

Online Search-based Collision-inclusive Motion Planning and Control for Impact-resilient Mobile Robots

Zhouyu Lu, Zhichao Liu, Merrick Campbell, and Konstantinos Karydis

Abstract—This paper focuses on the emerging paradigm shift of collision-inclusive motion planning and control for impact-resilient mobile robots, and develops a unified hierarchical framework for navigation in unknown and partially-observable cluttered spaces. At the lower-level, we develop a deformation recovery control and trajectory replanning strategy that handles collisions that may occur at run-time, locally. The low-level system actively detects collisions (via embedded Hall effect sensors on a mobile robot built in-house), enables the robot to recover from them, and locally adjusts the post-impact trajectory. Then, at the higher-level, we propose a search-based planning algorithm to determine how to best utilize potential collisions to improve certain metrics, such as control energy and computational time. Our method builds upon A* with jump points. We generate a novel heuristic function, and a collision checking and adjustment technique, thus making the A* algorithm converge faster to reach the goal by exploiting and utilizing possible collisions. The overall hierarchical framework generated by combining the global A* algorithm and the local deformation recovery and replanning strategy, as well as individual components of this framework, are tested extensively both in simulation and experimentally. An ablation study draws links to related state-of-the-art search-based collision-avoidance planners (for the overall framework), as well as search-based collision-avoidance and sampling-based collision-inclusive global planners (for the higher level). Results demonstrate our method’s efficacy for collision-inclusive motion planning and control in unknown environments with isolated obstacles for a class of impact-resilient robots operating in 2D.

I. INTRODUCTION

There has been an emerging paradigm shift in mobile robot motion planning and autonomous navigation whereby collisions with obstacles are not by default avoided but instead exploited to improve certain robot planning, control and navigation metrics [1]–[5]. Such *collision-inclusive* planning and control strategies capitalize on results demonstrating how some forms of collisions can in fact be useful in terms of sensing, localization, control, and agility [6]–[14]. Besides the benefits of embracing collisions, robot deployment in realistic (that is, dynamic, cluttered, and irregularly-shaped) environments may, at cases, make collision avoidance hard to achieve [15], [16]. For example, detecting all obstacles in the environment can be a challenge, especially when there

exist translucent and/or transparent obstacles, such as glass walls, or reflective surfaces [13]. At the same time, using a conservative local planner may fail finding a feasible path to the goal even if one exists [17]. Collision-inclusive motion planners can help address the aforementioned challenges.

Although research on collision-inclusive motion planning has already begun receiving attention, existing methods can be limited in their ways to apply in practical cases. On one hand, methods that evaluate the effect of collision within motion planning [3], [4] do not apply to online problems. On the other hand, existing online collision-inclusive planning methods [9], [11], [18] cannot decide how to use collisions optimally, which could help guide the robot to the goal. Our previous online planning method [2] can evaluate possible collisions in unknown space which lies outside the field-of-view (FoV) of the robot, but does not consider how to employ collisions optimally within the known (and/or visible) space.

In this paper, we propose a unified online collision-inclusive motion planning and control framework that evaluates the effect of possible collisions and decides when it might be preferred to collide with an obstacle (or a surface more broadly) instead of avoiding it. Our framework applies to impact-resilient robots with three core capabilities: 1) collision resilience, 2) collision identification, and 3) post-impact characterization. We design and fabricate in-house a custom omni-directional holonomic wheeled robot equipped with a collision ring that integrates Hall effect sensors along the arms holding the ring in place (Fig. 1); our robot satisfies all three core capabilities. The robot runs a Deformation Recovery control and trajectory Replanning (DRR) strategy [20] that enables it to recover from a collision and rapidly replan its post-impact trajectory using the information provided by the Hall effect sensors. The DRR strategy acts as the local replanner of the unified framework developed herein. We also propose and develop a global search-based planning algorithm based on the collision model generated from the DRR strategy. Similar to [21], our approach explores the space of trajectories using a set of short-duration motion primitives generated by solving an optimal control problem. Instead of pruning those primitives colliding with the obstacles in the global map, our proposed approach can adjust and evaluate them based on the collision model.

Succinctly, the paper’s contributions are as follows:

- We extend the DRR strategy to generate local trajectories when colliding with (non-)convex obstacles.

The authors are with the Dept. of Electrical and Computer Engineering, University of California, Riverside. Email: {zlu044, zliu017, mcamp077, karydis}@ucr.edu. We gratefully acknowledge the support of NSF #IIS-1910087, ONR #N00014-18-1-2252 and #N00014-19-1-2264, and ARL #W911NF-18-1-0266. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agencies.




Fig. 1: The evolution of our omni-directional holonomic robot prototypes used in our collision-inclusive motion planning and control research program. (a) The first iteration of the robot, inspired by the omnipuck robot [19], featured a passive collision ring and a single-board computer for motion control [1], [2]. This iteration has a radius of 0.12 m and weighs 0.6 kg. (b) The second iteration critically included an active collision ring-like structure that can sense collisions via embedded Hall effect sensors and a powerful onboard computer for online sensor data processing, decision making and motion control [20]. This iteration has a radius of 0.3 m and weighs 6 kg. A limitation of that prototype was the sparsity of its arms which reduced collision detection accuracy when the collision surface was not approximately perpendicular to any of the robot’s arms. (c) The most recent prototype developed in this present paper builds upon successful features of the second version and has significantly improved collision detection accuracy due to a complete redesign of the arms’ mechanical design and integration of more arms. This iteration has a radius of 0.3 m and weighs 8 kg.

- We develop a search-based planner to generate global trajectories and evaluate it in different benchmarks.
- We propose and evaluate a unified online collision-inclusive motion planning and control framework integrating the DRR strategy and search-based collision-inclusive planning while considering the robot’s FoV.

Our method is systematically evaluated via both simulated and real-world experiments using planar holonomic wheeled robot kinematics in environments that contain isolated convex and non-convex obstacles. We first test the DRR strategy experimentally to ensure its feasibility and safety when applied to the physical robot. Data collected from this process help identify parameters for the collision model which is necessary to test the search-based collision-inclusive algorithm in simulation. Comprehensive benchmark comparisons against state-of-art collision-avoidance and collision-inclusive methods demonstrate the differences, similarities and the utility of specific components, as well as of the overall proposed framework. Moreover, experimentation with the physical robot in a single corridor environment is conducted to validate the performance of our unified online collision-inclusive motion planning and control framework.

This paper builds upon and significantly extends previous results [2], [20]. The former [2] focuses only at the global planning level and evaluates possible collisions in the unknown (not yet observed) space. The latter [20] focuses only at the local control and planning level that utilizes DRR based on a-priori given waypoints. This paper, in contrast, develops the unified framework that combines the global planning and local control and planning levels together. In this newly-developed approach, the global planner can evaluate possible collisions both within and outside the robot’s FoV.¹

In what follows, we review related works in Sec. II and introduce our overall system’s structure in Sec. III. The deformation recovery control and post-impact trajectory replanning components are detailed in Sec. IV and Sec. V, respectively. The global search-based planner is discussed

in Sec. VI. Extensive benchmark (in simulation) and experimental results are given in Sec. VII. Sec. VIII discusses key findings and current limitations, and elaborates on future directions of research enabled by the proposed framework.

II. RELATED WORKS

Collision-free motion planning algorithms handle obstacle avoidance in distinct ways (e.g., [21]–[25]) to derive collision-free trajectories in real-time. Typically, such methods split the trajectory generation problem into two parts: 1) planning a collision-free geometric path or using motion primitives, and 2) optimizing the path locally to obtain a dynamically-feasible time-parameterized trajectory. When the environment is unknown (or partially-known), different strategies have been used based on those two-part framework. Many methods adopt the optimistic assumption [26], [27], which treats the unknown space as collision-free. This strategy improves the speed of reaching goals but may not guarantee safety. In contrast, other methods treat the unknown space as obstacle-occupied [28] and only allow for motions within the already known free space or FoV-observed free space [29]. Although these restrictions can help ensure safety, they tend to lead to conservative motion.

Tordesillas et al. [24] proposed a method that combines these two strategies by planning in both the known-free and unknown spaces. Instead of being overly optimistic about the unknown space, backup trajectories are also planned to enforce safety should the assumption about unknown space being free turns out to be wrong. While this method works well overall, it has put less emphasis on environment perception, lack of which may reduce safety or create over-conservative trajectories when the robot is tasked to operate at high speeds [30]. To this end, perception-aware strategies [23], [31], [32] have been proposed to predict unknown dangers and try to discover and avoid those dangers early on. However, prediction of unknown dangers does not necessarily ensure accuracy and usually requires additional computational effort which may limit online implementation.

