Safe Reinforcement Learning of Dynamic High-Dimensional Robotic Tasks: Navigation, Manipulation, Interaction

Puze Liu¹, Kuo Zhang¹, Davide Tateo¹, Snehal Jauhri¹, Zhiyuan Hu¹, Jan Peters¹⁻⁴ and Georgia Chalvatzaki^{1,3}



Fig. 1: *Top:* Sequence of real-world human-robot interaction. The human casually looks over the phone ignoring the robot that tries to deliver safely a cup of water. While the human reaches for an orange, the robot smoothly avoids collisions, maintaining the glass in an upward position. *Bottom:* The simulated digital twin of the real scene that illustrates the signed distance fields of 0.1m (red points) of the human, the robot, and the table, that are used as constraint models in our Safe Reinforcement Learning algorithm.

Abstract-Safety is a fundamental property for the realworld deployment of robotic platforms. Any control policy should avoid dangerous actions that could harm the environment, humans, or the robot itself. In reinforcement learning (RL), safety is crucial when exploring a new environment to learn a new skill. This paper introduces a new formulation of safe exploration for robotic RL in the tangent space of the constraint manifold that effectively transforms the action space of the RL agent for always respecting safety constraints locally. We show how to apply this approach to a wide range of robotic platforms and how to define safety constraints that represent dynamic articulated objects like humans in the context of robotic RL. Our proposed approach achieves state-of-the-art performance in simulated high-dimensional and dynamic tasks while avoiding collisions with the environment. We show safe real-world deployment of our learned controller on a TIAGo++ robot, achieving remarkable performance in manipulation and human-robot interaction tasks.

I. INTRODUCTION

Safe deployment of general-purpose robotic systems in the real world is an overarching goal of safe learning methods [1]. We envision robots learning complex highdimensional tasks in dynamic, unstructured environments

This work is supported by the China Scholarship Council (No. 201908080039), the DFG Emmy Noether Programme (CH 2676/1-1), the Daimler-Benz Foundation, and the Hessian Competence Center for High Performance Computing – funded by the Hessen State Ministry of Higher Education, Research and the Arts, and partially supported by the German Federal Ministry of Education and Research (BMBF) within the collaborative KIARA project (grant no. 13N16274).

¹ Computer Science Department, Technical University Darmstadt, ² German Research Center for AI (DFKI), Research Department: Systems AI for Robot Learning, ³ Hessian.AI, ⁴ Centre for Cognitive Science. puze@robot-learning.de, {davide.tateo, snehal.jauhri, jan.peters, georgia.chalvatzaki} @tu-darmstadt.de following the paradigm of Deep Reinforcement Learning (RL) [2]. In the standard RL setting, the agent interacts with an environment of unknown dynamics, collecting experience for learning a suitable behavior (policy). The exploration of a Markov Decision Process (MDP) for collecting experience can lead to hazardous situations that can damage the robot or the environment [3]. In this work, we investigate the problem of *safe exploration* in Deep RL for different robotic tasks with high-dimensional state and action spaces, such as robotic manipulation. We also focus on environments that can dynamically change, e.g., due to the dynamic behavior of a human in a Human-Robot Interaction (HRI) setting.

Safe learning for control and RL is a field of increasing interest in light of the different application areas where autonomous systems should operate [4]-[7]. Three main approaches for SafeRL are particularly relevant to our work. First is the Safe set approach, whose objective is to keep the agent in the set of states considered safe [3], [8]-[13]. The Safe set approaches rely either on a predefined safety set and backup policies or, starting from an initial safe policy, expand the safe set during learning. The second line of work formalizes the problem as a Constrained Markov Decision Processes (CMDP) [4]. The objective of CMDP is to learn an optimal policy where the expected (discounted) cumulative constraint violation does not exceed a threshold. To solve the problem under this formulation, many different approaches have been developed based on Lagrangian optimization [14]-[18], or using an additional reward signal to penalize constraint violations [19], [20]. More recent approaches [21], [22] use state augmentation, which transforms the CMDP problem into an unconstrained one by incorporating the constraint information into the policy.

Safe exploration methods address the safety problem at each step when interacting with the environment [23]. In this setting, we define as safe the behavior of an agent that complies with a set of constraints. These constraints are usually predefined, approximated, learned, or inferred from an initial safe policy. The algorithm directly samples a safe action or modifies one to enforce constraint satisfaction. These approaches are based on either Lyapunov functions [7], [24], [25], Hamilton-Jacobi reachability [26], Control Barrier Functions [27], [28], constraint linearization [29], [30] or planning [31], [32]. While some of these methods have already proven to be effective in some controlled tasks [26], [28], the objective of this work is to lay the foundations for extending the applicability of safe exploration techniques to the complexity of real-world constraints for safe robotic RL.

One of the most promising ways to impose safe learning is by Acting on the TAngent Space of the COnstraint Manifold (ATACOM) [30], a method that effectively maps the action space of an RL agent to the tangent space of the constraints, i.e., ensuring local constraint satisfaction. To allow such approach to be effectively used for safe exploration in RL of high-dimensional robotic tasks of different complexity, like navigation, manipulation, HRI, we extend ATACOM towards three key directions. (a) We generalize ATACOM to nonlinear affine control systems, describing a wide range of robotic systems. (b) We show how to define complex constraints that involve the uncontrollable environmental state in ATACOM to allow the agent to explore safely even with dynamic objects in the domain. (c) We provide an indepth discussion about the tangent space of the constraint manifold and show how to increase numerical stability.

