Definition 1 (High-Order Control Barrier Function [18]): Let $b : \mathbb{R}^n \rightarrow \mathbb{R}$ be r -times differentiable with relative degree r . Define recursively:

$\psi_0 := b$, $\psi_i := L_f \psi_{i-1} + \alpha_i(\psi_{i-1})$, $i = 1, \dots, r-1$, where each α_i is class- \mathcal{K} . Let $C_i := \{\mathbf{x} \mid \psi_{i-1}(\mathbf{x}) \geq 0\}$. Then b is a HOCBF if there exists $\alpha_r \in \mathcal{K}$ such that:

$$\sup_{\mathbf{u} \in \mathcal{U}} [L_f^r b + L_g L_f^{r-1} b \mathbf{u} + O(b) + \alpha_r(\psi_{r-1})] \geq 0, \quad \forall \mathbf{x} \in \bigcap_{i=1}^r C_i.$$

Definition 2 (High-Order Control Lyapunov Function): Let $V : \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable function of relative degree r , and define:

$\phi_0 := V$, $\phi_i := L_f \phi_{i-1} + \beta_i(\phi_{i-1})$, $i = 1, \dots, r-1$, with each $\beta_i \in \mathcal{K}$. Then V is a HOCLF if there exists $\beta_r \in \mathcal{K}_\infty$ such that:

$$\inf_{\mathbf{u} \in \mathcal{U}} [L_f^r V + L_g L_f^{r-1} V \mathbf{u} + O(V) + \beta_r(\phi_{r-1})] \leq 0, \quad \forall \mathbf{x} \neq \mathbf{0}_n.$$

Here, L_f, L_g denote Lie derivatives along f, g respectively, and $O(\cdot)$ collects all lower-order Lie and time derivatives up to degree $r-1$.

The HOCBF (1), and the HOCLF (2) can be integrated into a QP convex optimization problem:

$$\begin{aligned} \min_{\mathbf{u}, \delta} \quad & \|\mathbf{u}\|_2^2 + p \delta^2, \\ \text{s.t.} \quad & L_f^r b(\mathbf{x}) + L_g L_f^{r-1} b(\mathbf{x}) \mathbf{u} + O(b(\mathbf{x})) + \alpha_r(\psi_{r-1}(\mathbf{x})) \geq 0, \\ & L_f^r V(\mathbf{x}) + L_g L_f^{r-1} V(\mathbf{x}) \mathbf{u} + O(V(\mathbf{x})) + \beta_r(\phi_{r-1}(\mathbf{x})) \leq \delta. \end{aligned} \quad (4)$$

To keep the QP when several HOCBFs and HOCLFs contradict each other, we add a non-negative slack $\delta \geq 0$ with penalty $p > 0$. This slack relaxes safety constraints when strict enforcement is impossible due to the nominal input or mutual conflicts. Typically, HOCLFs capture performance objectives and HOCBFs safety and other constraints; assigning slack to lower-priority terms lets the controller trade performance for safety.

III. HIGH-ORDER CONTROL BARRIER AND LYAPUNOV FUNCTION FOR ENVIRONMENT-AWARE CONTROL

This section describes how we can design HOCBFs and HOCLFs for environment-aware soft robot control. Most importantly, we can ensure contact force limits, and with this, safety, by deploying HOCBFs while relying on HOCLFs for defining motion behavior and objectives. To enable this, we require access to differentiable algorithms that perform collision detection and resolution between the soft robot body and convex polygonal environment obstacles. We consider a

planar soft robotic arm operating within a two-dimensional workspace $\mathcal{W} \subset \mathbb{R}^2$, populated by n_{obs} known convex obstacles $\mathcal{W}_{\text{obs}} = \{\mathcal{O}_1, \dots, \mathcal{O}_{n_{\text{obs}}}\}$. We also assume the robot is fully actuated, with $m = 3N = \frac{n}{2}$.

A. Collision Detection

To facilitate collision detection and spatial reasoning, the soft robotic arm, though continuously deformable, is approximated as a discrete chain of convex polygonal parts, as illustrated in Fig. 2. Specifically, the robot is segmented into N_{srpoly} convex polygons $\mathcal{R} = (P_1, \dots, P_{N_{\text{srpoly}}})$, each defined by its vertices $\{\mathbf{v}_{i,1}, \dots, \mathbf{v}_{i,k_i}\}$, where k_i is the number of vertices of the i th part. These vertices are computed via the forward kinematics $\text{FK}(\mathbf{q}, s)$ defined earlier. Each polygon is represented as:

$$P_i = \left\{ \mathbf{x} \in \mathbb{R}^2 \mid \mathbf{x} = \sum_{j=1}^{k_i} \alpha_{i,j} \mathbf{v}_{i,j}, \alpha_{i,j} \geq 0, \sum_{j=1}^{k_i} \alpha_{i,j} = 1 \right\}. \quad (5)$$

The entire robot is then $\mathcal{R} = (P_i)_{i=1}^{N_{\text{srpoly}}}$.

This definition now allows us to detect collisions between the N_{srpoly} polygons approximating the robot body and the environment approximated by N_{obs} static convex polygons. In order to do so, we need to know the configuration-dependent distance $h_{i,j}(\mathbf{q}) : \mathbb{R}^{3N} \rightarrow \mathbb{R}$ between the i th soft robot's part P_i and the j th obstacle \mathcal{O}_j . Here, a positive $h_{i,j}(\mathbf{q})$ indicates separation, while a negative value indicates penetration. This distance is provided by a polygon distance metric $d(\cdot, \cdot)$ s.t. $h_{i,j}(\mathbf{q}) = d(P_i(\mathbf{q}), \mathcal{O}_j)$, $i = 1, \dots, N_{\text{srpoly}}$, $j = 1, \dots, n_{\text{obs}}$, between the i th soft robot's part P_i and the j th obstacle \mathcal{O}_j . For notational simplicity, we omit the explicit dependence on the configuration variable \mathbf{q} and write $h_{i,j}$ in the remainder of this section. In principle, any differentiable distance function $d(\cdot, \cdot)$ can be used. In Section IV-B, we present DCSAT, which is a differentiable, computationally efficient, and conservative version of the SAT algorithm.