Different from collision-avoidance, there have been efforts on designing impact-resilient robots that can withstand colli-

¹ To make this present paper self-contained, important methods and results from the previous related papers [2], [20] are included herein.

sions instead (e.g., [2], [4], [11], [13], [14], [18]–[20], [33]–[35]). With such robots as hand, one research direction has been to design characterization methods that can make the robot sense the collision and recover from the collision state. Most of such characterization methods have mainly focused on utilizing data from an onboard inertial measurement unit (IMU) [36]. However, IMUs are usually unable to distinguish collisions during aggressive maneuvers and to detect static contacts, resulting in low accuracy in collision detection. Sensors that could detect deformation of the robot during the collision process have been used in the past to provide more accurate collision detection [18], [37]. In related yet distinct previous work [14], we have implemented a passive quadrotor arm design with Hall effect sensors, making the robot able to detect and characterize collisions. The ability to sense and characterize collisions has led to various different methods to replan the local trajectory once the collision is detected [9], [11], [18]. Planning methods using motion patterns, e.g., to move forward in straight lines until collision with environment boundaries and then rotate in place and move forward again, have also been proposed [38]–[40]. Such methods can run online in environments with non-convex, polygon-shaped obstacles. A different trajectory generation method can be achieved by assuming the robot maintains contact with the obstacle [12]. Although these methods increase robustness and safety of post-impact trajectories, they cannot determine where the robot should collide with the environment to help it redirect toward the globally-planned goal.

Related works [3], [4] propose methods to evaluate and design possible collision spots of the global trajectory. Mote et al. [3] have introduced an empirical algebraic collision model by directly relating pre- and post-impact velocities with no thrust commanded. Then, a mixed-integer planning method based on that model is used to compute collision-inclusive trajectories in a known environment. However, integer constraints are hard to create, and solving a mixed-integer programming problem is usually time-consuming, making it impossible to run online for planning in an unknown environments. Further, the approach [3] has been demonstrated with a specific pair of objects over a relatively limited range of conditions (obstacles need to be line segments). Zha and Mueller [4] have proposed a rapidly exploring random tree (RRT) based planning method to plan global trajectories with collisions. The impulsive model is used to create the post-collision state once the pre-collision state is generated. Findings from [4] suggest that a collision-inclusive sampling-based planner is likely to find better trajectories in cluttered environments (such as narrow tunnels) as compared to environments that contain isolated obstacles. In addition, such algorithms remain limited in their use for online planning in unknown (or partially-known) environments.

Compared to our previous work [2], which developed a global path planner that explicitly trades-off between risk and collision exploitation only in unknown space, this paper proposes a new search-based global planner using a set of short-duration motion primitives which exploit possible collisions in the environment. The planner treats the un-

known space as collision-free. Our planner can generate waypoints with explicit information about possible collisions and more reasonable time allocation for the local trajectory generator. Contrary to collision avoidance methods with hard constraints that generate trajectories only in conservative local space [22], [24], our method utilizes gradient-based trajectory optimization (GTO) [23], which typically formulates trajectory generation as a nonlinear optimization problem and incorporates the artificial potential field (APF) to ensure safety. However, since GTO does not guarantee the robot will avoid all possible collisions, especially in unknown environments, we utilize the DRR strategy [20] for local trajectory generation; once a collision is sensed and characterized, DRR can ensure that the robot will recover from the collision and keep progressing toward its global goal.

III. SYSTEM OVERVIEW

A. Overall Framework

Our overall system architecture is shown in Fig. 2. Novel contributions relate to the low-level planner (Sec. IV and V) and the high-level planner (Sec. VI) in partially-known environments. The robot may collide with obstacles that were not detected at any time instant that the map (e.g., provided via LiDAR scans) refreshes. Instead of stopping when sensing the collision, the robot locally refines the trajectory and continues to explore the unknown space. To avoid repeated collisions with an obstacle (reminiscent of stacking into local minima) if another collision occurs while the robot follows the locally-revised trajectory part, the robot will then stop and invoke the high-level global planner to make more substantial refinements to the trajectory. Both processes run online.




Fig. 2: Overview of our unified framework for collision-inclusive motion planning and control. The method builds upon two novel components developed in this work; a global collision-inclusive planner and a local trajectory generator.

Contrary to collision avoidance algorithms, we do not impose any obstacle-related constraints in trajectory generation, nor we run a geometric collision check once a trajectory is generated at the low level. Instead, we directly generate a




Fig. 3: The DRR strategy for this work’s low-level planner.

trajectory based on given waypoints.² If a collision occurs, the robot receives a signal that a collision has occurred from any of the Hall effect sensors embedded between the main chassis and its deflection surfaces and activates a deformation recovery controller. The controller (Sec. IV) makes the robot detach from the collision surface by recovering from the deformation, and determines a post-collision state for the robot so as to facilitate post-impact trajectory replanning. The replanner (Sec. V) refines the initial trajectory since collisions change the second-order continuity of the trajectory followed before collision. To do so, the replanner uses the post-collision state computed by the recovery controller as initial state for refined trajectory generation. The procedure repeats as new collisions occur in the future, in a reactive and online manner. Figure 3 shows the DRR strategy, along with specific implementation components for experimentation.

We select GTO for post-impact recovery and global trajectory refinement based on [23] that revealed that GTO-based methods are particularly effective for local replanning, which is key for high-speed online motion planning in unknown environments. One drawback of GTO is the presence of local minima that may lead to undesirable solutions. Specifically, GTO may yield a trajectory that intersects with the obstacles in the environment [23]. Our DRR strategy can resolve this issue by offering a way to run a quick replan locally after the collision happens to ensure post-impact consistency.

B. Problem Assumptions and Notation

The proposed approach applies under the following:

- The boundary of the environment is known.
- Operating environments attain the form of confined corridors with isolated convex and non-convex obstacles, and only planar collisions obstacles are considered.
- During deformation and until a collided arm recovers its initial length, the tip of the arm remains in contact with the collision surface but does not rotate about the z axis, and the wheels of the robot contact the ground.
- The Hall effect sensor can return the information of collision state timely.

Key notation used in this paper is shown in Table I.

² The list of waypoints can be computed via any path planning method. It is independent from our proposed collision-inclusive planning algorithm.

TABLE I: Key notation.

l_0	neutral length vector of the spring
l_s	pre-tensioned spring length (arm not compressed)
l_e	length at maximum spring load following Hooke’s law
l	current spring length (deformation vector)
p	position vector of the robot
v	velocity vector of the robot
a	acceleration vector of the robot
s^-	state vector prior to the collision and recovery
s^+	state vector post to the collision and recovery
s_d	state vector of the robot (point of mass model)
k_e	spring constant of the arm.
k_d	damping coefficient of the arm.
τ	time interval
t_c	time instance when the sensing collision
T_r	time horizon of deformation recovery
T_{rep}	time horizon of replanning
${}^w_b R$	Rotation matrix from body frame to world frame
${}^w_c R$	Rotation matrix from collision frame \mathcal{F}_c to world frame
F	state transition matrix of deformation controller
G	input to state matrix of deformation controller
C	mapping matrix of polynomial coefficient η
Q	cost matrix of smoothness term
A_f	state transition matrix in free space
B_f	input to state matrix related to s_d in free space
J_s	smoothness term objective function
J_o	objective function of the clearance
J_v	penalty on velocity
J_a	penalty on acceleration

IV. DEFORMATION RECOVERY CONTROL

The purpose of our proposed deformation controller is to make the robot recover from a collision and reach a post-impact state that can facilitate recovery trajectory replanning (which we discuss in the next section).

A. Problem Setting

Consider a holonomic mobile robot (Fig. 1c), modeled as a point mass m . The robot’s main chassis is connected to its deflection surfaces via visco-elastic prismatic joints (Fig. 4). Note that the springs inside each joint are pre-tensioned. The robot’s compliant arms can both protect the robot from collision damage, and generate an external force driving it away from obstacles. External forces along each arm are caused via visco-elastic deformations assumed to follow the Voigt model; k_e and k_d denote the spring constant and damping coefficient, respectively. Hall effect sensors are used to measure the amount of deformation along each arm, and to signal collision detection when a user-tuned arm compression threshold is exceeded.³ The arm design with the bump sensor mechanism is similar to the button mechanism [41] and helps protect the robot from damage caused by collision as well as sense the collision in real time. Collision detection accuracy is related to the number of arms on the robot.

We consider four key quantities related to spring lengths: neutral l_0 , pre-tensioned l_s , maximum-load l_e , and current l (also referred to as deformation vector). These quantities play a significant role in the deformation recovery controller; they are also summarized in Table I, along with other key

³ The threshold is tuned based on the sensitivity of the Hall effect sensors.




Fig. 4: (Left) Model of our wheeled robot equipped with compliant arms. (Right) Close-up view of the assembly of the visco-elastic prismatic joint and Hall effect sensor.

notation. In single-arm collisions, current spring length vector l is aligned with the unit vector along the colliding arm, pointing from the tip of the arm to the center of robot along the compliant prismatic joint. For clarity of presentation, we consider in the following single-arm collisions. In multi-arm collisions we compute individual contributions from each colliding arm's spring and then consider their vector sum as the compound deformation vector used in lieu of l .