Thanks to the generalized formulation of ATACOM, we can transform the action space of RL agents for acting on the constraint manifold of both hand-designed (e.g., joint limits) and learned constraints (e.g., for avoiding contact with humans), resolving various classes of constraints for different tasks (e.g., differential drive constraint), while ensuring the deployability of the learned policy at test time, since ATA-COM is always active. We provide experimental evaluations on different robotic simulated tasks for learning navigation, manipulation, and HRI, and we show the superior performance of the generalized ATACOM agent in terms of safety and task success-rate against representative baselines from the safe exploration literature. Crucially, we demonstrate the real-world applicability of ATACOM in manipulation and HRI tasks with complex learned constraint manifolds (e.g., human manifold in Fig. 1), showcasing the ability to preserve local safety during deployment without significant performance losses due to sim-to-real gaps.

Related Work in Safe Learning for Robotics

Both safety and adaptability concern all robotic tasks, from manipulation [33], [34], & navigation [35], [36], to locomotion [37] & HRI [38]–[41]. These two requirements are the main focus of the vast literature on SafeRL [1]. Many different formulations and solutions have been proposed to face this problem, e.g., uncertainty-aware model-based RL in robot navigation tasks [42]–[44], offline learning for finding unsafe zones and recovery policy for a surgical robot is proposed in [45] and for locomotion in [46]. Many works adopt the idea of adding a safety layer on the RL-policy, that can adapt an unsafe action to a safe one in conjunction with reachability analysis [32], [47]. An optimization layer was used in [48] to train a robot-reaching task in the real world. An ensemble of policies is used in [49], from which the most likely safe policy is transferred to a real robot playing air hockey through episodic interaction. An RL framework that filters suboptimal actions in the domain of HRI is introduced in [50], but was not demonstrated on a real-world task. In this paper, we introduce a generalized framework for safe exploration in RL of robotic tasks that operates on the tangent space of the constraint manifold [30] satisfying both handcrafted and learned constraints [30]. We showcase real-world performance in challenging tasks, like HRI, demonstrating that our framework can enable safe robot learning of various tasks across different application areas. Problem Statement

Safe Reinforcement Learning (SafeRL) applies to problems modeled as a CMDP [4] defined by the tuple $\langle S, A, P, \gamma, R, C \rangle$, where S is the state space, A is the action space, $P : S \times A \times S \rightarrow [0,1]$ is the transition kernel, $\gamma \in (0,1]$ is the discount factor, $R : S \times A \rightarrow \mathbb{R}$ is the reward function, and $C := \{c_i : S \rightarrow \mathbb{R} | i \in \mathbb{N}\}$ is a set of constraint functions. We approach the problem of safety in RL through the glance of *safe exploration* that prevents constraint violations throughout the learning process. Therefore, we formalize our problem as follows,

$$\begin{split} \max_{\pi} \quad & \mathbb{E}_{\tau \sim \pi} \left[\sum_{t}^{T} \gamma^{t} r(\boldsymbol{s}_{t}, \boldsymbol{a}_{t}) \right] \\ \text{s.t.} \quad & c_{i}(\boldsymbol{s}_{t}) \leq 0, \quad i \in \{1, 2, ..., N\}, \ t \in \{0, 1, \cdots, T\} \end{split}$$

where $\tau = \{s_0, a_0, \dots, s_T, a_T\}$ is the trajectory under policy π , and $a_t \sim \pi(\cdot, s_t)$ is the action sampled from policy π . The objective is to maximize the discounted cumulative reward, while satisfying all constraints at each step.

II. NOVEL FORMULATION OF ATACOM FOR MOBILE ROBOTS AND MANIPULATORS

In this section, we provide a novel and more general formulation of the ATACOM method. Our proposed formulation allows applying ATACOM for learning a wide variety of mobile robotics and manipulation tasks. Moreover, we handle some critical aspects of the original ATACOM, improving its numerical stability, learning performance, and safe-space structure. Finally, we show how to model complex realworld collision-avoidance constraints using learned Signed Distance Function (SDF)s, that allow safe learning of complex tasks, as in the domain of HRI. We first introduce the original design of ATACOM. Then, we introduce our reformulation to it that allows its generalization to a broader class of problems.

A. Original ATACOM Formulation

ATACOM [30] is a method for safe exploration in the tangent space of the constraints' manifold. It converts the

constrained RL problem to a typical unconstrained one, while handling both equality and inequality constraints. This method allows us to utilize any model-free RL algorithm, while maintaining the constraints below a designated tolerance. In ATACOM, the state space S is separated into two sets, the controllable state space $Q \subset \mathbb{R}^n$ and the uncontrollable state space $\mathcal{X} \subset \mathbb{R}^m$, i.e., $s = [q^{\mathsf{T}} x^{\mathsf{T}}]^{\mathsf{T}} \in$ \mathbb{R}^{n+m} . ATACOM prescribes that all k constraints are defined on the controllable variable $c(q) \leq 0$, where $c : \mathbb{R}^n \to \mathbb{R}^k$ is differentiable. ATACOM constructs a constraint manifold by introducing the slack variable $\mu \in \mathbb{R}^k$ into the constraint

$$\mathcal{M} = \left\{ (\boldsymbol{q}, \boldsymbol{\mu}) : \bar{\boldsymbol{c}}(\boldsymbol{q}, \boldsymbol{\mu}) = \boldsymbol{c}(\boldsymbol{q}) + \frac{1}{2}\boldsymbol{\mu}^2 = \boldsymbol{0}, \boldsymbol{q} \in \mathcal{Q}, \boldsymbol{\mu} \in \mathbb{R}^k \right\} \quad (1)$$