B. Collision Resolution

The distance metric $h(\mathbf{q})$ now allows us to resolve the collision, which means that we project the collision forces onto the soft robot dynamics and vice versa. In this work, we specifically, without loss of generality, use a linear spring-damper model to capture the collision characteristics. Future work might explore more advanced contact models here. The collision force $F_c \in \mathbb{R}_{\geq 0}$ is given as

$$F_{c_{i,j}}(h_{i,j}, \dot{h}_{i,j}) = \begin{cases} 0, & \text{if } h_{i,j} > 0 \\ -k_c h_{i,j} - c_c \dot{h}_{i,j}, & \text{if } h_{i,j} \leq 0 \end{cases}$$

where k_c the contact stiffness, and c_c the damping coefficient. To recover differentiability, we approximate the collision dynamics using softplus-based smoothing with $\varepsilon \in \mathbb{R}_{\geq 0}$ as

$$F_{c_{i,j}}(h_{i,j}, \dot{h}_{i,j}) = k_c \ln \left(1 + e^{-h_{i,j}/\varepsilon} \right) + c_c \ln \left(1 + e^{-\dot{h}_{i,j}/\varepsilon} \right). \quad (6)$$

This contact force can now easily be applied to both the environment and the soft robot by projecting it along the contact surface vector. Specifically, for a given surface normal $\mathbf{n}_{c_{i,j}} \in \mathbb{R}^2$ pointing from the obstacle to the soft robot body with $\|\mathbf{n}_{c_{i,j}}\|_2 = 1$ and a contact position $\mathbf{p}_{c_{i,j}} \in \mathbb{R}^2$, the

generalized contact torque onto the soft robot can be described by

$$\boldsymbol{\tau}_c = \sum_{i=1}^{N_{\text{srpoly}}} \sum_{j=1}^{N_{\text{obs}}} \mathbf{J}_{c_{i,j}}^\top(\mathbf{q}) F_{c_{i,j}}(h_{i,j}, \dot{h}_{i,j}) \mathbf{n}_c \in \mathbb{R}^{3N}, \quad (7)$$

where $\mathbf{J}_{c_{i,j}} \in \mathbb{R}^{2 \times 3N}$ is the positional Jacobian of the contact point on the surface of the soft robot that is:

$$\dot{\mathbf{p}}_c = \underbrace{\left(\mathbf{J}_{x,y} + \text{diag}(-1, 1) (\mathbf{p}_{c_{i,j}} - \text{FK}_{x,y}(\mathbf{q}, s_{c_{i,j}}))^\top \mathbf{J}_\theta \right)}_{\mathbf{J}_{c_{i,j}}(\mathbf{q})} \dot{\mathbf{q}} \quad (8)$$

with $s_{c_{i,j}} \in (0, L]$ the point on the backbone of the robot closest to the contact position $\mathbf{p}_{c_{i,j}}$ and $\mathbf{J}_\theta(\mathbf{q}, s_{c_{i,j}}) \in \mathbb{R}^{1 \times 3N}$ and $\mathbf{J}_{x,y}(\mathbf{q}, s_{c_{i,j}}) \in \mathbb{R}^{2 \times 3N}$ the orientation and positional rows of the forward kinematics Jacobian $\mathbf{J}(\mathbf{q}, s_{c_{i,j}})$, respectively.

C. Ensuring Safety via High-Order CBFs

Below, we introduce a set of relative-degree-two HOCBFs dedicated to maintaining safety. We start with an HOCBFs that is standard in rigid-robotics applications and ensures complete avoidance of contact between the robot and its environment. While effective, these constraints can sharply curb performance, encourage overly cautious behavior, and block truly collaborative human–robot interactions.

To overcome such drawbacks, we highlight a HOCBF limiting the contact force to $F_{c,\text{max}} \in \mathbb{R}_{>0}$. This approach allows controlled contact between the soft robot and its surroundings while guaranteeing that such contact remains safe [10]. Our formulation is inspired by extensive studies on injury-severity criteria for rigid collaborative robots [2] and by ISO/TS 15066:2016 [15], which specifies body-part-dependent force thresholds as proxies for injury risk.

- 1) **Contact Avoidance HOCBF.** To ensure safety with respect to a forbidden region $\mathcal{A} \subset \mathcal{W}$, we define a HOCBF based on the signed distance between the robot segment P_i and \mathcal{A} . Let $r_{\text{safe}} \geq 0$ be a prescribed safety margin. Using the previously defined smooth distance metric $h_{i,\mathcal{A}}(\mathbf{q}) = d(P_i(\mathbf{q}), \mathcal{A})$, we define

$$b_{i,\mathcal{A}}(\mathbf{x}) = h_{i,\mathcal{A}}(\mathbf{q})^2 - r_{\text{safe}}^2. \quad (9)$$

Then $b_{i,\mathcal{A}}(\mathbf{x}) \geq 0$ guarantees that P_i maintains a distance of at least $r_{\text{safe}} \geq 0$ from \mathcal{A} .

- 2) **Contact Force Limit HOCBF.** For each obstacle $\mathcal{O}_j \in \mathcal{W}_{\text{obs}}$, define

$$b_{i,j}(\mathbf{x}) = F_{c,\text{max},j} - F_{c_{i,j}}(h_{i,j}, \dot{h}_{i,j}), \quad (10)$$

where $F_{c_{i,j}}(h_{i,j}, \dot{h}_{i,j}) \geq 0$ is the contact force, for example, stemming from a linear spring-damper contact model as defined in (6), and $F_{c,\text{max},j}$ is maximum allowable contact force in static settings as, e.g., defined in ISO/TS 15066:2016 [15]. This barrier function ensures that the contact force between the i th soft robot part and the j th obstacle remains below the threshold.

