

Provably Safe Deep Reinforcement Learning for Robotic Manipulation in Human Environments

Jakob Thumm and Matthias Althoff

Abstract—Deep reinforcement learning (RL) has shown promising results in the motion planning of manipulators. However, no method guarantees the safety of highly dynamic obstacles, such as humans, in RL-based manipulator control. This lack of formal safety assurances prevents the application of RL for manipulators in real-world human environments. Therefore, we propose a shielding mechanism that ensures ISO-verified human safety while training and deploying RL algorithms on manipulators. We utilize a fast reachability analysis of humans and manipulators to guarantee that the manipulator comes to a complete stop before a human is within its range. Our proposed method guarantees safety and significantly improves the RL performance by preventing episode-ending collisions. We demonstrate the performance of our proposed method in simulation using human motion capture data.

I. INTRODUCTION

In recent years, researchers solved many complex manipulation tasks using deep reinforcement learning (RL), such as operating door handles [1], playing table tennis [2], stacking boxes [3], and controlling multiple robotic arms [4]. Furthermore, recent work [5]–[7] demonstrated that RL-controlled manipulators could successfully maneuver in environments with dynamic obstacles, dodging moving obstacles, and reaching the goal state consistently. Despite these promising results, one open challenge in RL is to formally guarantee the safety of surrounding humans due to their unpredictable movement and many degrees of freedom (DoF). The lack of safety assurances makes applying any vanilla RL method in human working environments irresponsible.

To overcome these restrictions, we propose a *safety shield* concept for RL that guarantees human safety at all times, as shown in Fig. 1. In essence, our safety concept is *speed and separation monitoring* (SSM) according to DIN EN ISO 10218–1 2021, 5.10.3 [8] warranting that the robot is in a completely stopped state before any collision with a human could occur. We need to verify the robot trajectory at a high frequency to achieve this strong safety criterion while still being able to maneuver close to humans. However, for the RL agent it suffices to output actions at a low frequency to make long-term decisions like moving around a human to reach its goal. Our safety shield combines the low-frequency RL agent with our high-frequency formal verification. The RL agent used in this work is a state-of-the-art soft actor-critic (SAC) [9] with hindsight experience

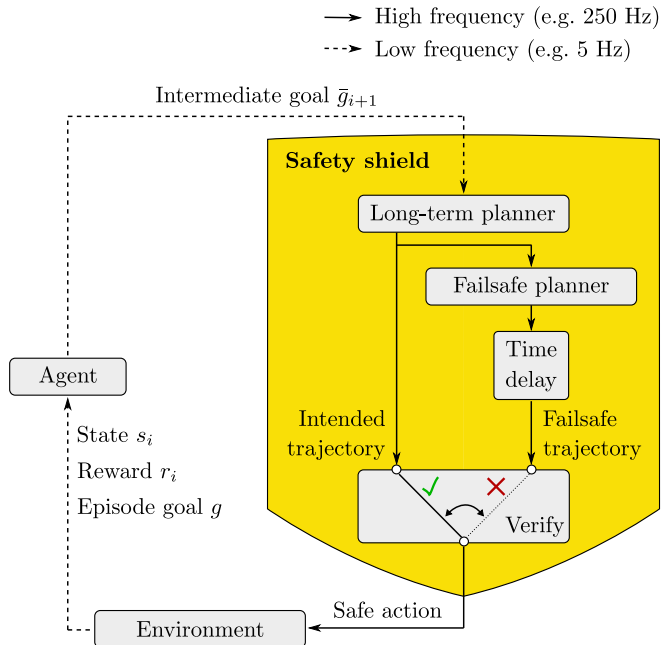


Fig. 1. Structure of the proposed RL agent with safety shield.

replay (HER) [10]. However, our safety shield can be used with any online RL agent. This work presents the first provably safe robot manipulator control based on deep RL in human environments.

A. Related work

Gu et al. [1] were one of the first to show that various complex manipulation tasks can be learned using only off-policy deep RL methods. They presented a method to train multiple real robots in parallel to solve a challenging door-opening task. Shortly later, [10] proposed HER to achieve randomized manipulation goals. The combination of HER with off-policy RL contributed to many remarkable results in pushing, pick and place, and throwing tasks as presented in [10]–[12]. Recently, HER was combined with state-of-the-art RL algorithms like SAC to solve complex manipulation tasks such as controlling multiarm manipulators [4].