The tangent-space bases of the constraint manifold are determined by computing the null space $N(q, \mu) \in \mathbb{R}^{(n+k) \times n}$ of the Jacobian matrix $J(q, \mu) = \left[\frac{\partial}{\partial q}\bar{c}(q, \mu)^{\mathsf{T}}, \frac{\partial}{\partial \mu}\bar{c}(q, \mu)^{\mathsf{T}}\right] \in \mathbb{R}^{k \times (n+k)}$. To simplify the notation, we use \bar{c} , N, and J, without explicitly writing the dependency on the input. The velocity of the controllable state can be determined by

$$\begin{bmatrix} \dot{q} \\ \dot{\mu} \end{bmatrix} = N\alpha - K_c J^{\dagger} \bar{c}$$
(2)

with the action $\alpha \sim \pi(\cdot | \boldsymbol{q}, \boldsymbol{x})$ sampled from the policy. The second term on the right-hand side, with the pseudoinverse of the Jacobian J^{\dagger} and the gain K_c , is the error correction term that forces the agent to stay on the manifold, and it is necessary when using time discretization.

B. Safe Exploration with Generalized ATACOM on Nonlinear Affine Control Systems

The original formulation of ATACOM assumes a holonomic system, i.e., it assumes that we can set an arbitrary derivative of each generalized coordinate describing the mechanical system. However, many robotics systems of interest are subject to non-holonomic constraints, i.e., nonintegrable constraints, such as the differential drive and the bicycle model, which prevent imposing arbitrary velocities or accelerations on the system's state variables.

To extend the applicability of ATACOM to a broader class of systems, we reformulate it for nonlinear affine control systems [51]. In this setting, we assume that the system's velocity of the generalized coordinates can be expressed as

$$\dot{\boldsymbol{q}} = \boldsymbol{f}(\boldsymbol{q}) + \boldsymbol{G}(\boldsymbol{q})\boldsymbol{a},\tag{3}$$

with the control action vector a, and two arbitrary (nonlinear) vector functions of the current state variable f(q), G(q).

Following the original ATACOM derivation, we consider the constraint with uncontrollable state, $c(q, x) \leq 0$. The constraint manifold can be defined as $\mathcal{M} = \{(q, x, \mu) : \overline{c}(q, x, \mu) = 0\}$. Assuming that the velocity of the state variables x involved in the constraints are known or estimated, and using the dynamical system in Eq. (3), we write the time derivative of the constraint function $\overline{c}(q, x, \mu)$ as

$$\begin{aligned} \frac{a}{dt}\bar{c}(q,x,\mu) &= J_q \dot{q} + J_x \dot{x} + J_\mu \dot{\mu} \\ &= J_q f(q) + J_q G(q) a + J_x \dot{x} + J_\mu \dot{\mu} \\ &= F(q,x,\dot{x},\mu) + J_G a + J_\mu \dot{\mu}, \end{aligned} \tag{4}$$

with $J_q = \frac{\partial}{\partial q} \bar{c}(q, x, \dot{x}, \mu)$, $J_x = \frac{\partial}{\partial x} \bar{c}(q, x, \mu)$, and $J_\mu = \frac{\partial}{\partial \mu} \bar{c}(q, x, \mu)$ the Jacobian matrices w.r.t. the q, x, and μ variables respectively, $F(q, x, \dot{x}, \mu) = J_q f(q) + J_x \dot{x}$, and $J_G = J_q G(q)$. Again, we drop the explicit dependency of the variables q, μ and x to simplify the notation.

We can now compute the safe action by imposing zero velocity of constraint violation. Setting the right-hand side of Eq. (4) to 0 and solving for a and $\dot{\mu}$ we obtain

$$\begin{bmatrix} a\\ \dot{\mu} \end{bmatrix} = N_{[G,\mu]}\alpha - J^{\dagger}_{[G,\mu]}F(q,x,\dot{x},\mu), \qquad (5)$$

where $J_{[G,\mu]}$ is the concatenation of the J_G and J_{μ} matrices, and $N_{[G,\mu]}$ is the null space of $J_{[G,\mu]}$. As done in ATACOM, it is also possible to add an error correction term $-K_c J^{\dagger} \bar{c}$ to the applied control action computed in (5): this term makes the system more responsive and able to deal with equality constraints. By the special choice of slack variable function discussed in Section II-C, we can ensure $J_{[G,\mu]}$ is full rank and invertible when μ is well-defined. We will introduce how the null space matrix is obtained in Section II-C.

With the extended formulation in Eq. (5), we can compute the control action for a wide variety of common robotics kinematics. As an example, we will consider the differential drive kinematics, for which the mobile robot can only move forward/backward along the heading direction and rotate around its center axis. The differential drive kinematics can be cast as a nonlinear affine control system as follows

$$\boldsymbol{q} = \begin{bmatrix} x \\ y \\ \theta \end{bmatrix} \quad \boldsymbol{f}(\boldsymbol{q}) = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \quad \boldsymbol{G}(\boldsymbol{q}) = \begin{bmatrix} \cos \theta & 0 \\ \sin \theta & 0 \\ 0 & 1 \end{bmatrix} \quad \boldsymbol{a} = \begin{bmatrix} v \\ \omega \end{bmatrix} \quad (6)$$

with the Cartesian coordinates x and y, the current yaw angle θ , the angular velocity ω , and the speed v along heading direction. Given this definitions, it is easy to derive a safe control action using Eq. (5). A similar derivation holds for other kinematics models, e.g. the bicycle kinematics.

Our newly proposed formulation presents a clean and general way to handle systems controlled in velocity/acceleration/jerk or any other arbitrary derivative: this objective can be easily achieved by adding all non-controlled derivatives as state variable q and defining appropriately the f(q) and G(q) functions.