D. Achieving Effective Motion Behavior via Higher-Order Control Lyapunov Functions

With the barrier conditions established, the next step is to define motion objectives that encourage task completion. Although the HOCLF framework affords substantial flexibility in

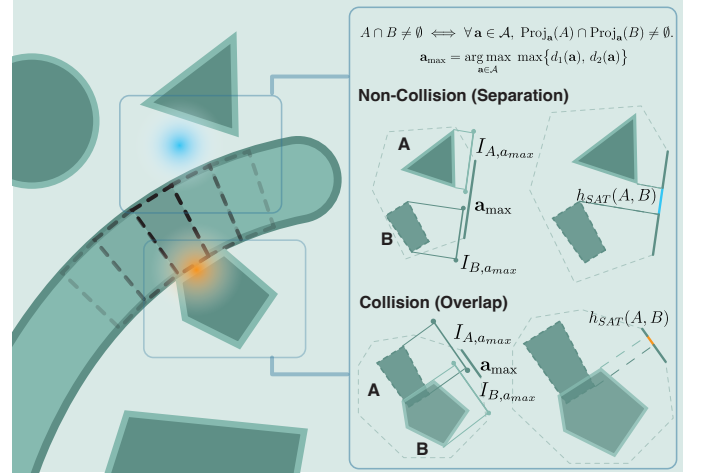


Fig. 2: **Illustration of SAT Polygon Distance Metrics.** Visualization of SAT-based polygon distance metrics used for collision detection between the soft robot body and convex polygonal obstacles. Specifically, we illustrate the definition of the signed distance $h_{\text{SAT}}(A, B)$ between convex polygons A and B . The right panel shows projection intervals $I_{A,a}$ and $I_{B,a}$ along the maximizing axis a_{max} in both separation (top) and overlap (bottom) scenarios.

defining task objectives, this letter concentrates on operational-space regulation; comparable objectives could also be formulated for trajectory tracking, configuration-space regulation and contact-force control.

Operational Space Regulation HOCLF. Let $\mathbf{p}_i(\mathbf{q}) \in \mathbb{R}^2$ be the Cartesian pose of the tip of the i th segment and $\mathbf{p}_{\text{goal}} \in \mathbb{R}^2$ the desired target. Then the HOCLF function

$$V_{\text{tsr},i}(\mathbf{x}) = \|\mathbf{p}_i(\mathbf{q}) - \mathbf{p}_i^d\|_2^2 \quad (11)$$

encourages convergence of the tip of the i th segment toward the target position.

IV. DIFFERENTIABLE POLYGON DISTANCE METRIC

For HOCBF constructions requiring an r th-order continuously differentiable distance metric, the classical SAT [25], despite its efficiency for convex polygons and widespread use over support mapping methods such as the Gilbert-Johnson-Keerthi (GJK) algorithm [26], is fundamentally unsuitable due to its reliance on non-smooth min and max operations. This lack of differentiability prevents its direct application in force- or distance-based high-order control formulations. To address this limitation, prior work has introduced a differentiable variant of SAT, named SSAT, that approximates these non-differentiable operations using multi-level smoothing techniques to achieve C^∞ continuity [21]. However, the SSAT overestimates the polygon separation distance, which can lead to a violation of the safety constraint in downstream applications. Alternatively, methods like DCOL [27] offer differentiable collision detection for convex primitives, but they cannot quantify penetration depth and are therefore unsuitable for collision resolution (e.g., computing the contact force). Similarly, the randomized smoothing approach [28] and the accelerated optimization-based method [29] focus on differentiable collision checking and computational speed-up, respectively, but do not return differentiable signed distance or

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

penetration depth, which limits their use in differentiable contact dynamics. In this letter, we propose a new differentiable variant of SAT, coined DCSAT, which replaces the layered smoothing pipeline of SSAT [21] with a single LogSumExp (LSE) approximation while providing a conservation metric for calculating distance between convex polygons. DCSAT preserves the C^∞ differentiability required for high-order control while significantly simplifying implementation and reducing computational overhead.

A. Definition of SAT

Before defining our distance metrics, we briefly revisit the SAT framework, which forms the foundation of convex polygon–distance computations based on orthogonal projections.

Definition 3 (Separating axis): Let $A, B \subset \mathbb{R}^d$ be convex sets. A unit vector $\mathbf{a} \in \mathbb{R}^d$ is called a *separating axis* for A and B if the projections of the sets onto \mathbf{a} are disjoint; that is, if either

$$\max_{x \in A} \mathbf{a}^\top x < \min_{y \in B} \mathbf{a}^\top y \quad \text{or} \quad \max_{y \in B} \mathbf{a}^\top y < \min_{x \in A} \mathbf{a}^\top x.$$

Theorem 1 (Separating Axis Theorem [25]): Two convex sets A and B in \mathbb{R}^d are disjoint if and only if there exists a separating axis between them.

Lemma 1 (Sufficiency of Edge Normals in \mathbb{R}^2): Let $A, B \subset \mathbb{R}^2$ be convex polygons. Then it suffices to test for separation along the set of directions orthogonal to the edges of A and B . Specifically, let $S(A)$ and $S(B)$ denote the sets of edge directions of A and B , respectively. Define $\mathcal{A} = \{\ell^\perp : \ell \in S(A) \cup S(B)\}$, where ℓ^\perp denotes the unit vector orthogonal to edge ℓ . Then \mathcal{A} is a complete set of candidate separating axes.