There have been many approaches that introduce safety constraints to the exploration process of RL agents. Simple reward shaping has shown to be insufficient to guarantee safety as it is infeasible for universally deciding between fulfilling the safety constraints and rapidly accomplishing the goal [13], [14]. Many recent works use constrained policy optimization techniques [13]–[15] to reduce the number of

safety-critical interactions with the environment. Although some of these approaches show a very low number of safety constraint violations, none can warrant safety at all times.

In many scenarios, especially when interacting with humans, a very low chance of collision is still unacceptable. Therefore, recent works [6], [16]–[18] in the field of provably safe RL (often referred to as shielding, safety layers, sandboxing, or constraining bolts) focus on provable safety guarantees. The core idea of these methods is to ensure that only safe actions are sent to the environment. As a result, unsafe actions are either replaced with safe actions, projected to the safe action space, or entirely prevented by limiting the action space. Many provably safe RL approaches have been proposed for simple environments such as Atari games [19], grid worlds [20], [21], or board games [22]. Most notably, Hunt et al. [17] exhibited how safe end-to-end learning from image data could be achieved in discrete action spaces. Unfortunately, these methods rely on a deterministic environment to assure safety, so they are not directly applicable to a complex task like human-robot collaboration or coexistence. For the use case of RL for manipulators, Pham et al. [6] proposed a differentiable safety layer (OptLayer) that projects any unsafe action to the closest safe action that satisfies the given constraints. Despite promising results, this method comes with two flaws. First, the OptLayer is only formally correct if the constraints of the underlying quadratic program are precisely fulfilled, and the optimizer finds an exact solution in the given time. Second, it is difficult to formalize the safety constraints for unknown and complex human motions. The reachable set-based safety layer of [16] provides formal safety for autonomous driving highway scenarios by masking out all nonsafe actions. Training an RL agent with this safety layer leads to faster convergence while all collisions can be avoided. This method was designed for a discrete action space and is therefore not directly transferable to the manipulator application. The approach closest to ours is the reachability-based trajectory safeguard of [18]. They use reachability analysis to predict if the agent could collide with static obstacles on the current trajectory and replan a safe trajectory if necessary. However, this approach is only designed and tested for static environments.

B. Contributions

We present a novel safety shield¹ that replaces unsafe actions from the agent with provably safe actions from a high-frequency safe trajectory planner (failsafe planner). To the best of our knowledge, our method is the first safety measure for RL that provides provable safety for continuous action spaces in high-dimensional state-spaces and hard-to-predict dynamic environments. Our safety shield can easily be applied to a large variety of manipulation tasks. Compared to all previously mentioned provably safe RL approaches, we sample from the entire set of actions (instead of a safe

¹Our code and models are publicly available at https://github.com/JakobThumm/safe_rl_manipulators.




Fig. 2. Simulation setup with the robot mounted on a desk and the human performing a task. The provided observations are listed on the right hand side.

subset) and check the safety of the current action during execution at high frequency. This allows for quick reactions to highly dynamic human motion. To summarize, our safety shield guarantees that the robot stops before a collision with a human could occur while still allowing the maximal freedom of movement under the given safety constraints.

C. Article structure

Section II summarizes our problem and presents the RL basics and notation used in the article. Next, Section III presents our proposed safety shield and RL agent. We then discuss the experimental setup and results in Section IV and present our conclusion in Section V.

II. PRELIMINARIES

A. Problem statement

In our RL setting, a six DoF modular robot, as described in [23], is mounted on a working table, where a human repeatedly performs a task. In each episode, the agent has to reach a randomized goal joint position, further referred to as *episode goal* g , from a fixed initial position and evade the human. We only consider episode goals that are collision-free with the static environment. This paper aims to provide formal safety guarantees for nearby humans with unknown motion behavior while consistently reaching the goal. Hereby, the RL agent receives the following observations: the current joint position and velocity, episode goal, current Cartesian end-effector position, and the relative Cartesian positions of the human wrists and head in relation to the end-effector. Since the only static obstacle in our scene is the table, we do not add static obstacle information to the observation space. Fig. 2 shows our setup and illustrates the observations.