We use three coordinate systems. The world and body frames (${}^w_b R$ denotes the rotation matrix from body to world frames while ${}^b l$ denotes the deformation vector expressed in the body frame), and a (local) collision frame \mathcal{F}_c . This frame is defined at the time instant a collision occurs, t_c , and remains fixed for throughout the collision recovery process, T_r . Its origin coincides with the origin of the robot when a collision is detected. Basis vector $\{n, t, k\}$ of \mathcal{F}_c are defined normal, tangent and upwards with respect to the deformation vector l . Let θ be the angle of deformation vector l in \mathcal{F}_c .⁴

The (frame-agnostic) robot collision dynamics is given by $m\ddot{l} + k_d\dot{l} + k_e(l - l_0) = m\mathbf{a}_{in}$, where \mathbf{a}_{in} is the robot's body acceleration input as provided by the robot's motors.

B. Deformation Controller

The deformation recovery controller's task is to steer the post-impact state of the robot to a desired one within a time period of $[t_c, t_c + T_r]$. The time horizon T_r is an important hyper-parameter tuned by the user. Typically, longer T_r means the robot will recover from collision with longer time and smoother motion pattern. Through a preliminary calibration phase we selected $T_r = 0.5s$.

⁴ Note that inability to define the deformation vector may make the collision frame ill-defined. There are three special cases for this to happen. One is when two opposite arms deform exactly equally. In this case, there are two possible solutions to define the direction of the deformation vector along the line connecting the two arms. However, our algorithm still works as it prioritizes motion along the tangent to the collision vector (this would be the case of going through a very narrow straight corridor). The second case contains asymmetric collisions with three or more arms such that the vector sum is still zero. Then, one can define the collision frame based on the most dominant (in terms of magnitude) individual collision vector. Our algorithm can still work, though it is possible that more collisions will occur as the robot tries to navigate through (this would be the case of going through a very narrow curvy corridor). The last case is when there is an even (four or greater) number of symmetric collisions of exactly the same magnitude. The collision frame can no longer be defined. However, this case can only happen if the robot is radially pressed (entrapped) so that motion is completely restricted, which is not expected to occur during normal operation.

The deformation controller operates with respect to the local, collision frame \mathcal{F}_c . Let the state variable be ${}^c s = [{}^c p_x \ {}^c p_y \ {}^c \theta \ {}^c v_x \ {}^c v_y]^\top$. The control input is $\mathbf{u} = [u_x \ u_y \ u_\theta]^\top$, where $u_x = ({}^c \mathbf{a}_{in} - \frac{k_e}{m}({}^c \mathbf{l}_s - {}^c \mathbf{l}_0)) \cdot {}^c \mathbf{n}$, $u_y = ({}^c \mathbf{a}_{in} - \frac{k_e}{m}({}^c \mathbf{l}_s - {}^c \mathbf{l}_0)) \cdot {}^c \mathbf{t}$, and $u_\theta = {}^c \boldsymbol{\omega} \cdot {}^c \mathbf{k}$ with ${}^c \boldsymbol{\omega}$ being the angular velocity of the robot in the collision frame. Note that position control terms include compensation for the force caused by the spring being pre-tensioned when the robot's arm is at its rest length. Then, the state space model of the robot recovering from collision can be expressed as

$$\begin{cases} \dot{p}_x = v_x \\ \dot{p}_y = v_y \\ \dot{\theta} = u_\theta \\ \dot{v}_x = -\frac{k_e}{m}p_x - \frac{k_d}{m}v_x + u_x \\ \dot{v}_y = -\frac{k_e(\mu \text{sign}(v_y) + \tan \theta)p_x + f_0}{m} - \frac{k_d(\mu \text{sign}(v_y) + \tan \theta)v_x}{m} + u_y \end{cases} \quad (1)$$

where $f_0 = \mu k_e \text{sign}(v_y)({}^c \mathbf{l}_s - {}^c \mathbf{l}_0) \cdot {}^c \mathbf{n}$.

Since the robot is holonomic, we can decouple orientation from position control.⁵ The orientation and angular velocity errors during recovery time $t \in [t_c, t_c + T_r]$ are $e_R(t) = \frac{1}{2}(R_d^\top R - R^\top R_d)^\vee$ and $e_{\dot{R}}(t) = \boldsymbol{\omega} - R^\top R_d \boldsymbol{\omega}_d$, respectively.⁶ Index d denotes desired quantities; these are $R_d = R(t_c)$ and $\boldsymbol{\omega}_d = [0 \ 0 \ 0]^\top$. (All terms are with respect to collision frame \mathcal{F}_c .) Then,

$$u_\theta = -K_r e_{R,z}(t) - K_\omega e_{\dot{R},z}(t). \quad (2)$$

Note that since this is a planar collision problem, the collision recovery orientation controller considers only the z -components of orientation and angular velocity errors.

Regarding collision recovery position control, note that the translation-only motion in (1) is affine. Thus, we can apply feedback linearization. The linearized system matrix F is

$$F = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -\frac{k_e}{m} & 0 & -\frac{k_d}{m} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

with state vector $\mathbf{s}_d = [p_x \ p_y \ v_x \ v_y]^\top$. The control input matrix is $G = I_{2 \times 2}$ with control input vector $\boldsymbol{\nu} = [\nu_x \ \nu_y]$ given by

$$\begin{cases} \nu_x = u_x \\ \nu_y = u_y - \frac{k_e(\mu \text{sign}(v_y) + \tan \theta)p_x + f_0}{m} - \frac{k_d(\mu \text{sign}(v_y) + \tan \theta)v_x}{m} \end{cases} \quad (3)$$

We formulate an optimal control problem with fixed time horizon T based on the linearized system $\dot{\mathbf{x}} = F\mathbf{x} + G\boldsymbol{\nu}$.

⁵ In our approach we seek to make the robot keep the same orientation it has at the instant it collides throughout the collision recovery process. We follow this approach because it can simplify the overall deformation recovery control problem without sacrificing optimality.

⁶ The vee map \vee is the inverse of a skew-symmetric mapping.

Using the change of variable $\tau = t - t_c$,⁷ we seek to solve

$$\min_{\mathbf{s}_d} \int_0^{T_r} (\mathbf{s}_d(\tau)^\top \Gamma \mathbf{s}_d(\tau) + \boldsymbol{\nu}^\top(\tau) H \boldsymbol{\nu}(\tau)) d\tau \quad (4a)$$

$$\text{subject to } \dot{\mathbf{s}}_d = F \mathbf{s}_d + G \boldsymbol{\nu}, \quad (4b)$$

$$- \|\mathbf{l}_e - \mathbf{l}_s\| \cos \theta \leq p_x \leq 0. \quad (4c)$$

$$\mathbf{s}_d(0) = [p_{0,x} \ 0 \ v_{0,x} \ v_{0,y}]. \quad (4d)$$

$$\mathbf{s}_d(T_r) = [0 \ p_{T,y} \ v_{T,x} \ v_{T,y}]. \quad (4e)$$

Matrices $\Gamma = \gamma \begin{bmatrix} I_{2 \times 2} & 0 \\ 0 & 0_{2 \times 2} \end{bmatrix}$ and $H = h I_{2 \times 2}$ penalize the displacement during the recovery process and the control input, respectively. There is a trade-off between the displacement and the control input of the robot. Tuning parameters γ and h balance this trade-off to select the controller with minimal control energy and displacement.

Constraint (4c) dictates that the robot should be in contact with the collision surface until the colliding arm's spring has recovered its original, pre-tensioned length \mathbf{l}_s (i.e. the arm is no longer compressed) without compressing beyond its linear region \mathbf{l}_e . Constraints (4d) and (4e) enforce initial and terminal position and velocity conditions, respectively. In detail, $p_{0,x}$ is determined by the colliding arm's Hall effector sensor reading. Since the vector form of the sensor's reading (that is, ${}^b \mathbf{l} - {}^b \mathbf{l}_s$) is expressed in the body frame, we need transform it to the collision frame \mathcal{F}_c as per

$$p_{0,x} = -[1 \ 0] {}^w R_c^\top {}^w R_b ({}^b \mathbf{l} - {}^b \mathbf{l}_s). \quad (5)$$

The velocity components at the collision instant $v_{0,x}$ and $v_{0,y}$ are expressed in frame \mathcal{F}_c and are estimated at runtime.⁸ Post-impact, the arm needs to be uncompressed (hence $p_{T,x}$ is set to 0), but $p_{T,y}$ is treated as an unconstrained free variable. Post-impact terminal velocity components $v_{T,x}$ and $v_{T,y}$ are also expressed in \mathcal{F}_c and can be set freely. In Sec. V, we discuss how to generate $v_{T,x}$ and $v_{T,y}$ based on the preplanned trajectory. We discretize the linearized system in (4b) with sampling frequency $f = 10$ Hz using the Euler method, and solve the corresponding quadratic program with CVXOPT. The process is summarized in Alg. 1.

Computed control inputs (4) and (2) make the robot detach from the collision surface and help bring it to a temporary post-collision state which can be used as the initial condition for post-impact trajectory generation. We discuss this next.

V. POST-IMPACT TRAJECTORY REPLANNING

A. Problem Formulation

We formulate the post-impact trajectory generation problem as a quadratic program with equality constraints, i.e.