Proof. By the Separating Axis Theorem (Theorem 1), if A and B are disjoint, there exists a direction $\mathbf{n} \in \mathbb{S}^1$ such that their projections onto \mathbf{n} do not overlap:

$$\max_{x \in A} \mathbf{n}^\top x < \min_{y \in B} \mathbf{n}^\top y \quad \text{or} \quad \max_{y \in B} \mathbf{n}^\top y < \min_{x \in A} \mathbf{n}^\top x.$$

Since the support function of a convex polygon is piecewise linear and attains its extrema at vertices, any separating direction must be orthogonal to some edge of A or B . Therefore, it suffices to test separation along directions in \mathcal{A} . If no such direction yields separation, A and B must intersect. \square

By combining Theorem 1 and Lemma 1, separation testing—and more specifically, the computation of separating distances—can be reduced to a finite set of one-dimensional projections. These observations underlie the projection-based metrics (SAT, SSAT, DCSAT) used throughout this letter.

Following by the above statements, we could derive the distance between two polygons A and B .

Projections. Let $\mathbf{A}_i, i \in \mathcal{I}_A, \mathbf{B}_j, j \in \mathcal{I}_B$ denote the vertices of convex sets A, B , respectively, where \mathcal{I}_A and \mathcal{I}_B denote the index sets of the vertices of A, B respectively. For each axis $\mathbf{a} \in \mathcal{A}$, the scalar projections of polygon vertices \mathbf{A}_i for $i \in \mathcal{I}_A$, and \mathbf{B}_j for $j \in \mathcal{I}_B$, onto \mathbf{a} are defined as $A_{i,a} = \mathbf{a}^\top \mathbf{A}_i$ and $B_{j,a} = \mathbf{a}^\top \mathbf{B}_j$.

Per-axis distance. The signed separation (positive if separated, negative if overlapping) between the projected intervals along axis \mathbf{a} is given by

$$g_{\text{SAT}}(\mathbf{a}) = \max\{d_1(\mathbf{a}), d_2(\mathbf{a})\}, \quad (12)$$

where

$$d_1(\mathbf{a}) = \min_{j \in \mathcal{I}_B} B_{j,a} - \max_{i \in \mathcal{I}_A} A_{i,a}, \quad d_2(\mathbf{a}) = \min_{i \in \mathcal{I}_A} A_{i,a} - \max_{j \in \mathcal{I}_B} B_{j,a}. \quad (13)$$

Global metric. The overall SAT-based signed distance between polygons A and B is then defined as

$$h_{\text{SAT}}(A, B) = \max_{\mathbf{a} \in \mathcal{A}} g_{\text{SAT}}(\mathbf{a}). \quad (14)$$

This value is positive when the polygons are separated, zero if they touch, and negative when they overlap.

B. Differentiable Conservative SAT (DCSAT)

To enable safe and differentiable distance evaluation for use in HOCBF+HOCLF controllers, and resolve the underestimation of the polygon extents present in SSAT [21], we introduce the DCSAT metric, which provides a conservative estimate of the polygon separation/penetration distance at an increased computational efficiency compared to SSAT. Unlike SSAT, which smooths each intermediate geometric operation and is primarily designed for quadrilateral shapes, DCSAT operates directly on global signed separation distances. It applies a single LSE operation to obtain a C^∞ approximation of the SAT metric that naturally extends to arbitrary convex polygons. Furthermore, DCSAT consistently underestimates the true separation distance, ensuring that safety constraints enforced via CBFs remain valid even under model uncertainty or near-contact conditions. This makes it particularly suitable for collision-aware control of systems with complex geometries.

Definition 4 (Differentiable Conservative SAT (DCSAT)): Let $A, B \subset \mathbb{R}^2$ be convex polygons, and let \mathcal{A} denote the set of separating axes. For each axis $\mathbf{a} \in \mathcal{A}$, define the separation terms $d_1(\mathbf{a})$ and $d_2(\mathbf{a})$ as in (12). Let

$$\mathcal{D} = \{d_m(\mathbf{a}) \mid \mathbf{a} \in \mathcal{A}, m \in \{1, 2\}\}.$$

Then, the DCSAT distance is defined as

$$h_{\text{DCSAT}}(A, B) := \frac{1}{\alpha_{\max}} \log \left(\sum_{d \in \mathcal{D}} e^{\alpha_{\max} d} \right) - \frac{\log(2|\mathcal{A}|)}{\alpha_{\max}}.$$

Lemma 2 (Bounds for DCSAT):

$$-\frac{\log(2|\mathcal{A}|)}{\alpha_{\max}} \leq h_{\text{DCSAT}}(A, B) - h_{\text{SAT}}(A, B) \leq 0.$$

Proof. Let $M = \max \mathcal{D} = h_{\text{SAT}}(A, B)$, and define $e := \frac{1}{\alpha_{\max}} \log \left(\sum_{d \in \mathcal{D}} e^{\alpha_{\max} d} \right)$.

Lower bound: One term in the sum is $e^{\alpha_{\max} M}$, so

$$\begin{aligned} \sum e^{\alpha_{\max} d} &\geq e^{\alpha_{\max} M} \Rightarrow e \geq M \\ &\Rightarrow h_{\text{DCSAT}}(A, B) - M \geq -\frac{\log(2|\mathcal{A}|)}{\alpha_{\max}}. \end{aligned}$$

Upper bound: Since $|\mathcal{D}| = 2|\mathcal{A}|$, we have

$$\begin{aligned} \sum e^{\alpha_{\max} d} &\leq 2|\mathcal{A}| \cdot e^{\alpha_{\max} M} \Rightarrow e \leq M + \frac{\log(2|\mathcal{A}|)}{\alpha_{\max}} \\ &\Rightarrow h_{\text{DCSAT}}(A, B) - M \leq 0. \end{aligned}$$