B. RL

For RL, we use the notation in [9] and consider the Markov decision process defined by the tuple $(\mathcal{S}, \mathcal{A}, p, r)$ with both continuous state space \mathcal{S} and action space \mathcal{A} . For simplicity, we assume the state to be fully observable. The transition function $p : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ denotes the




Fig. 3. Illustration of intended and failsafe trajectories. In time steps $t_0 \dots t_2$, the intermediate goal for the intended trajectories is \bar{g}_i . Before time step t_0 , a failsafe trajectory is calculated for time step t_1 . In between t_0 and t_1 , an intended trajectory for t_1 followed by a failsafe trajectory for t_2 is calculated. The intended trajectory in t_1 is not safe, therefore the robot executes the failsafe trajectory. A new intended trajectory for t_2 followed by a failsafe trajectory for t_3 is calculated during execution of the failsafe trajectory between t_1 and t_2 . At t_2 , the intended trajectory is verified as safe and can be followed. Before the trajectory calculation for t_3 , the agent outputs a new intermediate goal \bar{g}_{i+1} , so that the new intended trajectory is calculated from t_3 to \bar{g}_{i+1} . Since the intended trajectory to \bar{g}_{i+1} is not verified as safe at t_3 , the robot follows the previous failsafe trajectory. A new intended trajectory from t_4 to \bar{g}_{i+1} is calculated followed by a failsafe trajectory. This intended trajectory is safe, so that the robot switches to it.

probability density function of reaching the next state \mathbf{s}_{i+1} when choosing action \mathbf{a}_i in state \mathbf{s}_i . After each transition, the agent receives a reward from the environment according to the reward function $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$. The agent learns a stochastic policy $\pi(\mathbf{a}_i | \mathbf{s}_i)$ for action \mathbf{a}_i given state \mathbf{s}_i .

The optimal policy is learned using the online off-policy SAC algorithm first presented in [9]. Our specific implementation extends the `spinningup` [24] version of SAC with HER and is further described in Section III-B. The goal of SAC is to automatically balance exploration and exploitation by maximizing the accumulated reward and the information content of the policy function. Therefore, the entropy \mathcal{H} (see e.g. [25, Eq. 4.2]) of the policy is used as part of the optimization objective (from [9, Eq. 1])

$$J(\pi) = \sum_{t=0}^T \mathbb{E}_{(\mathbf{s}_i, \mathbf{a}_i) \sim \rho_\pi} [r(\mathbf{s}_i, \mathbf{a}_i) + \alpha \mathcal{H}(\pi(\cdot | \mathbf{s}_i))]. \quad (1)$$

In many robotic applications, one encounters the problem of sparse rewards, where the agent needs to get to a random and hard-to-reach goal. If the agent rarely reaches the goal, almost no transitions with goal encounters are added to the replay buffer. In order to create more of these goal transitions, [10] proposed HER with the idea to sample additional transitions with fictional goals, where the goal g of an episode is added to state \mathbf{s}_i , written as $\mathbf{s}_i || g$. The goal-conditioned transition tuples can be written as $(\mathbf{s}_i || g, \mathbf{a}_i, r_i, \mathbf{s}_{i+1} || g)$. For each transition, additional fictional episode goals \tilde{g}^h , $h = 1 \dots k_{\text{HER}}$ are created from states that were actually reached in the episode of the transition, leading to k_{HER} (in our case $k_{\text{HER}} = 4$) new transitions $(\mathbf{s}_i || \tilde{g}^h, \mathbf{a}_i, \tilde{r}_i^h, \mathbf{s}_{i+1} || \tilde{g}^h)$. By sampling new goals from states that have actually been reached, many successful goal transitions are added. This method enables us to use a sparse reward, where $r_i = 0$ if the goal was reached and $r_i = -1$ otherwise. Since we terminate an episode if a collision occurs, e.g., the human walks into the stationary robot, a

negative reward for collisions is not strictly necessary.

III. METHODOLOGY

A. Safety shield

Our proposed safety shield is based on the provably safe trajectory planner, called *failsafe planner* in [26]. The shield functionality is illustrated in Fig. 1 and further elaborated in this subsection. Please refer to [23], [26]–[28] for a more detailed description of the trajectory planning and verification process.

In our application, the actions of the RL agent are *intermediate goal* joint positions \bar{g}_i . These intermediate goals \bar{g}_i are not to be confused with the episode goal g or the fictional episode goals for HER \tilde{g}^h . In each episode, the agent outputs multiple intermediate goals \bar{g}_i to reach the single episode goal g . The HER goals \tilde{g}^h are only relevant for training the agent. Furthermore, we would like to emphasize that the RL agent outputs the intermediate goals at a lower frequency than our safety shield is operated. Each RL action is executed for a user-defined time interval ΔT , or until all joints are within a small ϵ -range of the intermediate goal. We refer to this as one *RL step* and use the index i . The user-defined time between two safety shield verification steps is Δt with $\Delta t < \Delta T$ and the index k is used.