C. Robust Tangent Space Bases

An ATACOM agent explores the tangent space of the constraint manifold. Therefore, obtaining smooth and consistent bases of the tangent space is essential for training the RL policy. The original method [30] constructs the constraint manifold by introducing the slack variables in quadratic form and determines the unique tangent space bases by QR decomposition of the Jacobian matrix and Reduced Row Echlon Form (RREF). However, this approach may derive inconsistent bases and suffers numerical stability issues. In this work, we further investigate the numerical problems and introduce a new type of slack variable and a new way of determining the tangent space bases that keep the consistency. a) Projected Tangent Space Bases: The tangent space bases are usually determined by computing the null space of the Jacobian matrix using QR [52] or SVD [53] decomposition. One desired property is to have continuously varying tangent space bases. However, there is no continuous function that generates the null space [54]. QR/SVD-based methods use the Householder transformation, which applies a sign function $sgn(\cdot)$ during the Tridiagonalization. The null space bases can flip the direction when a diagonal element changes sign. However, [54] proved the continuity of the projection-based method to generate tangent space bases under certain conditions. Following their idea, the projected null space of the Jacobian matrix is determined by

$$N = (I - J^{\mathsf{T}} (J J^{\mathsf{T}})^{-1} J) Z$$
(7)

where J is the Jacobian of the constraint manifold, $Z = [I_n \mathbf{0}_k]^{\mathsf{T}} \in \mathbb{R}^{(n+k) \times n}$ are the augmented bases that combine the normalized action-space bases $I_n \in \mathbb{R}^{n \times n}$ with $\mathbf{0}_n \in \mathbb{R}^{n \times k}$. Eq. (7) projects the action bases onto the tangent space of the constraint manifold. Notice that the first ndimensions of the augmented state correspond to the original action space, and that the project bases are not orthogonal.

b) The SoftCorner Slack Variable: In the original manifold construction, the mapping between the constraint c(q)and μ in Eq. (1) is not unique, causing an ambiguity in the tangent space bases. As an example, Fig. 2 (right) illustrates the constraint manifold $q + \mu^2 = 0$ for the inequality constraint q < 0, where the arrow shows the tangent space basis. The basis of the tangent space at configuration q_0 leads to an opposite direction along the q-axis because the slack variable μ is different, as shown by the points (q_0, μ_1) and (q_0, μ_2) . However, as the slack variable is not observed by the agent, this inconsistency of the basis will cause failures during training. Instead, We can formulate a bijective mapping between the original constraint and the slack variable, such as the exponential form $\exp(\beta \mu)$. Here, β is a positive scalar controlling how strong the original action space shrinks as the constraint function approaches the limit: lowering the β parameter makes the action space more sensitive to the value of the constraint, as the green curve shown in Fig. 2 (left). However, this formulation morphs the



Fig. 2: Comparison of different types of slack variables. **Left:** The constraint manifold defined by different type of slack variables for q < 0. The quadratic slack variable suffers ambiguity issues at point (q_0, μ_1) and (q_0, μ_2) . The tangent space bases (orange arrow) leads to opposite directions along q-axis. **Right:** The morphing of the action space with different slacks. The red dashed lines define the original action limits -1 < a < 1 and black dashed lines are constraints -1 < q < 0. The shaded area defines the projected action space for each state q following Eq. (7)

action space even in states that are far away from constraint violations, as shown in Fig. 2 (right). To avoid this issue, we introduce the *SoftCorner* slack variable parametrization

$$\boldsymbol{c}(\boldsymbol{q},\boldsymbol{x}) + \frac{1}{\beta}\log\left(1 - \exp(\beta\boldsymbol{\mu})\right) = \boldsymbol{0}, \quad (8)$$

where β here has the same interpretation as in the exponential parametrization. In addition, the SoftCorner slack variable maintains the original action space when the states are far from violating the constraint. This slack definition is not sensitive to the metric spaces in which the constraints are defined, such as the combination of the joint space and Cartesian space constraints.

D. Collision Avoidance with Learned SDFs

Collision avoidance constraints are formulated as maintaining a sufficient distance margin between two objects. Typical approach approximate obstacle with primitive shapes, such as spheres. However, this approach is not feasible for complex or dynamic shapes, such as a shelf or a human, as the computation load grows quadratically with the number of primitives. The SDF is a prominent representation for expressing distance w.r.t. a given surface by defining a function that precomputes the distance of an arbitrary query point in the Cartesian space. As SDFs provide a smooth differentiable function, we will discuss how to employ them together with ATACOM.

We rely on Regularized Deep Signed Distance Fields (ReDSDF) [55], which approximates the distance fields of objects with complex shapes or even articulations, such as humans or robots, and provides distances in a regularized form. ReDSDF uses a neural network to approximate the distance of a query point p w.r.t the center of the articulated object, specified by the joint configuration q_o . Unlike DeepSDF [56] that focuses on reconstructing the object's surface as close as possible to the zero level-set, the training in ReDSDF integrates an intuitive distance inductive bias, allowing the learning of distance fields at any scale.

The ReDSDF approximates the closest distance of a query point $p \in \mathbb{R}^3$ in Cartesian space w.r.t the object. To integrate the collision avoidance constraints in ATACOM, we define multiple Point of Interest (PoI) located at relevant positions of the robot. We can compute the Cartesian position of the PoIs p_i given a robot configuration q using forward kinematics $p_i = FK_i(q), i \in [1, 2, ..., N]$. We formulate the collision avoidance constraints as

$$c_i(q, q_o): \delta_i - d(p_i(q), q_o) \le 0, \quad i \in (0, 1, ..., N),$$
 (9)

where δ_i are thresholds assigned for each PoI and $d(\mathbf{p}_i(\mathbf{q}), \mathbf{q}_o)$ is the distance computed by the ReDSDF model. A key advantage of using ReDSDF is that we can compute the gradient w.r.t the robot configuration \mathbf{q} as $\nabla_{\mathbf{q}} \mathbf{c}_i = -\nabla_{\mathbf{p}_i} d \cdot \nabla_{\mathbf{q}} \mathbf{p}_i$, with $\nabla_{\mathbf{p}_i} d$ the gradient of the distance field and $\nabla_{\mathbf{q}} \mathbf{p}_i$ the Jacobian of the PoI w.r.t the robot configuration, which can be computed based on the adjoint matrix [57]. Note that we don't include the velocity for the distance constraint in (9) as the velocity of the obstacles is considered in (5).