⁷ We employ this change of variable for clarity. Problem (4) resets every time a new collision occurs; this gives rise to an LTI system, hence the change of variable can apply.

⁸ In the experiments conducted in this work, velocity measurements are provided via a motion capture camera system, but the method applies as long as velocity estimates are available, e.g., via optical flow.

Algorithm 1: Deformation recovery controller

input : Displacement in body frame ${}^b \mathbf{l} - {}^b \mathbf{l}_s$ via Hall effect sensors readings; collision time instant $\tau_c \in [0, \Delta t_{i_c}]$; position in world frame at collision instant, ${}^w \mathbf{p}_{\tau_c}$; velocity in world frame at collision instant, ${}^w \mathbf{v}_{\tau_c}$; rotation matrix ${}^w R_b$; rotation matrix ${}^c R$; next waypoint point in world frame, ${}^w \mathbf{p}_{next}$.

output : Control input \mathbf{u}

parameter: Maximum velocity of the robot v_{max}

1 **Function** RECOVERYCONTROLLER (${}^b \mathbf{l} - {}^b \mathbf{l}_s, \tau_c, \Delta t_{i_c}, {}^w \mathbf{p}_{\tau_c}, {}^w \mathbf{p}_{next}, {}^w \mathbf{v}_{\tau_c}, {}^w R, {}^c R$):

2 ${}^w \mathbf{v}_T \leftarrow \frac{{}^w \mathbf{p}_{next} - {}^w \mathbf{p}_{\tau_c}}{\Delta t_{i_c} - \tau_c}$

3 ${}^c \mathbf{v}_T \leftarrow {}^c R^\top {}^w \mathbf{v}_T$

4 **if** ${}^c v_{T,x} < 0$ **then**

5 | ${}^c v_{T,x} \leftarrow 0$

6 **end**

7 **if** $\|{}^c \mathbf{v}_T\| \geq v_{max}$ **then**

8 | ${}^c \mathbf{v}_T \leftarrow v_{max} \text{normalize}({}^c \mathbf{v}_T)$

9 **end**

10 Calculate $p_{0,x}$ based on (5) with \mathbf{l}_b

11 $p_{0,y} \leftarrow 0$

12 Calculate u_x and u_y based on (4) and (3) with given \mathbf{v}_T and \mathbf{p}_0

13 Calculate u_θ based on (2)

14 $\mathbf{u} \leftarrow [u_x \ u_y \ u_\theta]^\top$

15 **return** \mathbf{u}

$$\min_{\boldsymbol{\eta}} J_s(\boldsymbol{\eta}) = \sum_{i=i_c}^{N_I} \int_0^{\Delta t_i} \| {}^w \mathbf{p}_i^{(q)}(t) \| dt \quad (6a)$$

$$\text{subject to } C_{0,i_c,\beta}^{(0)} \boldsymbol{\eta}_{i_c,\beta} = {}^w \mathbf{p}_{r,\beta}, \quad (6b)$$

$$C_{0,i_c,\beta}^{(1)} \boldsymbol{\eta}_{i_c,\beta} = {}^w \mathbf{v}_{r,\beta}, \quad (6c)$$

$$C_{\Delta t_{N_I}, N_I, \beta}^{(\alpha)} \boldsymbol{\eta}_{N_I, \beta} = {}^w \mathbf{d}_{\Delta t_{N_I}, N_I, \beta}^{(\alpha)},$$

$$\alpha = \{0, 1 \dots q-1\}, \quad (6d)$$

$$C_{\Delta t_i, i, \beta}^0 \boldsymbol{\eta}_{i, \beta} = {}^w \mathbf{p}_{i+1, \beta}, \quad (6e)$$

$$C_{\Delta t_i, i+1, \beta}^{(\alpha)} \boldsymbol{\eta}_{i, \beta} = C_{0, i+1, \beta}^{(\alpha)} \boldsymbol{\eta}_{i+1, \beta},$$

$$\alpha = \{1, 2 \dots q-1\}. \quad (6f)$$

For polynomial segments, we can rewrite J_s as

$$J_s = \sum_{\beta \in \{x, y\}} \sum_{i=i_c}^{N_I} \boldsymbol{\eta}_{i, \beta}^\top Q_\beta^q(\Delta t_i) \boldsymbol{\eta}_{i, \beta},$$

where i_c is the segment where the collision happens and N_I is the number of trajectory segments. Superscript q denotes the derivative order; for example, $q = \{1, 2, 3, 4\}$ correspond to min-velocity, min-acceleration, min-jerk and min-snap trajectories, respectively. Subscript $\beta \in \{x, y\}$ indicates the x and y component of the trajectory, and Δt_i is the time duration for i^{th} polynomial segment. Parameter $\boldsymbol{\eta}_{i, \beta}$ is the vector of coefficients of i^{th} polynomial. $C_{0, i, \beta}^{(\alpha)}$ maps

the coefficients to α^{th} order derivative of the start point in segment i , while $C_{\Delta t_i, i, \beta}^{(\alpha)}$ maps the coefficients to α^{th} order derivative of the end point in segment i .

Constraints (6b) and (6c) impose the initial values for the 0^{th} and the 1^{st} order derivatives to match the position and velocity values attained via the collision recovery controller, respectively. Constraint (6d) imposes that the α^{th} order derivatives of the end position are fixed. Constraint (6e) imposes that the trajectory will pass through desired waypoints after i_c . Constraint (6f) is imposed to ensure α^{th} continuity among polynomial segments.

We solve this quadratic programming (QP) problem given initial (post-collision) and end states, and intermediate waypoints. Then, we perform time scaling as in [22] to reduce the maximum values for planned velocities, accelerations and higher-order derivatives as appropriate, and thus improve dynamic feasibility of the refined post-impact trajectory.

The solution of the QP problem serves as the initial value for GTO [42], where we change the objective function to

$$\min \quad \lambda_s J_s + \lambda_o J_o + \lambda_d (J_v + J_a) , \quad (7)$$

where J_o is the cost to avoid collisions, and J_v and J_a are the penalties when candidate velocity and acceleration solutions exceed the dynamic feasibility limit, respectively. Weight parameters λ_s , λ_o and λ_d trade off between smoothness, trajectory clearance and dynamical feasibility, respectively.

Similar to [42], we use an exponential cost function. At a position with distance d to the closest obstacle, the cost $c_o(d)$ is written as

$$c_o(d) = \alpha_o \exp(-d - d_o)/\gamma_o , \quad (8)$$

where α_o is the magnitude of the cost function, d_o is the threshold where the cost starts to rapidly rise, and γ_o controls the rate of the function's rise. Then, J_o can be computed as

$$\begin{aligned} J_o &= \sum_{i=i_c}^{N_I} \int_0^{\Delta t_i} c_o(\mathbf{p}(t)) \|\mathbf{v}(t)\| dt \\ &= \sum_{i=i_c}^{N_I} \sum_{k=0}^N c_o(\mathbf{p}(t_k)) \|\mathbf{v}(t_k)\| \delta t . \end{aligned} \quad (9)$$

J_v can be computed in a similar manner, whereby $c_v(v)$ is the cost function applied on the velocity and attains the same form as in (8). We can then obtain

$$\begin{aligned} J_v &= \sum_{\beta \in \{x, y\}} \sum_{i=i_c}^{N_I} \int_0^{\Delta t_i} c_v(v_\beta(t)) \|a_\beta(t)\| dt \\ &= \sum_{\beta \in \{x, y\}} \sum_{i=i_c}^{N_I} \sum_{k=0}^N c_v(v_\beta(t_k)) \|a_\beta(t_k)\| \delta t \end{aligned} \quad (10)$$

The formulation of J_a is similar to (10). The cost function of the acceleration constraint $c_a(a)$ is also an exponential function similar to $c_o(d)$ and $c_v(v)$, since it can penalize when close to or beyond acceleration bounds while staying flat when away from the bounds. We apply a similar Newton trust region method as in [42] to optimize the objective.




Fig. 5: Waypoint adjustment process when the collision is sensed.

B. Waypoint Adjustment

In some cases, it may be necessary to adjust the waypoints given by a preplanned trajectory with the information obtained from the collision, and then solve the aforementioned problem in Sec. V-A with the adjusted waypoints. Such cases occur when there is no direct line of sight between the collision state and the waypoint at the end of the immediately next trajectory segment following collision recovery. By enabling such waypoint adjustment, the algorithm promotes exploration and in certain cases prevents the robot from being trapped in a local minima in which repeated collisions at the same (or very close-by) place could otherwise occur.

With reference to Alg. 2, we express in the local collision frame \mathcal{F}_c the next waypoint `waypoint_list`[$i_c + 1$] (lines 2–4). In line 5, we adjust `waypoint_list` with the information we get from collision. Details of this process are shown in Fig. 5. We add an additional waypoint \mathbf{p}_{add} to create a path detouring the collided obstacle. Then we select the shortest path among all the possible paths toward the next waypoint \mathbf{p}_{next} that was originally in the list before collision. Possible \mathbf{p}_{add} waypoints are generated by either using a path generation algorithm (e.g., jump point search) when the complete collision surface can be perceived, or by searching along the y -axis of collision frame \mathcal{F}_c by a (user-defined) exploration distance $\epsilon_{explore}$ when the complete collision surface cannot be reliably perceived (e.g., via LiDAR measurements).