a) Long-term and failsafe planning: This paragraph describes the trajectory planning in our safety shield and is accompanied by a simplified example in Fig. 3. To reach the intermediate goal \bar{g}_i , we calculate an *intended trajectory* from the current joint state \mathbf{x}_k to \bar{g}_i using a long-term planner. We guarantee that a failsafe trajectory always exists, which brings the robot to an inherently safe state, i.e., a full stop, because a dynamic obstacle can block the robot's path during the execution of this intended trajectory. Formal safety can now be guaranteed by induction, where it is assumed that the robot starts in a safe state (stopped). The first intended trajectory is executed starting from t_0 if no collision with

a human is possible between t_0 and t_1 and if a collision-free failsafe trajectory exists starting from t_1 . During the execution of a trajectory between t_k and t_{k+1} , we calculate a failsafe trajectory starting from t_{k+2} . If this failsafe trajectory cannot be verified as safe, we execute the previously verified failsafe trajectory starting from t_{k+1} . Since each failsafe trajectory ends in a safe state, we can guarantee safety for an infinite time horizon. Similarly, one time step of the intended trajectory between t_k and t_{k+1} is only executed if it is verified as safe; otherwise, the failsafe trajectory is executed. In each time step, a new intended and failsafe trajectory is planned.

In our RL scenario, the agent repeatedly updates the intermediate goal. If there is a new intermediate goal \bar{g}_{i+1} at time step t_k , the intended trajectory is calculated from the state at time t_{k+1} to \bar{g}_{i+1} , and a new failsafe trajectory is calculated from the new intended trajectory starting at time t_{k+2} . If the new intended and failsafe trajectories can be verified as safe, the robot changes to the new intended trajectory. Otherwise, it follows the last verified failsafe trajectory at t_{k+1} .

b) Adaption to manipulator control: In our use case, the 3D trajectory planning can become computationally expensive, especially in regards to the high-frequency requirement. Therefore, we do not recompute the intended trajectory in every time step. Instead, we only compute intended trajectories for every new intermediate goal \bar{g}_{i+1} until an executable one is found and use velocity scaling for the failsafe trajectory as described in [26]. In short, a failsafe trajectory is path-consistent with the intended trajectory and ends in the complete stop of the robot. We use the synchronous *Type IV* online trajectory planner of [29] for the intended and the failsafe planner. Thus, the intended trajectory planning adheres to velocity, acceleration, and jerk limits (v_{\max}^{traj} , a_{\max}^{traj} , and j_{\max}^{traj}), which are gentle for the manipulator joints. To achieve fast-braking, we increase the acceleration and jerk limits for the failsafe planning ($a_{\max}^{\text{failsafe}}$ and $j_{\max}^{\text{failsafe}}$) to the physical limits given by the robot manufacturer.

c) Verification: The trajectory verification is based on collision checking of the reachable sets of all possible human motions and the intended robot trajectory. It is assumed that no human joint moves faster than 2 m/s as defined in DIN EN ISO 13855:2010 [30], and all human joint positions are reliably measurable within a specified error bound, e.g., by using a motion capture system. The safety verification also works with other perception methods like light curtains, as presented in [31], but the robot’s movement would become more conservative. The entire space a human can occupy within the time interval $[t_k, t_{k+b}]$ is defined as the reachable occupancy $\Gamma([t_k, t_{k+b}])$, where b is the number of time steps needed to bring the robot to a complete stop, and is calculated using the task space approach presented in [23], [27].

The human and robot occupancies are modeled with capsules, as shown in Fig. 4, to achieve fast computation times. A capsule comprises a cylinder with half-spheres at both ends and is defined by a line segment $l(p_1, p_2)$




Fig. 4. Example of a human–robot capsule pair in collision.

with endpoints p_1 and p_2 and a radius r . The robot occupancy is described as a set of capsules enclosing each robot link’s movement between time t_k and t_{k+b} . Hereby, we follow the approach described in [26] to guarantee over approximation of the link occupancies. This approach assumes that the controller follows the desired trajectory exactly. If additionally deviation occurring due to low-level control should be considered, the approach presented in [31] can be used. A robot motion is verified as safe if no robot capsule intersects with any human capsule for all times before reaching the resting position.