Fig. 3: RL environments, from left to right: TableEnv, ShelfEnvSim, ShelfEnvReal, NavEnv

III. EXPERIMENTAL EVALUATION

To evaluate the performance of generalized ATACOM, we designed a series of tasks with different complexity. First, we designed static goal-reaching tasks in which the robot has to avoid collisions with objects of different geometric complexity, e.g., a table and a shelf, as shown in Fig. 3. We then designed two dynamic tasks for navigation and for HRI. In the robot navigation scenario, we assume a differentialdrive mobile robot and test the performance of ATACOM in respecting the control constraints while avoiding collisions with a blindly moving mobile robot. In the HRI task, we consider a shared-workspace scenario in which a human moves dynamically while the robot attempts to deliver a cup appropriately to the desired goal location without colliding with the human or "spilling" the content of the cup. In the following, we provide our empirical evaluation against baselines. All algorithms and environments are implemented within the MushroomRL framework [58]; technical and implementation details can be found on the website https: //irosalab.com/saferobotrl/.

A. Manipulation Tasks

The first two experiments are goal-reaching tasks in the proximity of the obstacles, defined as a table TableEnv or a shelf ShelfEnvSim (Fig. 3). To ensure safety, we define 9 query-PoI distributed along the links of the moving arm for computing the distance constraints with ReDSDF w.r.t. the tables or shelf. We also define distance constraints for avoiding the ground, self-collisions, and joint limits. Both the TableEnv and the ShelfEnvSim contain 34 constraints. The reward is defined as $r(s_t, a_t) = -\rho_d d_{\text{goal}} - \rho_{\angle} d_{\angle} - \rho_d d_{\text{goal}} - \rho_d d_{A_d}$ $\rho_a \| \boldsymbol{a}_t \| + \mathbb{I}_{\boldsymbol{s}}$, with the distance to the goal position d_{goal} and orientation d_{\perp} , the action penalty $||a_t||$, the scaling factor $\rho_d, \rho_{\angle}, \rho_a$, and a task indicator \mathbb{I}_s . The task indicator encourages the robot to stay on top of the table or inside the shelf. For the manipulation task, we apply velocity control $u = \dot{q}$ on the joints with control frequency 30 Hz. Each episode contains 500 steps. Each episode terminates if a collision is detected, with a penalty of $r_{\text{term}} = -1000$.

We compare our approach with three baselines, i.e., vanilla SAC [59], SafeLayer [29], and a hard-coded Linear Attractor (LA). LA-ATACOM applies ATACOM on top of the action computed by LA. Learning the constraint as proposed in [29] is very challenging; thus, we use the constraint functions defined by our approach for a fair comparison. We conducted a hyperparameter search for each task on the learning rates with 5 random seeds, and then, ran 25 seeds with the



Fig. 4: Experimental comparison of ATACOM and baseline methods in different tasks. We report discounted reward, the final distance reached across evaluations, the number of early terminations due to damage events (collision or violation of joint limits), and the average steps per episode (higher means agent living longer through a training episode).

best hyperparameters. We report the final discounted reward and the distance to the target before episode termination. We also compare the total number of episodes with early terminations due to damage events (collisions or violation of joint limits) and average episode steps throughout 2 million training steps, as shown in Fig. 4a and Fig. 4b. Training is not required in LA and LA-ATACOM, and only the final performance is reported. From the boxplots it is clear that learning with ATACOM high-dimensional manipulation tasks achieves better performance, as LA-ATACOM gets stuck in local minima due to joint limits and complex obstacle shapes.

B. Navigation: Differential-Drive Robot

In the NavEnv, we designed a navigation task for the differential-drive TIAGo++ robot (white) that moves in a room while avoiding the Fetch robot (Blue), as shown in Fig. 3. The Fetch robot moves to its randomly assigned target goal, using a hand-crafted policy that ignores TIAGo, and serves as a dynamic obstacle. We aim to train TIAGo to reach its randomly generated goal while avoiding collisions with the Fetch and the walls. The control actions are the

linear and angular velocities of the robot base. We follow a reward structure similar to the one in Section III-A. However, the orientation reward is defined as -sigmoid(30(|d|-0.2)). $|\mathcal{L}_{\text{goal}} - \theta|$, with the heading angle of the robot θ and the yaw angle to the goal \angle_{goal} . As shown in Fig. 4c, the ATACOM-TIAGo can reach the goal with lower error while experiencing significantly fewer collisions. Notably, the performance of SafeLayer in NavEnv is significantly lower compared to the static TableEnv and ShelfEnvSim. Indeed, the SafeLayer approach only corrects the actions when the constraints are violated. To ensure safety in a dynamic task, SafeLayer requires a larger safety threshold, limiting task performance. Instead, ATACOM reduces the feasible action space when approaching the constraint's boundary. In this task, LA-ATACOM outperforms the learned policy as the moving obstacle prevents local minima and the hard-coded policy is a strong and efficient bias to reach a fixed target.