\square

Theorem 2 (Conservative Approximation via DCSAT): Let

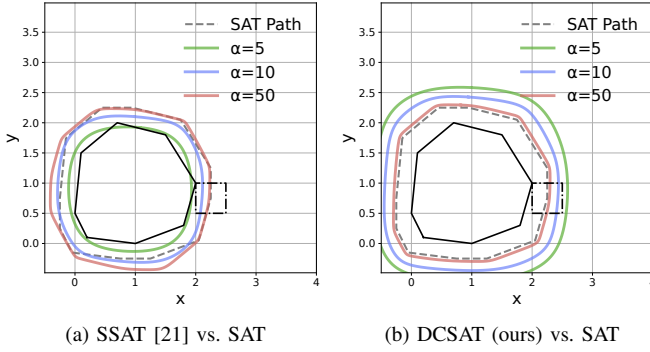


Fig. 3: **Qualitative Benchmarking of DCSAT.** Comparison of zero-level contours between smoothed / differentiable polygon distance metrics and the classical SAT. Specifically, we compare the position of the square polygon at a zero distance with the 8-sided polygon according to the respective distance metric. In both cases, we report the resulting contours for various sharpness parameters α , and the gray dashed curve represents the true zero-level set of the classical, but not differentiable SAT (i.e., the ground-truth), where the centroid of polygon B is in contact with polygon A .

$$e := \frac{1}{\alpha_{\max}} \log \left(\sum_{d \in \mathcal{D}} e^{\alpha_{\max} d} \right), h_{\text{DCSAT}}(A, B) := e - \frac{\log(2|\mathcal{A}|)}{\alpha_{\max}}.$$

Then:

- (a) $h_{\text{DCSAT}}(A, B) \in C^\infty$,
- (b) $h_{\text{DCSAT}}(A, B) \leq h_{\text{SAT}}(A, B)$ for all A, B ,
- (c) If $h_{\text{DCSAT}}(A, B) = 0$, then the true SAT distance satisfies

$$0 \leq h_{\text{SAT}}(A, B) < \frac{\log(2|\mathcal{A}|)}{\alpha_{\max}}. \quad (15)$$

Proof. (a) Since e is composed of exponential and logarithmic operations over a finite sum of smooth terms, it follows that $e \in C^\infty$. Subtracting a constant preserves smoothness, hence $h_{\text{DCSAT}} \in C^\infty$. (b) Follows directly from Lemma 2. (c) If $h_{\text{DCSAT}}(A, B) = 0$, then by definition, $e = \frac{\log(2|\mathcal{A}|)}{\alpha_{\max}}$.

From Lemma 2, it follows that

$$0 \leq h_{\text{SAT}}(A, B) \leq \frac{\log(2|\mathcal{A}|)}{\alpha_{\max}},$$

since $h_{\text{DCSAT}} \leq h_{\text{SAT}}$, with $h_{\text{DCSAT}} = 0$. This confirms the conservative approximation of the real distance: if the smoothed metric DCSAT hits zero, the true SAT distance remains non-negative. \square

The DCSAT procedure is summarized in Algorithm 1.

Algorithm 1 Differentiable Conservative SAT (DCSAT)

Require: Convex polygons $A = \{\mathbf{A}_i\}$, $B = \{\mathbf{B}_j\}$; sharpness $\alpha > 0$

- 1: $\mathcal{A} \leftarrow \{\ell^\perp : \ell \in S(A) \cup S(B)\}$
- 2: **for all** $\mathbf{a} \in \mathcal{A}$ **do**
- 3: $d_1 \leftarrow \min_j(\mathbf{a}^\top \mathbf{B}_j) - \max_i(\mathbf{a}^\top \mathbf{A}_i)$
- 4: $d_2 \leftarrow \min_i(\mathbf{a}^\top \mathbf{A}_i) - \max_j(\mathbf{a}^\top \mathbf{B}_j)$
- 5: Add d_1, d_2 to \mathcal{D}
- 6: **end for**
- 7: $h \leftarrow \frac{1}{\alpha} \log \sum_{d \in \mathcal{D}} e^{\alpha d} - \frac{\log(2|\mathcal{A}|)}{\alpha}$
- 8: **return** h

TABLE I: Average runtime (RT) in milliseconds over 1,000 trials for a batch of 32 four-sided robot segment–obstacle polygon pairs, where each obstacle polygon has N sides. All polygons are aligned to be just in contact. Speedup is measured relative to SSAT. Distance errors are computed as the mean of the deviation from the SAT metric normalized by the obstacle size. Additionally, we report the minimum and maximum observed across all pairs.

N	SSAT RT ↓	DCSAT RT ↓	SAT RT (ours) ↓	Speedup ↑	SSAT Err. [min, max]	DCSAT Err. [min, max]
4	0.046 ms	0.031 ms	0.025 ms	1.51	-12.0% [-37.6%, -0.1%]	-2.8% [-3.5%, -1.9%]
8	0.064 ms	0.041 ms	0.033 ms	1.56	-1.1% [-21.9%, +1.2%]	-2.6% [-3.2%, -1.8%]
16	0.096 ms	0.053 ms	0.039 ms	1.81	+0.2% [-2.3%, +1.2%]	-2.9% [-3.4%, -2.0%]
32	0.129 ms	0.064 ms	0.046 ms	2.00	+0.4% [-2.0%, +1.0%]	-2.8% [-3.3%, -2.1%]
64	0.165 ms	0.076 ms	0.063 ms	2.17	+0.5% [-0.4%, +1.1%]	-2.8% [-3.4%, -2.3%]

TABLE II: Average runtime per segment–polygon pair (ms) with increasing batch size, the polygons have eight sides. Each value is averaged over 1000 trials. Speedup is measured relative to SSAT.