B. Safe RL

Our policy network has a tanh activation function for the output layer and thus outputs actions $a_i = [a_i^0, a_i^1, \dots, a_i^N]$, $a_i^n \in [-1; 1] \forall n \in 1 \dots N$ with N being the number of joints. We convert these actions to intermediate goal joint positions with $q_{\bar{g}, i+1}^n = q_i^n + a_i^n \Delta q_{\max}$, where q_i^n is the current position of joint n and Δq_{\max} is the maximum joint position difference per RL step. In recent literature, it is more common to use the relative change in 3D position of the robot end-effector as action space and calculate q_{i+1}^n using inverse kinematics (e.g. in [3]) since it leads to a lower dimensional action space. However, we use joint positions as actions to have full control over each joint. Nevertheless, both approaches are equally compatible with our proposed safety shield. The choice of Δq_{\max} depends on the maximum velocity of the robot v_{\max}^{traj} and the execution time of each RL step ΔT . Choosing $\Delta q_{\max} \geq v_{\max}^{\text{traj}} \Delta T$ allows the agent to output a new action before the current action is finished. This is desirable because each trajectory ends in a stopped state and we do not want the robot to stop after each action. If the robot collides with the static environment in its intermediate goal, a new action is generated randomly with a uniform distribution over the entire action space until a collision-free intermediate goal is determined. An alternative to this method is to project the unsafe action to the safe action space with respect to the static environment. However, this is nontrivial in the case of a high-dimensional manipulator.

In our training procedure, which is described in Algorithm 1, an episode is *done* if the episode goal is reached, the robot is in collision, or by timeout when the maximum number of RL steps $T_{\max, \text{episode}}$ is reached. We consider the episode goal as reached, if $|q_{i+1}^n - q_g^n| < \epsilon_g, \forall n \in 1 \dots N$, where q_{i+1}^n is the position of joint n after the environment step, q_g^n is the episode goal

Algorithm 1: Training procedure

```

1 Initialize:
   • soft actor-critic agent A
   • environment E
   • replay buffer  $R$  and local buffer  $L$ 
   •  $T_{\text{total}} \leftarrow 0$ ,  $T_{\text{last\_update}} \leftarrow 0$ 
   • randomly initialize weights of  $\pi_b$ 
2 for  $j \leftarrow 0$  to  $n_{\text{epochs}}$  do
3   for  $l \leftarrow 0$  to  $n_{\text{episodes\_per\_epoch}}$  do
4     reset E and clear  $L$ ;
5     for  $i \leftarrow 0$  to  $t_{\text{max\_episode}}$  do
6       if  $T_{\text{total}} \geq k_{\text{start\_steps}}$  then
7          $a_i \leftarrow \pi_b(s_i || g)$ ;
8       else
9          $a_i \leftarrow$  random action;
10       $(s_{i+1}, r_i, \text{done}) \leftarrow \mathbf{E}.\text{step}(a_i)$ ;
11       $L[i] \leftarrow (s_i || g, a_i, r_i, s_{i+1} || g)$ ;
12       $T_{\text{total}} \leftarrow T_{\text{total}} + 1$ ;
13      if done then break;
14       $l_{\text{episode}} \leftarrow i$ ;
15      for  $i \leftarrow 0$  to  $l_{\text{episode}} - 1$  do
16         $(s_{id} || g, \dots) \leftarrow L[id]$ ;
17         $R.\text{store}(s_i || g, a_i, r_i, s_{i+1} || g)$ ;
18        remaining_ids
19         $\leftarrow (i + 1) \dots (l_{\text{episode}} - 1)$ ;
20        randomly select  $k_{\text{HER}}$  random_ids from
21        remaining_ids;
22        for id in random_ids do
23           $(s_{id} || g, \dots) \leftarrow L[id]$ ;
24           $\tilde{g} \leftarrow s_{id}$ ;
25           $\tilde{r}_i \leftarrow \mathbf{E}.\text{ComputeReward}(s_{i+1} || \tilde{g})$ ;
26           $R.\text{store}(s_i || \tilde{g}, a_i, \tilde{r}_i, s_{i+1} || \tilde{g})$ ;
27      if  $T_{\text{total}} \geq T_{\text{update\_after}}$  and
28       $(T_{\text{total}} - T_{\text{last\_update}}) \geq T_{\text{update\_every}}$  then
29        for  $m \leftarrow 0$  to  $T_{\text{total}} - T_{\text{last\_update}}$  do
30           $\mathbf{A}.\text{Update}(R.\text{sample\_batch}());$ 
31         $T_{\text{last\_update}} \leftarrow T_{\text{total}}$ ;

```

position of joint n , and ϵ_g is a small user-defined tolerance for reaching the goal. At the beginning of each episode, the environment is reset (Algorithm 1 Line 4) by moving the robot back to its initial position and resetting the human animation. The action of a single RL step is executed in Line 10 by the function $\mathbf{E}.\text{step}(a_i)$. To achieve a high exploration rate in the beginning of the training, we choose random actions instead of the policy for the first $k_{\text{start_steps}}$ RL steps (Line 6). After each episode, HER samples are added to the replay buffer as described in Line 15 to Line 24. The agent networks are first updated after $T_{\text{update_after}}$ RL steps to improve exploration at the start of training and then every $T_{\text{update_every}}$ RL steps from Line 25 to Line 27. All relevant hyperparameters for RL training and operation are given in Table I.