C. Human Robot Interaction

In the last task, we test our approach in a HRI-Sim environment (Fig. 1) with reward function similar to Section III-A. TIAGo should deliver a cup of water vertically to a target point while the human operates in a shared workspace. Like in TableEnv, we define constraints for avoiding collisions w.r.t the table, the robot, the ground, and the human using the pre-trained ReDSDF. To simulate human motion, we record a human using a motion capture system (~18.000 data points). During the recording, the human moves near the table arbitrarily without considering the robot. We convert the human motion to the SMPL representation [60]. To train a robust policy, we randomly initialize the human motion at the beginning of each episode. Note that given the inferior performance of the baselines in the previous tasks, we do not compare against them in this far more challenging task to save computational and energy resources. We evaluate the final policy with 10000 steps, the converged ATACOM policy achieves a final distance to target 0.14 ± 0.02 m and the number of episodes in which a spill happens is 4 ± 3.16 . Throughout the training process (2 million steps), the total number of collisions were 9.2 ± 8.11 . The small number of constraint violation during training confirm the effectiveness of the safe exploration strategy of ATACOM, that ends up learning a safe interaction policy.

D. Real Robot Validation

Given the encouraging results of ATACOM in our simulated results, we transfer our safe policies to the real world^{*}. We create the ShelfEnvReal, in which we place two objects on the shelves of a bookcase (Fig. 3). We use motion capture to perceive the objects' poses. The robot starts from a random configuration and reaches the two targets sequentially. The maximum reaching time for each target is 16.67s. Once the distance between the robot grasping frame and the object is smaller than 10cm for 2s, we count the task as a success. We conduct 25 trials with different combinations of

TABLE I: ShelfEnvReal Experiments

Target	Final Error (m)	Suc. Rate	# Collisions	Time (s)
1	0.038 ± 0.010	90%	0	6.58 ± 1.01
2	0.047 ± 0.024	90%	0	6.47 ± 1.73

object placements and initial arm configuration. As seen in our results in Table I, our policy is 100% safe.

We also conduct real HRI experiments (Fig. 1). We rely on a motion capture system to get an accurate human pose estimate for querying the human-ReDSDF. In our experimental scenario, the human is instructed to ignore the robot (i.e., the subject looks at the phone) while reaching for an object. The robot, in the meantime, delivers a cup of water while avoiding all possible collisions smoothly, exactly as in the simulated task. We conducted 25 trials in which the human behaved differently. The experiment resulted in a **minimum distance** of (0.185 ± 0.021) m a **success rate** of 96% and 0 **collisions**. Note that a task is considered successful when the episode ends with no spills (here, we used granulated sugar instead of water). Our ATACOM-agent trades off a small percentage of task failure to ensure 100% safety.

Limitations While, in principle, it would be possible to use ATACOM to train real robots from scratch, it is still problematic for various reasons. First, the high computational complexity of RL algorithms for solving challenging robotic tasks would require an impractical amount of time to operate the real robot. Moreover, the complexity of resetting the state of a real robot at the end of each episode renders learning from scratch impractical. Additionally, the Gaussian exploration model may be problematic for robotic actuators. Our approach also has some other practical limitations. To ensure safety, we require a high-performance tracking controller and perfect perception e.g., from motion capture. Furthermore, the computational requirements can become an issue if we use multiple ReDSDF networks to compute the constraints. Finally, in this work, we use a simplified model for HRI i.e., a simulated human unaware of the robot's existence. In reality, the human can behave very differently and would probably try to avoid the robot while interacting with it, at least to some extent.

IV. CONCLUSIONS

This paper introduces a novel and more general formulation of Safe Exploration in RL by Acting on the TAngent Space of the COnstraint Manifold (ATACOM), i.e., by transforming the agent's action space by mapping it on the nullspace of the constraints. We extend the original approach to non-linear control affine systems, hence, treating various robotic platforms and a wide variety of tasks, such as manipulation and navigation. Furthermore, we show how to integrate learned constraints with ReDSDF: this allows imposing arbitrary shapes and articulated structures as safety constraints. This flexibility in the constraint definition makes learning safe interactions with humans possible, even when deploying the learned policies in the real world, without sacrificing performance. The proposed method paves the way for possible breakthroughs in the deployment of RL methods in dynamic real-world tasks.

^{*}Videos of the real-world experiments can be found in: https://irosalab.com/saferobotrl/