Algorithm 2: Post-impact waypoint adjustment

input : Position after the collision recovery in world frame, ${}^w\mathbf{p}_r$; waypoint list of preplanned trajectory; `waypoint_list`; wR ; trajectory segment i_c where the collision happens.

output : waypoint list after adjustment
`waypoint_list`

parameter: Robot radius r_{rob}

```

1 Function WAYPOINTADJUSTMENTLINE (  ${}^w\mathbf{p}_r$ ,
  waypoint_list,  ${}^wR$ ,  $i_c$  ) :
2    ${}^w\mathbf{p}_{next} \leftarrow \text{waypoint\_list}[i_c + 1]$ 
3   Transfer  ${}^w\mathbf{p}_{next}$  into  $\mathcal{F}_c$  frame to get  ${}^c\mathbf{p}_{next}$ 
4   Transfer  ${}^w\mathbf{p}_r$  into  $\mathcal{F}_c$  frame to get  ${}^c\mathbf{p}_r$ 
5   Adjust waypoint_list as Fig. 5
6 return waypoint_list

```

As the robot progresses and reaches the additional waypoint \mathbf{p}_{add} that was added following the collision, then it replans based on latest information provided from the perception module. This happens when the robot either reaches

the added waypoint (to ensure that the next waypoint is in free space) or it senses another collision from the deformation sensor. In this case, the original \mathbf{p}_{next} will change as well. Note that this process runs online. In the case that the robot senses a collision before reaching \mathbf{p}_{add} , then it will recover and stop (instead of running the fast replanning approach listed above) and call the global planner to revise larger parts of the trajectory. If a new waypoint is inserted in the list, we map the path generated by ${}^w p_r$ and waypoints in the list after $i_c + 1$ into time domain using a trapezoidal velocity profile. If no new waypoint is inserted, we set the time duration of i_c segment in (6) as $\Delta t_{i_c} = t_{i_c+1} - t_c$, where t_{i_c+1} is the time reaching next waypoint p_{i_c+1} in the preplanned trajectory.

VI. SEARCH-BASED COLLISION-INCLUSIVE PLANNING

In this section, we propose the main algorithm to generate the waypoint list and trajectory segments that serve as the input to the DRR strategy.

A. Problem Formulation

Let the system state $\mathbf{s}_d(t) \in \mathcal{S} \subset \mathbb{R}^{n \times q}$ contain the configuration and the $(q-1)^{th}$ -order derivatives in 2D (i.e. $n = 2$). The free state space, $\mathcal{S}^{free} \subset \mathcal{S}$, contains both obstacle-free configurations, \mathcal{P}^{free} , as well as the system's dynamical constraints, \mathcal{D}^{free} , which include minimum and maximum bounds on velocity $[v_{min}, v_{max}]$, acceleration $[a_{min}, a_{max}]$, jerk $[j_{min}, j_{max}]$, and other higher-order derivatives. We can then write $\mathcal{S}^{free} = \mathcal{P}^{free} \times \mathcal{D}^{free} = \mathcal{P}^{free} \times [v_{min}, v_{max}] \times [a_{min}, a_{max}] \times \dots$. $\mathcal{P}^{obs} = \mathcal{P} \setminus \mathcal{P}^{free}$ and $\mathcal{S}^{obs} = \mathcal{P}^{obs} \times \mathcal{D}^{free}$ defines the obstacle region.

The differential flatness of some mobile robot systems (e.g., [21]) helps design control inputs from 1D time-parameterized polynomial trajectories independently for each of the n positions. Hence, $\mathbf{s}_d(t) = [\mathbf{p}_D(t)^\top, \dot{\mathbf{p}}_D(t)^\top, \dots, \mathbf{p}_D^{(q-1)}(t)^\top]^\top$ where $\mathbf{p}_D(t) = \sum_{i=0}^q \mathbf{d}_i \frac{t^i}{i!}$ and $\mathbf{D} = [\mathbf{d}_0, \dots, \mathbf{d}_q] \in \mathbb{R}^{n \times (q+1)}$, and $\mathbf{d}_i = [d_{i,x}, d_{i,y}]^\top$ in (6). To simplify the notation, we re-express the derivatives as $\mathbf{v}(t) = \dot{\mathbf{p}}_D^\top(t)$, $\mathbf{a}(t) = \ddot{\mathbf{p}}_D^\top(t)$, $\mathbf{j}(t) = \dddot{\mathbf{p}}_D^\top(t)$, etc., and drop subscript D .

We can construct the polynomial trajectories via $\mathbf{p}_D^{(q)}(t) = \mathbf{u}(t)$ with controls $\mathbf{u}(t) \in \mathcal{U} = [-u_{max}, u_{max}]^n \subset \mathbb{R}^n$. In state space form this yields $\dot{\mathbf{s}}_d(t) = A_f \mathbf{s}_d(t) + B_f \mathbf{u}(t)$, with

$$A_f = \begin{bmatrix} 0 & I_n & 0 & \dots & 0 \\ 0 & 0 & I_n & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \dots & \dots & 0 & I_n \\ 0 & \dots & \dots & 0 & 0 \end{bmatrix}, B_f = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ I_n \end{bmatrix}. \quad (11)$$

In collision-inclusive planning, we consider a smoothness cost $J_s(\mathbf{D}) = \sum_{k=1}^K \int_0^{T_k} \|\mathbf{u}(t)\|^2 dt = \sum_{k=1}^K \int_0^{T_k} \|\mathbf{p}_{D,k}^{(q)}(t)\|^2 dt$. The trajectory is not q th order differentiable as it would be in collision avoidance. The smoothness of the entire trajectory is the sum of its q th order differentiable segments. We consider

two additional costs. First, $T_g = \sum_{k=1}^K T_k + (K-1)T_r$ penalizes the overall trajectory duration. Then,

$$J_c = \begin{cases} \frac{(|{}^c v_x^+| - |{}^c v_x^-|)^2 + \sum_{\beta \in \{y,z\}} \Delta E_\beta}{T_r} & \forall \zeta(t) = 1 \\ 0 & otherwise \end{cases} \quad (12)$$

evaluates the effect of a collision in changing the direction of motion of the robot. $\Delta E_\beta = ({}^c v_\beta^+ - {}^c v_\beta^-)^2$, where ${}^c v_\beta^+ = {}^c v_\beta(t + T_r)$ and ${}^c v_\beta^- = {}^c v_\beta(t)$. ${}^c v_\beta^+$ can be approximated via Alg. 4 if \mathbf{p}_{goal} is known. (We discuss Alg. 4 in detail in Sec. VI-D.) We also define an indicator function $\zeta(t) = \{0, 1\}$ that signals if the robot is colliding at time t .

We can then define the optimization problem

$$\min_{\mathbf{D}, T_g} J_s(\mathbf{D}) + \rho_t T_g + \rho_c J_c(t) \quad (13a)$$

$$\text{subject to } \dot{\mathbf{s}}_d(t) = A_f \mathbf{s}_d(t) + B_f \mathbf{u}(t), \forall \zeta(t) = 0, \quad \forall t \in [0, T_g] \quad (13b)$$

$$\mathbf{s}_d(t + T_r) = F_{DRR}(\mathbf{s}_d(t)), \forall \zeta(t) = 1, \quad \forall t \in [0, T_g - T_r] \quad (13c)$$

$$\zeta(t) \in \{0, 1\}, \forall t \in [0, T_g] \quad (13d)$$

$$\zeta(t) = 0, \text{ if } \mathbf{s}_d(t + \delta t) \in \mathcal{S}^{free}, \delta t \rightarrow 0, \quad \forall t \in [0, T_g] \quad (13e)$$

$$\zeta(t) = 1, \text{ if } \mathbf{s}_d(t + \delta t) \in \mathcal{S}^{obs}, \delta t \rightarrow 0, \quad \forall t \in [0, T_g - T_r] \quad (13f)$$

$$\mathbf{s}_d(0) = \mathbf{s}_{d0}, \mathbf{s}_d(T_g) \in \mathcal{S}^{goal}, \quad \zeta(0) = 0, \zeta(T_g) = 0 \quad (13g)$$

$$\mathbf{s}_d(t) \in \mathcal{S}^{free}, \mathbf{u}(t) \in \mathcal{U}, \forall t \in [0, T_g] \quad (13h)$$

$${}^c v_x(t) \in \mathcal{V}^c, \text{ if } \zeta(t) = 1. \quad (13i)$$

Parameters $\rho_t > 0$ and $\rho_c > 0$ regulate the relative importance of trajectory smoothness, duration, and amount of collisions that switch the direction of motion. Conditions (13e) and (13f) determine how the value for $\zeta(t)$ is being set. In (13i), $\mathcal{V}^c = [-v_{max,c}, v_{max,c}]$; $v_{max,c}$ indicates the maximum collision velocity which, if exceeded, will lead to the robot flipping over. Thus, we set the pre-collision velocity component along x axis of \mathcal{F}_c as $v_x(t) \in \mathcal{V}^c$.