Batch Size	SSAT (ms) ↓	DCSAT (ours) (ms) ↓	Speedup ↑
32	0.0708	0.0463	1.53
64	0.0946	0.0514	1.84
128	0.1330	0.0653	2.04
256	0.2356	0.0985	2.39

C. Benchmarking DCSAT

While the original SSAT formulation [21] was designed specifically for rectangles or axis-aligned quadrilaterals, its underlying principle—approximating projection half-extents using smooth absolute value functions—can be generalized to arbitrary convex shapes. To enable a fair comparison across general geometries, we implemented this natural extension and evaluated it against our proposed DCSAT and the classical SAT [25]. Fig. 3 visualizes the zero-level contours of both smooth metrics relative to the classical SAT boundary. Notably, DCSAT consistently generates a conservative underestimation of the true separation distance, as verified by the consistently negative values in the DCSAT error column of Table I, ensuring that the smoothed constraint remains valid under model error or near-contact conditions. In contrast, SSAT [21] tends to overestimate the true separation at low sharpness levels, as can be seen in the positive values in the SSAT error column of Table I, which can compromise safety-critical guarantees in barrier-based control.

Although both smoothed variants are implemented efficiently in JAX, DCSAT achieves comparable or better computational performance in large batch collision detections compared to SSAT [21] despite its global formulation. As shown in Table I, it maintains fast execution across polygon sizes while preserving full C^∞ smoothness. To evaluate scalability, we benchmark the runtime of our method (DCSAT) against SSAT with increasing batch sizes (Table II). Each batch contains multiple robot segment–obstacle polygon pairs, where each obstacle polygon has $N = 8$. Results show that DCSAT consistently outperforms SSAT in runtime, and the relative speedup grows with the batch size, from 1.53× at batch size 32 to 2.39× at batch size 256. Taken together, these results suggest that DCSAT offers a robust and practical surrogate to classical SAT for use in applications that require differentiability and a conservative approximation of the polygon separation distance.

V. EXPERIMENTS

A. Baseline Methods

a) *Safety Unaware HOCLF*: This baseline captures the typical approach found in current soft-robotic control research [22]: controlling the system without explicitly enforcing safety constraints. Here, we drop the HOCBF terms from the QP and retain only the HOCLF component—specifically, the Operational-Space Regulation HOCLF—within the QP constraints.

b) *Contact Avoidance Artificial Potential Field*: The purpose of this baseline is to represent the scenario that is widespread in the rigid robotics literature (e.g., collision avoidance [2]) - the aim to fully avoid contact - in this case via application of a repulsive Artificial Potential Field (APF) approach [30]. We adopt an integral-free (i.e., $K_i = 0$) variant of operational-space impedance controller from (2) to operate on positions of the soft robot segment tips with $\mathbf{o} = [\mathbf{p}_1^\top, \dots, \mathbf{p}_N^\top]^\top \in \mathbb{R}^{2N}$. Combined with the repulsive artificial potential force \mathbf{f}_{rep} , the resulting control law is:

$$\mathbf{u} = \mathbf{u}_o + \sum_{i=1}^{N_{\text{srpoly}}} \sum_{j=1}^{N_{\text{obs}}} \mathbf{J}_{c_{i,j}}^\top \mathbf{f}_{\text{rep}_{i,j}}, \quad (16)$$

$$\mathbf{f}_{\text{rep}_{i,j}} = \begin{cases} k_{\text{rep}} \frac{h_{i,j}^{-1} - r_{\text{safe}}^{-1}}{h_{i,j}^2} \mathbf{n}_{c_{i,j}}, & \text{if } h_{i,j} \leq r_{\text{safe}} \\ \mathbf{0}_2, & \text{if } h_{i,j} > r_{\text{safe}} \end{cases}$$

Please note that the proportional term of \mathbf{u}_o in (2) corresponds to force stemming from an attractive APF.

c) *Contact Avoidance HOCBF+HOCLF*: The purpose of this baseline is similar to the last one, but instead of relying on an artificial potential field to avoid contact and a classic operational space controller for incorporating the task objective, we rely here on the HOCBF+HOCLF framework by combining a *Contact Avoidance HOCBF* with a *Operational Space Regulation HOCLF*.

d) *Contact Force-Limit HOCBF Filter*: In this baseline, we replace the HOCLF objectives with a nominal operational-space controller whose commands are filtered for safety by solving the QP with *Contact Force Limit HOCBF* constraints. Specifically, we employ $\mathbf{u}_o(\mathbf{x})$ from (2), where the operational-space is defined as $\mathbf{o} = [\mathbf{p}_1^\top, \dots, \mathbf{p}_N^\top]^\top \in \mathbb{R}^{2N}$, representing the positions of the segment tips.

B. Implementation & Simulation Details

We build on the `CBFpy` [31] package that offers an easy-to-use and high-performance implementation of (high-order) CBF+CLF in JAX while leveraging analytical gradients obtained via autodifferentiation. Our simulations consider a planar, two-segment PCS soft robot ($N = 2$) implemented in the `JSRM` package [20]. Each segment is 0.13 m long with a backbone radius of 0.02 m; the elastic modulus is 2 kPa and the shear modulus is 1 kPa. The actuation matrix is set to the identity, $\mathbf{A}(\mathbf{q}) = \mathbb{I}_{3N}$. No state or input bounds are enforced. The closed-loop ODE is integrated with a numerical solver implementing Tsitouras' 5/4 method.

For the contact model, we choose $k_c = 3000 \text{ N/m}$, $c_c = 0$ and set the contact force limit to $F_{c,\text{max},j} = 15 \text{ N}$ and, in contact-avoidance scenarios, impose a minimum safety clearance of $r_{\text{safe}} = 0.005 \text{ m}$, and a repulsive artificial potential

field stiffness of $k_{\text{rep}} = 3000 \text{ N/M}$. The ε value for softplus-based smoothing force is 2×10^{-4} .