REFERENCES

- L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, pp. 411–444, 2022.
- [2] J. Ibarz, J. Tan, C. Finn, M. Kalakrishnan, P. Pastor, and S. Levine, "How to train your robot with deep reinforcement learning: lessons we have learned," *The International Journal of Robotics Research*, vol. 40, no. 4-5, pp. 698–721, 2021.
- [3] M. Pecka, K. Zimmermann, and T. Svoboda, "Safe exploration for reinforcement learning in real unstructured environments," in *Proc. of* the Computer Vision Winter Workshop, 2015.
- [4] E. Altman, Constrained Markov decision processes: stochastic modeling. Routledge, 1999.
- [5] T. M. Moldovan and P. Abbeel, "Safe exploration in markov decision processes," in *International Conference on Machine Learning (ICML)*, 2012, pp. 1711–1718.
- [6] H. B. Ammar, R. Tutunov, and E. Eaton, "Safe policy search for lifelong reinforcement learning with sublinear regret," in *International Conference on Machine Learning*. PMLR, 2015.
- [7] F. Berkenkamp, M. Turchetta, A. P. Schoellig, and A. Krause, "Safe Model-based Reinforcement Learning with Stability Guarantees," in *Conference on Neural Information Processing Systems (NIPS)*, 2017.
- [8] A. Hans, D. Schneegaß, A. M. Schäfer, and S. Udluft, "Safe Exploration for Reinforcement Learning," in *European Symposium on Artificial Neural Networks (ESANN)*, 2008.
- [9] J. Garcia and F. Fernandez, "Safe Exploration of State and Action Spaces in Reinforcement Learning," *Journal of Artificial Intelligence Research*, vol. 45, pp. 515–564, 2012.
- [10] A. K. Akametalu, J. F. Fisac, J. H. Gillula, S. Kaynama, M. N. Zeilinger, and C. J. Tomlin, "Reachability-based safe learning with gaussian processes," in 53rd IEEE Conference on Decision and Control. IEEE, 2014, pp. 1424–1431.
- [11] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu, "Safe Reinforcement Learning via Shielding," in AAAI Conference on Artificial Intelligence (AAAI), 2018.
- [12] A. Wachi, Y. Sui, Y. Yue, and M. Ono, "Safe Exploration and Optimization of Constrained MDPs Using Gaussian Processes," in AAAI Conference on Artificial Intelligence (AAAI), 2018.
- [13] T. Koller, F. Berkenkamp, M. Turchetta, and A. Krause, "Learning-Based Model Predictive Control for Safe Exploration," in *IEEE Conference on Decision and Control*, 2018.
- [14] E. Altman, "Constrained Markov Decision Processes with Total Cost Criteria: Lagrangian Approach and Dual Linear Program," *Mathemati*cal methods of operations research, vol. 48, no. 3, pp. 387–417, 1998.
- [15] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained Policy Optimization," in *International Conference on Machine Learning* (*ICML*), 2017.
- [16] A. Stooke, J. Achiam, and P. Abbeel, "Responsive Safety in Reinforcement Learning by PID Lagrangian Methods," in *International Conference on Machine Learning (ICML)*, 2020.
- [17] D. Ding, X. Wei, Z. Yang, Z. Wang, and M. R. Jovanovic, "Provably Efficient Safe Exploration via Primal-Dual Policy Optimization," in *International Conference on Artificial Intelligence and Statistics (AIS-TATS)*, vol. 130, 2021.
- [18] A. I. Cowen-Rivers, D. Palenicek, V. Moens, M. A. Abdullah, A. Sootla, J. Wang, and H. Bou-Ammar, "Samba: Safe model-based & active reinforcement learning," *Machine Learning*, pp. 1–31, 2022.
- [19] Y. Liu, J. Ding, and X. Liu, "Ipo: Interior-point policy optimization under constraints," in AAAI Conference on Artificial Intelligence (AAAI), vol. 34, no. 04, 2020, pp. 4940–4947.
- [20] C. Tessler, D. J. Mankowitz, and S. Mannor, "Reward Constrained Policy Optimization," in *International Conference on Learning Rep*resentations (ICLR), 2019.
- [21] A. Sootla, A. I. Cowen-Rivers, J. Wang, and H. B. Ammar, "Enhancing safe exploration using safety state augmentation," *arXiv preprint* arXiv:2206.02675, 2022.
- [22] A. Sootla, A. I. Cowen-Rivers, T. Jafferjee, Z. Wang, D. Mguni, J. Wang, and H. Bou-Ammar, "Saute rl: Almost surely safe reinforcement learning using state augmentation," *arXiv preprint* arXiv:2202.06558, 2022.
- [23] F. Berkenkamp, "Safe exploration in reinforcement learning: Theory and applications in robotics," Ph.D. dissertation, ETH Zurich, 2019.