Herein we show that, similar to the collision avoidance motion planning problem [21], safety constraints may be addressed by reformulating problem (13) into a deterministic shortest path one with $(n \times q)$ state \mathcal{S} and n control \mathcal{U} . Since the dimensionality of \mathcal{U} remains at n , search-based planning (e.g., A* [43]) that discretizes \mathcal{U} using motion primitives can be an effective way to determine in finite-time resolution-complete and optimal (in the discretized space) trajectories.

B. Motion Primitives

Choosing a number of samples $r \in \mathbb{Z}^+$ along each axis $[-u_{max}, u_{max}]$, which defines a discretization step: $du = \frac{u_{max}}{r}$ and results in $M = (2r + 1)^n$ motion primitives, is one technique to acquire the discretization \mathcal{U}_M . Given initial state $\mathbf{s}_{d0} = [\mathbf{p}_0^\top, \mathbf{v}_0^\top, \dots]^\top$, we generate a motion primitive

of duration $\tau > 0$ that applies piece-wise constant control

$$\tilde{\mathbf{u}}_m(t) = \begin{cases} \mathbf{u}_m & \mathbf{p}_D^{(q-1)}(t) \in [\mathbf{p}_{D,min}^{(q-1)}, \mathbf{p}_{D,max}^{(q-1)}], \\ \mathbf{0} & \mathbf{p}_D^{(q-1)}(t) \notin [\mathbf{p}_{D,min}^{(q-1)}, \mathbf{p}_{D,max}^{(q-1)}] \end{cases} \quad (14)$$

where $\mathbf{u}_m \in \mathcal{U}_M$ for $t \in [0, \tau]$. Given initial conditions,

$$\mathbf{p}_D(t) = \tilde{\mathbf{u}}_m(t) \frac{t^q}{q!} + \dots + \mathbf{v}_0 t + \mathbf{p}_0 \quad (15)$$

is a piece-wise function. The resulting trajectory of (11) is

$$\mathbf{s}_d(t) = \underbrace{e^{A_f t}}_{A_{df}(t)} \mathbf{s}_{d0} + \underbrace{\int_0^t e^{A_f(t-\sigma)} B_f \tilde{\mathbf{u}}_m(\sigma) d\sigma}_{B_{df}(t) \mathbf{u}_m} \quad (16)$$

By beginning at \mathbf{s}_{d0} and applying all primitives to acquire the M possible states after $\tau \in [0, \tau_f]$ (Alg. 3), we can create a graph representation of the attainable system states. There will be M^2 potential states at time 2τ if all primitives are applied to each of the M states once again. The set of reachable states \mathcal{S}^r is finite given the free space \mathcal{S}^{free} is bounded. These enable the construction of a graph the states of which are connected by a motion primitive $e = (\tilde{\mathbf{u}}_m, \tau, \xi)$ with ξ being an integer (discussed in Sec. VI-D).

We construct the graph to explore the free state space \mathcal{S}^{free} using Alg. 3. Given the constant time upper-bound τ_f and the fully specified state \mathbf{s} , the primitive is derived in line 4 using the control input \mathbf{u}_m ; lines 5–23 check whether the primitive intersects with the obstacles and then modify those primitives intersecting with the obstacles. This step will be further discussed Sec. VI-D. In lines 6–13, we evaluate the end state of a valid primitive not intersecting with the obstacles and we add it to the set of successors of the current node; meanwhile, we estimate the edge cost from the corresponding primitive. In lines 16–21, we modify the end state of the primitive and add it to the set of successors of the current node; meanwhile, we estimate the edge cost related to the corresponding modified primitive. Line 19 shows that we consider T_r for the robot recovering from the collision using DRR in the cost function. Further modification of the cost function about estimating the cost related to J_c part will be discussed in Sec. VI-D. The nodes in the successor set $\mathcal{R}(\mathbf{s}_d)$ are added to the graph after we have checked all the primitives in the finite control input set. Finally, the graph keeps growing until we reach the goal is reached.

C. Deterministic Shortest Trajectory

We can re-formulate (13) as a graph-search problem using the set of motion primitives \mathcal{U}_M and the induced discretization. To do so, we introduce additional constraints for the control input $\mathbf{u}(t)$ in (13) to be piecewise-constant. We introduce an additional variable $N \in \mathbb{Z}^+$, so that $T_g = \sum_{k=0}^{N-1} (\tau_k + \zeta_{k+1} T_r)$, and $\tilde{\mathbf{u}}_k$ is computed by (14) with $\mathbf{u}_k \in \mathcal{U}_M$ for $k = 0, \dots, N-1$ and a constraint in (13h):

$$\mathbf{u}(t) = \sum_{k=0}^{N-1} \tilde{\mathbf{u}}_k \mathbb{1}_{t \in [T_k, T_{k+1}]} \quad (17a)$$

Algorithm 3: Collision-inclusive motion primitive generation

input : Initial state $\mathbf{s}_d \in \mathcal{S}^r \subset \mathcal{S}^{free}$; motion primitive set \mathcal{U}_M , upper-bound of duration τ_f

output : Reachable set $\mathcal{R}(\mathbf{s}_d)$ from \mathbf{s}_d in one step; costs set $\mathcal{C}(\mathbf{s}_d)$; duration set $\mathcal{T}(\mathbf{s}_d)$; collision states set $\mathcal{Z}(\mathbf{s}_d)$

parameter: Time interval δt ; recover time T_r in DRR

- 1 **Function** GETMOTIONPRIMITIVE ($\mathbf{s}_d, \mathcal{U}_M, \tau_f$) :
- 2 $\mathcal{R}(\mathbf{s}_d) \leftarrow \emptyset, \mathcal{C}(\mathbf{s}_d) \leftarrow \emptyset, \mathcal{T}(\mathbf{s}_d) \leftarrow \emptyset, \mathcal{Z}(\mathbf{s}_d) \leftarrow \emptyset$
- 3 **for all** $\mathbf{u}_m \in \mathcal{U}_M$ **do**
- 4 Calculate edge $e_m(t)$ according to (16) for $t \in [0, \tau_f]$
- 5 **if** $e(t) \in \mathcal{S}^{free}$ **for all** $t \in [0, \tau_f]$ **then**
- 6 $\zeta_m \leftarrow 0$
- 7 $\tau_m \leftarrow \tau_f$
- 8 $\mathbf{s}_{d,m} \leftarrow e_m(\tau_f)$
- 9 $\mathcal{R}(\mathbf{s}_d) \leftarrow \mathcal{R}(\mathbf{s}_d) \cup \{\mathbf{s}_{d,m}\}$
- 10 $J_D \leftarrow \int_0^{\tau_f} \|\tilde{\mathbf{u}}_m(t)\|^2 dt$
- 11 $\mathcal{C}(\mathbf{s}_d) \leftarrow \mathcal{C}(\mathbf{s}_d) \cup \{J_D + \rho_t \tau_f\}$
- 12 $\mathcal{T}(\mathbf{s}_d) \leftarrow \mathcal{T}(\mathbf{s}_d) \cup \{\tau_m\}$
- 13 $\zeta(\mathbf{s}_d) \leftarrow \mathcal{Z}(\mathbf{s}_d) \cup \{\zeta_m\}$
- 14 **else**
- 15 $\zeta_m \leftarrow 1$
- 16 Generate $\mathbf{s}_{d,m}, \tau$ and calculate J_c or prune this primitive (discussed in VI-D and VI-F).
- 17 $\mathcal{R}(\mathbf{s}_d) \leftarrow \mathcal{R}(\mathbf{s}_d) \cup \{\mathbf{s}_{d,m}\}$
- 18 $J_D \leftarrow \int_0^{\tau_f} \|\tilde{\mathbf{u}}_m(t)\|^2 dt$
- 19 $\mathcal{C}(\mathbf{s}_d) \leftarrow \mathcal{C}(\mathbf{s}_d) \cup \{J_D + \rho_t(\tau + T_r) + J_c\}$
- 20 $\mathcal{T}(\mathbf{s}_d) \leftarrow \mathcal{T}(\mathbf{s}_d) \cup \{\tau_m\}$
- 21 $\zeta(\mathbf{s}_d) \leftarrow \mathcal{Z}(\mathbf{s}_d) \cup \{\zeta_m\}$
- 22 **end**
- 23 **end**
- 24 **return** $\mathcal{R}(\mathbf{s}_d), \mathcal{C}(\mathbf{s}_d), \mathcal{T}(\mathbf{s}_d), \mathcal{Z}(\mathbf{s}_d)$