TABLE I
HYPERPARAMETERS

Parameter	Value	Parameter	Value
Δt	4 ms	Hidden layers (all networks)	3
ΔT	200 ms	Hidden units per layer	64
$v_{\text{traj}}^{\text{max}}$	2 rad/s	Minibatch size	128
$a_{\text{traj}}^{\text{max}}$	2 rad/s ²	Replay buffer size	10 ⁶
$J_{\text{traj}}^{\text{max}}$	15 rad/s ³	Discount factor γ	0.99
$a_{\text{failsafe}}^{\text{max}}$	10 rad/s ²	Entropy tradeoff α	0.2
$J_{\text{failsafe}}^{\text{max}}$	400 rad/s ³	$k_{\text{start_steps}}$	5000
Δq_{max}	0.4 rad	$T_{\text{update_after}}$	1000
		$T_{\text{update_every}}$	200
		n_{epochs}	200
		$n_{\text{episodes_per_epoch}}$	30

C. Software structure

Our software framework is based on Gazebo and ROS Noetic. The robot simulation uses the Open Dynamics Engine (ODE) solver with a step size of 1 ms. We choose the ODE QuickStep method with 10 solver iterations. The safety shield is implemented in C++² to achieve fast calculation times with an average of 0.5 ms per verification step³. With the safety shield operating at 250 Hz, our simulation runs roughly five times faster than real-time.

IV. RESULTS

A. Experimental setup

We conduct two different experiments. In the first *randomized-goal experiment*, the episode goal is randomly and uniformly sampled over the entire joint space, and the start state is fixed with all joints at their zero position. Notably, the human may block the episode goal. For the second *human-evasion experiment*, we manually define the start and episode goal position so that the robot collides with the human if it takes the shortest path. The agent must learn to evade the human by taking a longer path or waiting until the human moves away from the table. In this scenario, the episode goal is only slightly randomized so that the human never blocks the goal state. The two experiments are also shown in our accompanying video⁴.

For the simulation, the human motion is taken from real-world motion capture data provided by CMU⁵. Currently, we are using animation 62_01, containing a human walking to the working table, where they perform a wrenching action. This motion pattern is diverse and complex and can effectively validate our concept. To prevent overfitting to the motion data, we randomize the x- and y-position and start time of the human animation in each episode uniformly in the range of -0.2 m to 0.2 m and 0 s to 1 s.

In all experiments, the maximum episode length is set to 100 RL steps per episode to give the robot adequate

²Part of our code is based on the open source reachability analysis tool SaRA (available at <https://github.com/Sven-Schepp/SaRA>) [32].

³Run on a ThinkPad P15 Gen 1 with an Intel(R) Core(TM) i7-10750H CPU and 32 GB DDR4 RAM.

⁴Available at <https://youtu.be/Lzrs2HQUIOc>

⁵Data publicly available at <http://mocap.cs.cmu.edu/>




Fig. 5. Evaluation of the randomized-goal (upper row) and human-evasion (lower row) experiment. Figures (a) and (d) compare the rate of safety-critical collisions per episode of the baseline and the safe agent. The success rate of both experiments is displayed in (b) and (e). The reason for episode endings using the safe agent are displayed in (c) and (f). Hereby, a safe collision occurs when the human touches the stopped robot.

time to move around the human to any goal. Since an episode may end earlier due to collisions and successful goal achievements, the number of RL steps per episode varies. Each training consists of 200 epochs with 30 episodes per epoch, leading to a maximum of 600 000 RL steps.