C. Navigation Results

We implemented multiple scenarios to validate our framework. Among them, the search & rescue task is the most representative for comparing different baselines and is, therefore, highlighted in this letter (see supplementary materials for additional scenarios). Specifically, we compare the behavior of the *Contact Force-Limit HOCBF+HOCLF* control strategy against the baselines. We present the results in Fig. 4 and Fig. 5 as sequences of stills and time evolution plots, respectively. Indeed, the results demonstrate that the *safety-unaware* controller (1) generates high contact forces that are potentially unsafe, (2) the *contact-avoidance* baselines is not able to complete the task as it exhibits an overly conservative behavior and the soft robot is not able to exploit its embodied intelligence, (3) a model-based operational-space regulator with HOCBF-based safety filter is not able to complete the task as the control strategy is not able to resolve the conflicts between task objective and safety constraints, and (4) the proposed *contact-force limit* exhibits very good task performance while ensuring safety by restricting the maximum contact force.

VI. CONCLUSION

This letter introduces a HOCBFs+HOCLFs framework [18] for the control of soft robots, grounded in a differentiable PCS model [19], [20] and a novel differentiable and conservative polygon distance metric DCSAT. Our method allows (1) soft robots to embrace contact with the environment but ensure safety via contact force-limit HOCBFs evaluated along the entire soft robot body, (2) the flexible and expressive specification of control objectives, such as shape and end-effector regulation or object manipulation, via HOCLFs. The DCSAT metric offers a C^∞ approximation of the classical SAT, yielding conservative signed distances with reduced computational overhead compared to existing baseline methods that overestimate the distance, leading to potentially unsafe behavior. Simulated experiments in navigation scenarios validate the framework's ability to maintain geometric safety and guide soft robots toward task objectives. Current limitations of the work include the focus on planar settings, the simplistic contact model, and the reliance on accurate proprioception and exteroception.

REFERENCES

- [1] A. K. Hall, U. Backonja *et al.*, "Acceptance and perceived usefulness of robots to assist with activities of daily living and healthcare tasks," *Assistive Technology*, 2019.
- [2] S. Haddadin, *Towards safe robots: approaching Asimov's 1st law*. Heidelberg: Springer Berlin, 2013, vol. 90.
- [3] O. Khatib, "A unified approach for motion and force control of robot manipulators: The operational space formulation," *IEEE Journal on Robotics and Automation*, vol. 3, no. 1, pp. 43–53, 1987.
- [4] A. Pupa and C. Secchi, "Efficient iso/ts 15066 compliance through model predictive control," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 17 358–17 364.
- [5] A. D. Ames, X. Xu *et al.*, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2016.

IEEE Robotics and Automation Letters (RA-L) paper, presented at ICRA 2026, Vienna, Austria. Cite as RA-L paper.

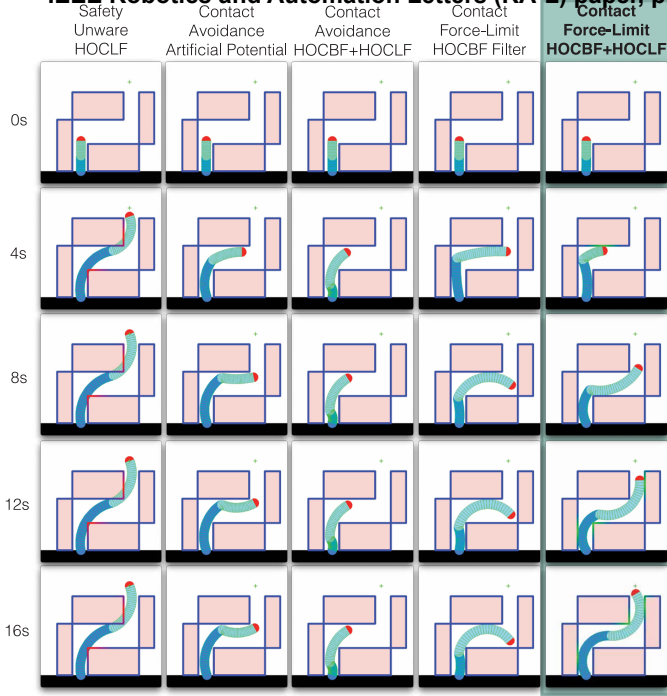


Fig. 4: **Search & Rescue Sequence of Stills.** Sequence of stills for the system evolution in the *search & rescue* scenario for five different control paradigms. Contact interaction states are visualized with color cues: blue indicates no contact, green denotes safe contact, and red highlights contact above the maximum allowable force. The green cross denotes the task-space goal for this task. *Safety Unaware HOCLF*: Only optimizing for the task objective encoded in the HOCLF and disregarding safety constraints. *Contact Avoidance Artificial Potential*: Fully avoiding contact, which is the common paradigm in rigid robotics, using an artificial potential field [30], *Contact Avoidance HOCBF+HOCLF*: Fully avoiding contact using the HOCBF+HOCLF framework. *Contact Force-Limit HOCBF Filter*: Embracing contact with the environment while ensuring that the contact forces remain within safe bounds via an HOCBF-based safety filter applied to a nominal operational-space controller [17], [23], [24]. *Contact Force-Limit HOCBF+HOCLF (proposed)*: Also ensuring that the contact forces remain within safe bounds, but instead of a safety filter, directly solving a QP with HOCBF and HOCLF constraints.