By letting $T_i = \sum_{k=0}^{i-1} \tau_k$ we can force the control trajectory to be a composition of the motion primitives in \mathcal{U}_M . Given an initial state $\mathbf{s}_{d0} \in \mathcal{S}^{free}$, a goal area \mathcal{S}^{goal} and a finite set of motion primitives \mathcal{U}_M with duration $\tau > 0$, we seek to select a series of motion primitives $\mathbf{u}_{0:N-1}$ of length N , such that

$$\min_{N, \mathbf{u}_{0:N-1}} \sum_{k=0}^{N-1} \|\mathbf{u}_k\|^2 + \rho_t(\tau_k + \zeta_{k+1} T_r) + \rho_c J_{c,k} \quad (17a)$$

$$\text{subject to } \mathbf{s}_d(\tilde{t}) = A_{df}(\tilde{t}) \mathbf{s}_{d,k} + B_{df}(\tilde{t}) \mathbf{u}_k \in \mathcal{S}^{free}, \quad \forall \tilde{t} \in [0, \tau_k] \quad (17b)$$

$$\zeta_k \in \{0, 1\}, \forall k \in \{0, 1, \dots, N-1\} \quad (17c)$$

$$\zeta_{k+1} = 0, \text{ if } \mathbf{s}_d(\tau_k + \delta t) \in \mathcal{S}^{free}, \delta t \rightarrow 0 \quad (17d)$$

$$\zeta_{k+1} = 1, \text{ if } \mathbf{s}_d(\tau_k + \delta t) \in \mathcal{S}^{obs}, \delta t \rightarrow 0 \quad (17e)$$

$$\mathbf{s}_{d,k+1} = \mathbf{s}_d(\tau_k), \forall \zeta_{k+1} = 0 \quad (17f)$$

$$\mathbf{s}_{d,k+1} = F_{DRR}(\mathbf{s}_d(\tau_k)), \forall \zeta_{k+1} = 1 \quad (17g)$$

$$\mathbf{s}_{d,0} = \mathbf{s}_{d0}, \mathbf{s}_{d,N} \in \mathcal{S}^{goal}, \zeta_0 = 0, \zeta_N = 0 \quad (17h)$$

$$\mathbf{u}_k \in \mathcal{U}_M \quad (17i)$$

$${}^c v_{k,x}^- \in \mathcal{V}^c, \text{ if } \zeta_k = 1. \quad (17j)$$

The optimal cost of (17) is an upper bound to the optimal cost of (13) because (17) is a constrained version of (13). The whole trajectory consists of a set of continuous and collision free primitives of τ_k duration and initial state $\mathbf{s}_{d,k}$. If the end state of the primitive $\mathbf{s}_d(\tau_k)$ is state which collides with an obstacle, we modify it based on (17g). We modify the final state based on DRR controller model. We make the modified final state as initial state of next primitive. If the end state of the primitive $\mathbf{s}_d(\tau_k)$ is collision free, we keep the final state similar the collision-avoidance planner making the final state as initial state of the next primitive as (17g). Reformulating into a discrete control problem enables the use of several motion planning methods that rely on search-based [44] or sampling-based [4] techniques. We choose to adopt an A* technique similar to [21] and concentrate on the creation of effective, guaranteed collision checking and post-collision behavior categorizing methods, as well as an accurate and consistent heuristic since the former can ensure limited time (sub-)optimality.⁹

D. Collision Checking and Post-collision Behaviors

For a computed edge $e(t) = [\mathbf{p}(t)^\top \mathbf{v}(t)^\top \mathbf{a}(t)^\top \dots]^\top$, in Alg. 3, we need to check if $e(t) \in \mathcal{S}^{free}$ for all $t \in [0, \tau_f]$. For $e(t) \in \mathcal{S}^{free} \wedge e(t + \delta t) \in \mathcal{S}^{obs}$ with $\delta t \rightarrow 0$ for all $t \in [0, \tau_f]$, we need to modify the edge $e(t)$ as in lines 16–21 in Alg. 3. We check collisions in the geometric space $\mathcal{P}^{free} \subset \mathbb{R}^n$ separately from enforcing dynamic constraints $\mathcal{D}^{free} \subset \mathbb{R}^{n \times (q-1)}$. An edge $e(t)$ is collision-free only if its geometric shape $\mathbf{p}_e(t) \in \mathcal{P}^{free}$ for all $t \in [0, \tau_f]$.

In general, determining collision points for each motion primitive can be very challenging. Herein we model \mathcal{P} as an occupancy grid map, \mathcal{M}_o . Other representations such as polyhedral maps [3], [22], [25] are also possible but often hard to obtain from a robot's FoV sensor data (e.g., from LiDAR) and hence not pursued herein. Let $\mathcal{P}_e = \{\mathbf{p}_e(t_i) \mid t_i \in [0, \tau_f], i = 1, \dots, I\}$ be a set of positions that the system traverses along the trajectory. For collision-free primitives we need $\mathbf{p}_e(t_i) \in \mathcal{P}^{free}$ for all $i \in \{0, \dots, I\}$. The duration of the collision-free trajectory is $\tau = \tau_f$. For the given polynomial $\mathbf{p}_e(t)$, $t \in [0, \tau_f]$, the positions $\mathbf{p}_e(t_i)$ are sampled by defining

$$t_i = \frac{i}{I} \tau_f \quad \text{such that} \quad \frac{\tau}{I} v_{max} \geq \epsilon_{map}, \quad (18)$$

where ϵ_{map} is the occupancy grid resolution, and $v_{max} = \max\{|v_{min}|, |v_{max}|\}$. This condition ensures that the maximum distance between two consecutive samples will not

exceed the map resolution. Since it is an approximation, some cells traversed by $\mathbf{p}_e(t)$ with a portion of the curve within the cell shorter than ϵ_{map} may be missed, but it guarantees the collision-free trajectory does not hit any obstacles.

In not collision-free $e(t)$, the estimated collision time instant t_i is when $\mathbf{p}_e(t_i) \in \mathcal{P}^{free} \wedge \mathbf{p}_e(t_i + \delta t) \in \mathcal{P}^{obs}$ with $\delta t \approx t_{i+1} - t_i$ for all $i \in \{0, \dots, I-1\}$. Then, we set the duration τ of the collision-inclusive motion primitives in Alg. 3 to t_i , and modify the end state $\mathbf{s}_{d,e}$ of $e(t)$ as $\mathbf{s}_{d,e} = \{\mathbf{s}_{d,e}^-, \mathbf{s}_{d,e}^+\}$ with $\mathbf{s}_{d,e}^- = e(t_i) = [\mathbf{p}_e(t_i)^\top \mathbf{v}_e(t_i)^\top \mathbf{a}_e(t_i)^\top \dots]^\top$. We set the duration of this edge $\tau = t_i$ and set $\zeta(\tau) = 1$. $\mathbf{s}_{d,e}^+ = F_{DRR}(\mathbf{s}_{d,e}^-)$ is the post-impact state recovered using the DRR strategy. We discuss how to set $\mathbf{s}_{d,e}^+$ shortly.

Since \mathbf{v}, \mathbf{a} and other higher-order derivatives are polynomial functions, we can compute their extrema within the time period $[0, \tau]$ to check if the respective maximum bounds are violated. The polynomials' order is less than 5 for $n \leq 3$, hence the extrema can be computed quickly in closed form. We eliminate the primitives that cannot be dynamically implemented (i.e. any bounds are exceeded). For the collision-inclusive primitives, we need to check the x component of the velocity v_e^- in \mathcal{F}_c corresponding ${}^c v_{e,x}^- \in \mathcal{V}^c$. We prune those with ${}^c v_{e,x}^- \notin \mathcal{V}^c$ to prevent the robot from flipping over after colliding.

To generate the frame \mathcal{F}_c required for evaluating the collision-inclusive primitives, we need to get the geometric information of each obstacle that the robot collides on. Given a current scan from the mapping sensor (e.g., a LiDAR) we identify all possible collision surfaces and use regression to fit curve equations to the possible collision surfaces. The value of doing so is that it enables a rapid calculation of the tangent and normal unit vectors at selected possible collision points on those collision surfaces. Basis vectors of \mathcal{F}_c are generated as discussed in Sec. IV-A whereas the origin of \mathcal{F}_c is set to be the estimated position of collision \mathbf{p}_e^- in $\mathbf{s}_{d,e}^-$.

After generating \mathcal{F}_c , we are able to generate $\mathbf{s}_{d,e}^+$ based on the map \mathcal{M}_o which we predict the robot will collide on when arriving at $\mathbf{s}_{d,e}^-$ with the given motion primitive. Given the goal position \mathbf{p}_{goal} , we are able to set $\mathbf{s}_{d,e}^+$ according to Alg. 4. This way, we can ensure the trajectory generated by the search-based algorithm respect constraint (17g). In Fig. 6, we show how to generate the intermediate waypoint \mathbf{p}_{add} based on the jump point search algorithm. If there is no feasible path to the goal, we prune this collision-inclusive motion primitive. Given ${}^c v_e^-$ and ${}^c v_e^+$, we can generate J_c of this collision-inclusive motion primitive according to (12).¹⁰ We set a lower bound to J_c, \underline{J}_c , to induce a cost if the robot tries to use collisions alone to steer. Tuning ρ_c help regulate collision-avoiding and collision-inclusive trajectories.