- [24] Y. Chow, O. Nachum, E. Duenez-Guzman, and M. Ghavamzadeh, "A Lyapunov-based Approach to Safe Reinforcement Learning," in *Conference on Neural Information Processing Systems (NIPS)*, 2018.
- [25] Y. Chow, O. Nachum, A. Faust, E. Duenez-Guzman, and M. Ghavamzadeh, "Lyapunov-based Safe Policy Optimization for Continuous Control," in *Reinforcement Learning for Real Life* (*RL4RealLife*) Workshop in the 36 th International Conference on Machine Learning, 2019.
- [26] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, "A general safety framework for learning-based control in uncertain robotic systems," *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2737–2752, 2018.
- [27] 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 AAAI Conference on Artificial Intelligence. AAAI Press, 2019, pp. 3387–3395.
- [28] A. Taylor, A. Singletary, Y. Yue, and A. Ames, "Learning for safetycritical control with control barrier functions," in *Learning for Dynamics and Control.* PMLR, 2020, pp. 708–717.
- [29] G. Dalal, K. Dvijotham, M. Vecerik, T. Hester, C. Paduraru, and Y. Tassa, "Safe Exploration in Continuous Action Spaces," *arXiv* preprint arXiv:1801.08757, 2018.
- [30] P. Liu, D. Tateo, H. B. Ammar, and J. Peters, "Robot reinforcement learning on the constraint manifold," in *Conference on Robot Learning*. PMLR, 2022, pp. 1357–1366.
- [31] L. Hewing, K. P. Wabersich, M. Menner, and M. N. Zeilinger, "Learning-Based Model Predictive Control: Toward Safe Learning in Control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, no. 1, pp. 269–296, 2020.
- [32] Y. S. Shao, C. Chen, S. Kousik, and R. Vasudevan, "Reachability-based trajectory safeguard (rts): A safe and fast reinforcement learning safety layer for continuous control," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3663–3670, 2021.
- [33] B. Sukhija, M. Turchetta, D. Lindner, A. Krause, S. Trimpe, and D. Baumann, "Scalable safe exploration for global optimization of dynamical systems," *arXiv preprint arXiv:2201.09562*, 2022.
- [34] D. Martínez, G. Alenya, and C. Torras, "Safe robot execution in model-based reinforcement learning," in 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2015, pp. 6422–6427.
- [35] K. Weerakoon, A. J. Sathyamoorthy, U. Patel, and D. Manocha, "Terp: Reliable planning in uneven outdoor environments using deep reinforcement learning," in 2022 International Conference on Robotics and Automation (ICRA). IEEE, 2022, pp. 9447–9453.
- [36] A. Bajcsy, S. Bansal, E. Bronstein, V. Tolani, and C. J. Tomlin, "An efficient reachability-based framework for provably safe autonomous navigation in unknown environments," in 2019 IEEE 58th Conference on Decision and Control (CDC). IEEE, 2019, pp. 1758–1765.
- [37] A. Marco, D. Baumann, M. Khadiv, P. Hennig, L. Righetti, and S. Trimpe, "Robot learning with crash constraints," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1439–1446, 2021.
- [38] Y. L. Pang, A. Xompero, C. Oh, and A. Cavallaro, "Towards safe human-to-robot handovers of unknown containers," in 2021 30th IEEE International Conference on Robot & Human Interactive Communication (RO-MAN). IEEE, 2021, pp. 51–58.
- [39] R. Pandya and C. Liu, "Safe and efficient exploration of human models during human-robot interaction," arXiv preprint arXiv:2208.01103, 2022.
- [40] J. Thumm and M. Althoff, "Provably safe deep reinforcement learning for robotic manipulation in human environments," in 2022 International Conference on Robotics and Automation (ICRA). IEEE, 2022, pp. 6344–6350.
- [41] S. R. Schepp, J. Thumm, S. B. Liu, and M. Althoff, "Sara: A tool for safe human-robot coexistence and collaboration through reachability analysis," in 2022 International Conference on Robotics and Automation (ICRA). IEEE, 2022, pp. 4312–4317.
- [42] G. Kahn, A. Villaflor, V. Pong, P. Abbeel, and S. Levine, "Uncertaintyaware reinforcement learning for collision avoidance," arXiv preprint arXiv:1702.01182, 2017.
- [43] B. Lütjens, M. Everett, and J. P. How, "Safe reinforcement learning with model uncertainty estimates," in 2019 International Conference on Robotics and Automation (ICRA). IEEE, 2019, pp. 8662–8668.
- [44] G. Chalvatzaki, X. S. Papageorgiou, P. Maragos, and C. S. Tzafestas, "Learn to adapt to human walking: A model-based reinforcement

learning approach for a robotic assistant rollator," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3774–3781, 2019.

- [45] B. Thananjeyan, A. Balakrishna, S. Nair, M. Luo, K. Srinivasan, M. Hwang, J. E. Gonzalez, J. Ibarz, C. Finn, and K. Goldberg, "Recovery rl: Safe reinforcement learning with learned recovery zones," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4915–4922, 2021.
- [46] T.-Y. Yang, T. Zhang, L. Luu, S. Ha, J. Tan, and W. Yu, "Safe reinforcement learning for legged locomotion," arXiv preprint arXiv:2203.02638, 2022.
- [47] J. F. Fisac, N. F. Lugovoy, V. Rubies-Royo, S. Ghosh, and C. J. Tomlin, "Bridging hamilton-jacobi safety analysis and reinforcement learning," in 2019 International Conference on Robotics and Automation (ICRA). IEEE, 2019, pp. 8550–8556.
- [48] T.-H. Pham, G. De Magistris, and R. Tachibana, "Optlayer-practical constrained optimization for deep reinforcement learning in the real world," in 2018 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2018, pp. 6236–6243.
- [49] R. Kaushik, K. Arndt, and V. Kyrki, "Safeapt: Safe simulation-toreal robot learning using diverse policies learned in simulation," *IEEE Robotics and Automation Letters*, 2022.
- [50] M. El-Shamouty, X. Wu, S. Yang, M. Albus, and M. F. Huber, "Towards safe human-robot collaboration using deep reinforcement learning," in 2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2020, pp. 4899–4905.
- [51] A. D. Ames, S. Coogan, M. Egerstedt, G. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in 2019 18th European control conference (ECC). IEEE, 2019, pp. 3420–3431.

- [52] S. S. Kim and M. J. Vanderploeg, "QR Decomposition for State Space Representation of Constrained Mechanical Dynamic Systems," *Journal* of Mechanisms, Transmissions, and Automation in Design, vol. 108, pp. 183–188, 1986.
- [53] R. P. Singh and P. W. Likins, "Singular Value Decomposition for Constrained Dynamical Systems," *Journal of Applied Mechanics, Transactions ASME*, vol. 52, no. 4, pp. 943–948, 1985.
- [54] R. H. Byrd and R. B. Schnabel, "Continuity of the Null Space Basis and Constrained Optimization," *Mathematical Programming*, vol. 35, no. 1, pp. 32–41, 1986.
- [55] P. Liu, K. Zhang, D. Tateo, S. Jauhri, J. Peters, and G. Chalvatzaki, "Regularized deep signed distance fields for reactive motion generation," arXiv preprint arXiv:2203.04739, 2022.
- [56] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, "Deepsdf: Learning continuous signed distance functions for shape representation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 165– 174.
- [57] K. M. Lynch and F. C. Park, *Modern robotics*. Cambridge University Press, 2017.
- [58] C. D'Eramo, D. Tateo, A. Bonarini, M. Restelli, and J. Peters, "Mushroomrl: Simplifying reinforcement learning research," *Journal* of Machine Learning Research, vol. 22, no. 131, pp. 1–5, 2021.
- [59] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel *et al.*, "Soft actor-critic algorithms and applications," *arXiv preprint arXiv:1812.05905*, 2018.
- [60] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, "Smpl: A skinned multi-person linear model," ACM TOG, 2015.