- [6] F. Ferraguti, M. Bertuletti *et al.*, “A control barrier function approach for maximizing performance while fulfilling to iso/ts 15066 regulations,” *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 5921–5928, 2020.
- [7] M. Han, Y. Tian *et al.*, “Reinforcement learning control of constrained dynamic systems with uniformly ultimate boundedness stability guaran-

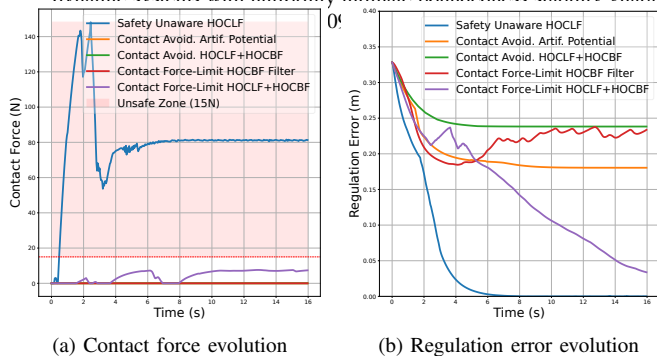


Fig. 5: **Search & Rescue Time Evolution.** The contact force value, regulation error value across the robot for different controllers during a *search & rescue* scenario. Refer to the caption of Fig. 4 for more details about the experiment and the considered controllers.

- [8] D. Rus and M. T. Tolley, “Design, fabrication and control of soft robots,” *Nature*, vol. 521, no. 7553, pp. 467–475, 2015.
- [9] M. Stölzle, N. Pagliarini *et al.*, “Soft yet effective robots via holistic co-design,” *arXiv preprint arXiv:2505.03761*, 2025.
- [10] A. K. Dickson, J. C. P. Garcia *et al.*, “Safe autonomous environmental contact for soft robots using control barrier functions,” *IEEE Robotics and Automation Letters*, vol. 10, no. 11, pp. 11 283–11 290, 2025.
- [11] D. A. Haggerty, M. J. Banks *et al.*, “Control of soft robots with inertial dynamics,” *Science robotics*, vol. 8, no. 81, p. eadd6864, 2023.
- [12] Z. J. Patterson, W. Xiao *et al.*, “Safe control for soft-rigid robots with self-contact using control barrier functions,” in *2024 IEEE 7th International Conference on Soft Robotics (RoboSoft)*. IEEE, 2024, pp. 151–156.
- [13] P. Rao, O. Salzman, and J. Burgner-Kahrs, “Towards contact-aided motion planning for tendon-driven continuum robots,” *IEEE Robotics and Automation Letters*, 2024.
- [14] F. Xu, X. Kang, and H. Wang, “Hybrid visual servoing control of a soft robot with compliant obstacle avoidance,” *IEEE/ASME Transactions on Mechatronics*, vol. 29, no. 6, pp. 4446–4455, 2024.
- [15] I. Standard, “Iso/ts 15066: 2016: Robots and robotic devices—collaborative robots,” *International Organization for Standardization: Geneva, Switzerland*, 2016.
- [16] G. Mengaldo, F. Renda *et al.*, “A concise guide to modelling the physics of embodied intelligence in soft robotics,” *Nature Reviews Physics*, vol. 4, no. 9, pp. 595–610, 2022.
- [17] C. Della Santina, R. K. Katzschmann *et al.*, “Model-based dynamic feedback control of a planar soft robot: trajectory tracking and interaction with the environment,” *The International Journal of Robotics Research*, vol. 39, no. 4, pp. 490–513, 2020.
- [18] W. Xiao, C. A. Belta, and C. G. Cassandras, “High order control lyapunov-barrier functions for temporal logic specifications,” in *2021 American Control Conference (ACC)*. IEEE, 2021, pp. 4886–4891.
- [19] F. Renda, F. Boyer *et al.*, “Discrete cosserat approach for multisection soft manipulator dynamics,” *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1518–1533, 2018.
- [20] M. Stölzle, D. Rus, and C. Della Santina, “An experimental study of model-based control for planar handed shearing auxetics robots,” in *Experimental Robotics*. Cham: Springer Nature Switzerland, 2024, pp. 153–167.
- [21] N. Takasugi, M. Kinoshita *et al.*, “Real-time perceptive motion control using control barrier functions with analytical smoothing for six-wheeled-telescopic-legged robot tachyon 3,” in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2024, pp. 6802–6809.
- [22] C. Della Santina, C. Duriez, and D. Rus, “Model-based control of soft robots: A survey of the state of the art and open challenges,” *IEEE Control Systems Magazine*, vol. 43, no. 3, pp. 30–65, 2023.
- [23] M. Stölzle, S. S. Baberwal *et al.*, “Guiding soft robots with motor-imagery brain signals and impedance control,” in *2024 IEEE 7th International Conference on Soft Robotics (RoboSoft)*. IEEE, 2024, pp. 276–283.
- [24] M. Stölzle, “Safe yet precise soft robots: Incorporating physics into learned models for control,” Dissertation (TU Delft), Mechanical Engineering, Delft University of Technology, 09 2025.
- [25] W. Bittle, “Sat (separating axis theorem),” <https://dyn4j.org/2010/01/sat/>, 2010, accessed: 2025-04-29.
- [26] E. G. Gilbert, D. W. Johnson, and S. S. Keerthi, “A fast procedure for computing the distance between complex objects in three-dimensional space,” *IEEE Journal of Robotics and Automation*, vol. 4, no. 2, pp. 193–203, 1988.
- [27] K. Tracy, T. A. Howell, and Z. Manchester, “Differentiable collision detection for a set of convex primitives,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3663–3670.
- [28] L. Montaut, Q. Le Lidec *et al.*, “Differentiable collision detection: a randomized smoothing approach,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3240–3246.
- [29] L. Montaut, Q. L. Lidec *et al.*, “Collision detection accelerated: An optimization perspective,” in *RSS 2022-Robotics: Science and Systems*, 2022.
- [30] O. Khatib, “The potential field approach and operational space formulation in robot control,” in *Adaptive and Learning Systems: Theory and Applications*. Springer, 1986, pp. 367–377.
- [31] D. Morton and M. Pavone, “Safe, task-consistent manipulation with operational space control barrier functions,” *arXiv preprint arXiv:2503.06736*, 2025, accepted at IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Hangzhou, 2025.