B. Randomized-goal experiment

The following discussion refers to the agent trained excluding and including a safety shield as *baseline* and *safe agent*, respectively. We trained the baseline agent with the same hyperparameters as the safe agent to achieve good comparability. In order to improve the performance of the baseline agent, we also tried to operate it at higher frequencies (50 Hz and 250 Hz) and provide it with a high negative reward for collisions. Unfortunately, none of the measures lead to successful training improvements, so we omit these results. First, we want to highlight the results for completely randomized goal positions. The collision rate comparison in Fig. 5a with and without a safety shield clearly shows the advantage of our proposed method. The baseline agent causes a safety-critical collision in approximately 20% of all episodes. Therefore, the baseline is too dangerous to be deployed in a real-world environment, especially at the beginning of the training. In contrast, the safety shield successfully prevents all safety-critical collisions. Both agents show a similar success rate in Fig. 5b. Interestingly, the baseline agent learns slightly faster than the safe agent despite its high collision rate. At the end of the training, both agents have approximately the same success rate of 65%. Fig. 5c evaluates how the remaining 35% of episodes end for the safe agent. Approximately 10% of episodes end in a safe collision caused by the static human animation walking into the stationary robot. The remaining 25% of episodes ended by exceeding $T_{\max, \text{episode}}$. A manual investigation of the unsuccessful cases showed that most of the goals were

not reachable because humans blocked the path to the goal.

C. Human-evasion experiment

The human-evasion experiment highlights another advantage of our proposed method. Since the human blocks the optimal path, the baseline agent collides with the human in almost every episode in the first half of the training as shown in Fig. 5d. In the second half of training, the baseline agent learned to avoid collisions with the human more frequently, but still collides in approximately 25% of all episodes. Again, our safety shield prevents all safety-critical collisions. Fig. 5e shows that the baseline agent cannot reach the goal because it does not learn how to evade the human. However, the safe agent learns to move around the human to reach the goal effectively.

V. CONCLUSIONS

Our results clearly show that our proposed safety shield is an effective method to avoid collisions with humans in the working environment. Contrary to existing RL methods for safe manipulator control, our safety shield provides formal safety guarantees in highly dynamic and prior unknown human environments. The human-evasion experiment also demonstrated that our safety shield could significantly impact the training success in scenarios with a high likelihood of a collision. We believe that our method enables the high-level control of manipulators with RL agents in real-world human environments while providing the necessary safety guarantees to any human operator. Our subsequent goals are to refine the RL agent further and train it on a more complex set of real-world human motions. Finally, we plan to test our safe RL agent on actual robots and perform tasks in a human working environment.

ACKNOWLEDGMENT

The authors gratefully acknowledge financial support by the Central Innovation Programme of the German Federal Government under grant ZF4086004LP7 and the Horizon 2020 EU Framework Project CONCERT under grant 101016007.

REFERENCES

- [1] S. Gu, E. Holly, T. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2017, pp. 3389–3396.
- [2] W. Gao, L. Graesser, K. Choromanski, X. Song, N. Lazic, P. Sanketi, V. Sindhwani, and N. Jaitly, "Robotic table tennis with model-free reinforcement learning," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2020, pp. 5556–5563.
- [3] R. Li, A. Jabri, T. Darrell, and P. Agrawal, "Towards practical multi-object manipulation using relational reinforcement learning," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2020, pp. 4051–4058.
- [4] E. Prianto, M. Kim, J.-H. Park, J.-H. Bae, and J.-S. Kim, "Path planning for multi-arm manipulators using deep reinforcement learning: Soft actor-critic with hindsight experience replay," *Sensors*, vol. 20, no. 20, pp. 1–22, 2020.
- [5] M. El-Shamouty, X. Wu, S. Yang, M. Albus, and M. F. Huber, "Towards safe human-robot collaboration using deep reinforcement learning," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2020, pp. 4899–4905.
- [6] T.-H. Pham, G. De Magistris, and R. Tachibana, "Optlayer - practical constrained optimization for deep reinforcement learning in the real world," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2018, pp. 6236–6243.
- [7] B. Sangiovanni, G. P. Incremona, M. Piastra, and A. Ferrara, "Self-configuring robot path planning with obstacle avoidance via deep reinforcement learning," *IEEE Control Systems Letters*, vol. 5, no. 2, pp. 397–402, 2021.
- [8] "Robotics - safety requirements - part 1: Industrial robots," International Organization for Standardization, Geneva, CH, Standard DIN EN ISO 10218-1:2021-09 DC, Sep. 2021.
- [9] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *Proc. of the 35th Int. Conf. on Machine Learning (ICML)*, Jul. 2018, pp. 1861–1870.
- [10] M. Andrychowicz, F. Wolski, A. Ray, J. Schneider, R. Fong, P. Welinder, B. McGrew, J. Tobin, P. Abbeel, and W. Zaremba, "Hindsight experience replay," in *Proc. of the 31st Int. Conf. on Neural Information Processing Systems (NeurIPS)*, 2017, pp. 5055–5065.
- [11] Z. Ren, K. Dong, Y. Zhou, Q. Liu, and J. Peng, "Exploration via hindsight goal generation," in *Proc. of the 33rd Int. Conf. on Neural Information Processing Systems (NeurIPS)*, 2019.
- [12] Z. Bing, M. Brucker, F. O. Morin, R. Li, X. Su, K. Huang, and A. Knoll, "Complex robotic manipulation via graph-based hindsight goal generation," *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–14, 2021, early access.
- [13] J. Achiam, D. Held, A. Tamar, and P. Abbeel, "Constrained policy optimization," in *Proc. of the 34th Int. Conf. on Machine Learning (ICML)*, 2017, pp. 22–31.
- [14] A. Ray, J. Achiam, and D. Amodei, "Benchmarking safe exploration in deep reinforcement learning," *arXiv preprint arXiv:1910.01708*, vol. 7, 2019.
- [15] 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.
- [16] H. Krasowski, X. Wang, and M. Althoff, "Safe reinforcement learning for autonomous lane changing using set-based prediction," in *Proc. of the IEEE 23rd Int. Conf. on Intelligent Transportation Systems (ITSC)*, 2020, pp. 1–7.
- [17] N. Hunt, N. Fulton, S. Magliacane, T. N. Hoang, S. Das, and A. Solar-Lezama, "Verifiably safe exploration for end-to-end reinforcement learning," in *Proc. of the 24th Int. Conf. on Hybrid Systems: Computation and Control (HSCC)*, 2021, article no. 14.
- [18] 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.
- [19] N. Jansen, B. Könighofer, S. Junges, A. Serban, and R. Bloem, "Safe reinforcement learning using probabilistic shields," in *Leibniz Int. Proc. in Informatics (LIPIcs)*, vol. 171, 2020, pp. 31–316.
- [20] M. Alshiekh, R. Bloem, R. Ehlers, B. Könighofer, S. Niekum, and U. Topcu, "Safe reinforcement learning via shielding," in *Proc. of the 32nd AAAI Conf. on Artificial Intelligence (AAAI)*, 2018.
- [21] B. Könighofer, J. Rudolf, A. Palmisano, M. Tappler, and R. Bloem, "Online shielding for stochastic systems," in *NASA Formal Methods*, 2021, pp. 231–248.
- [22] G. De Giacomo, L. Iocchi, M. Favorito, and F. Patrizi, "Foundations for restraining bolts: Reinforcement learning with LTLf/LDLf restraining specifications," *Proc. of the 29th Int. Conf. on Automated Planning and Scheduling (ICAPS)*, pp. 128–136, 2021.
- [23] M. Althoff, A. Giusti, S. B. Liu, and A. Pereira, "Effortless creation of safe robots from modules through self-programming and self-verification," *Science Robotics*, vol. 4, no. 31, pp. 1–14, 2019.
- [24] J. Achiam, "Spinning up in deep reinforcement learning," OpenAI, 2018.
- [25] B. D. Ziebart, "Modeling purposeful adaptive behavior with the principle of maximum causal entropy," Ph.D. dissertation, Carnegie Mellon University, 2010.
- [26] D. Beckert, A. Pereira, and M. Althoff, "Online verification of multiple safety criteria for a robot trajectory," in *Proc. of the IEEE 56th Conf. on Decision and Control (CDC)*, 2017, pp. 6454–6461.
- [27] A. Pereira and M. Althoff, "Calculating human reachable occupancy for guaranteed collision-free planning," in *Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2017, pp. 4473–4480.
- [28] —, "Overapproximative human arm occupancy prediction for collision avoidance," *IEEE Transactions on Automation Science and Engineering*, vol. 15, no. 2, pp. 818–831, 2018.
- [29] T. Kröger and F. M. Wahl, "Online trajectory generation: Basic concepts for instantaneous reactions to unforeseen events," *IEEE Transactions on Robotics*, vol. 26, no. 1, pp. 94–111, 2010.
- [30] "Safety of machinery - positioning of safeguards with respect to the approach speeds of parts of the human body," International Organization for Standardization, Geneva, CH, Standard DIN EN ISO 13855:2010-10 ST N, 2010.
- [31] A. Pereira and M. Althoff, "Safety control of robots under computed torque control using reachable sets," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2015, pp. 331–338.
- [32] S. Schepp, J. Thumm, S. B. Liu, and M. Althoff, "Sara: A tool for safe human-robot coexistence and collaboration through reachability analysis," in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2022, early access.