We also create an infeasible, \mathcal{P}^{inf} , area to link pruned collision-inclusive $\mathbf{p}_e^- \in \mathcal{P}^{inf}$. We apply \mathcal{P}^{inf} to prevent the robot from getting into areas where the collisions are difficult to detect using this arm design (i.e. when collision

⁹ We note here that in principle both a search-based (as herein) and a sampling-based global planner is possible. In Section VII-D we demonstrate the differences of the two within collision-inclusive motion planning.

¹⁰ We consider that most of the collision energy can be recovered by the robot via its compliant arms. In practice, precise computation of the dissipated energy is a challenge; however, the DRR strategy accommodates for collision energy losses without any explicit energy dissipation models.




Fig. 6: An example of performing collision and detouring away from a non-convex obstacle by generating a new waypoint between the point of collision and the goal.

surfaces reduce to almost a point, such as obstacle corners).

Algorithm 4: Post-collision state generation

input : Pre-collision state in world frame $s_{d,e}^-$; collision frame \mathcal{F}_c ; goal position in world frame p_{goal} ; \mathcal{M}_o .

output : Post-collision state in world frame $s_{d,e}^+$; behavior type after collision ξ .

parameter: Lower bound and upper bound of the velocity v_{min}, v_{max} ; upper bound of duration of each primitives $\tau_f, \delta t$.

1 **Function** GETPOSTCOLLISIONSTATE ($s_{d,e}^-, p_{goal}, \mathcal{M}_o$):

2 $v^+ \leftarrow \frac{p_{goal} - p_e^-}{\tau_f}$

3 Generate ${}^c v^+$ based on \mathcal{F}_c as what is shown in Fig. 6

4 **if** ${}^c v_x^+ < 0$ **then**

5 ${}^c v_x^+ \leftarrow 0$

6 $\xi \leftarrow 2$

7 We generate an intermediate waypoint p_{add} as what is shown in Fig. 6 given \mathcal{M}_o based on jump point search algorithm. The waypoint should be the last visible waypoint along the path. $v^+ \leftarrow \frac{p_{add} - p_e^-}{\tau_f}$

8 **else**

9 $\xi \leftarrow 1$

10 $p_{add} \leftarrow \emptyset$

11 **end**

12 We adjust each component of v^+ with a saturation function restricting upper bound and lower bound as v_{min} and v_{max} .

13 We set all derivatives of $s_{d,e}^+$ as $p_e^{(q)+} = 0$, for all $q \geq 2$, $p_e^{(1)+} = v^+$, $p_e^+ = p_e^-$.

14 **return** $s_{d,e}^+, \xi$

E. Heuristic Function Design

A heuristic function that is admissible, informative (i.e. provides a tight approximation of the optimal cost), and consistent (i.e. it can be inflated to obtain solutions with bounded sub-optimality efficiently) is required for efficient graph search to solve (17). Similar to [21], we solve a relaxed form of (13) and arrive at a reasonable heuristic function. The basic concept is to replace the difficult-to-satisfy $s_d(t) \in \mathcal{S}^{free}$ and $u(t) \in \mathcal{U}$ requirements in (13) with a constraint on time T . Next, we demonstrate that a relaxation of (13) that includes motion planning may be solved optimally and effectively. We add a constraint to ensure that the robot will travel through the recently added waypoint p_{add} , avoiding the obstacle it collided with and preventing repeated collisions with it if $p_{add} \neq \emptyset$.

1) *Lower Bound of Time:* Limits on maximum speed, acceleration, jerk, etc. imposed by \mathcal{S}^{obs} and \mathcal{U} can help create a lower bound on T . If $p_{add} = \emptyset$, the minimum time to reach the nearest state $s_{d,goal}$ in the goal region \mathcal{S}^{goal} is constrained by $\underline{T}_v = \frac{\|p_{goal} - p_0\|_\infty}{v_{max}}$. This is because the system's maximum velocity is bounded by v_{max} along each axis. The system's maximum acceleration is bounded by a_{max} , hence the state $s_{d,goal} = [p_{goal}^\top v_{goal}^\top]^\top$ cannot be reached faster than

$$\min_{a, \underline{T}_a} \underline{T}_a \quad (19a)$$

$$\text{subject to } \dot{s}_d(t) = A_f s_d(t) + B_f u(t), u(t) = a(t), \quad \forall t \in [0, \underline{T}_a] \quad (19b)$$

$$\|a(t)\| \leq a_{max} \quad (19c)$$

$$s_d(0) = [p_0^\top v_0^\top]^\top, s_d(\underline{T}_a) = [p_{goal}^\top v_{goal}^\top]^\top \quad (19d)$$

The above is a minimum-time optimal control problem with input constraints, which can be solved in closed form along each individual axis to obtain the lower bound $\underline{T}_a = \min\{\underline{T}_{a,x}, \underline{T}_{a,y}, \underline{T}_{a,z}\}$ [45, ch. 5]. This procedure applies for constraints in higher-order derivatives, but in practice the computed times are less likely to provide better bounds while requiring higher computational effort. Hence, even though we can define a lower bound on the minimum achievable time via $\underline{T} = \min\{\underline{T}_v, \underline{T}_a, \dots\}$, for computational expediency we use the efficiently-computed (but less tight) bound $\underline{T} = \underline{T}_v$. For those cases with $p_{add} \neq \emptyset$, we generate \underline{T}_1 and \underline{T}_2 for path segments $p_0 \rightarrow p_{add}$ and $p_{add} \rightarrow p_{goal}$ as $\underline{T}_1 = \underline{T}_{v,1} = \frac{\|p_{add} - p_0\|_\infty}{v_{max}}$ and $\underline{T}_2 = \underline{T}_{v,2} = \frac{\|p_{goal} - p_{add}\|_\infty}{v_{max}}$.

2) *Velocity Control Linear Quadratic Minimum Time Heuristic:* The lower bound \underline{T} can help relax (13) by replacing the state and input constraints. If $p_{add} = \emptyset$, then

$$\min_{D, T_g} J_s(D) + \rho_t T_g \quad (20a)$$

$$\text{subject to } \dot{s}_d(t) = A_f s_d(t) + B_f u(t), \forall t \in [0, T_g] \quad (20b)$$

$$s_d(0) = s_{d0}, s_d(T_g) \in \mathcal{S}^{goal} \quad (20c)$$

$$T_g \geq \underline{T} \quad (20d)$$

The relaxed problem (20) is in fact the classical Linear Quadratic Minimum-Time Problem [46]. The optimal cost

generated from (20) according to [22] is

$$h(\mathbf{s}_{d,0}) = \delta_T^\top W_T^{-1} \delta_T + \rho_t T_g . \quad (21)$$

We define $\delta_T = \mathbf{s}_{d,goal} - e^{A_f T_g} \mathbf{s}_{d,0}$ and the controllability Gramian $W_T = \int_0^{T_g} e^{A_f t} B_f^\top e^{A_f^\top t} B_f dt$.

Let us consider velocity control as an illustrative example of (21). Given T_g , $\mathbf{s}_{d,0} = \mathbf{p}_0$, $\mathbf{s}_{d,goal} = \mathbf{p}_{goal}$, we can rewrite the optimal cost of (20) shown in (21) as

$$h_v(\mathbf{s}_{d,0}) = C^*(T_g) = \frac{\|\mathbf{p}_{goal} - \mathbf{p}_0\|^2}{T_g} + \rho_t T_g . \quad (22)$$

By minimizing C^* in (22) with the constraint $T_g^* \geq \underline{T}$, we are able to obtain the ideal T_g^* . If the positive real root $root^+ \geq \underline{T}$, then the solution is the positive real root of $\frac{dC^*}{dT_g} = 0$. Otherwise, $T_g^* = \underline{T}$. Furthermore, the optimal cost is $C^*(T_g^*)$. For the case where $\mathbf{p}_{add} \neq \emptyset$, we modify (22) to

$$h_v(\mathbf{s}_{d,0}) = C^*(T_1, T_2) = \frac{\|\mathbf{p}_{add} - \mathbf{p}_0\|^2}{T_1} + \rho_t T_1 + \frac{\|\mathbf{p}_{goal} - \mathbf{p}_{add}\|^2}{T_2} + \rho_t T_2 \quad (23)$$

Similarly, we are able to derive the optimal T_1^* and T_2^* by minimizing C^* in (23) with constraints $T_1^* \geq \underline{T}_1$ and $T_2^* \geq \underline{T}_2$. We can get the solution of this optimization problem by solving the positive real root of $\frac{\partial C^*}{\partial T_1} = 0$ and $\frac{\partial C^*}{\partial T_2} = 0$. The optimal cost then is $C^*(T_1^*, T_2^*)$.

F. Jump Point-based Computation to Improve Efficiency

Previous analyses (Sec. VI-A to VI-E), yield the overall structure of our proposed collision-inclusive search-based motion planning algorithm, based on A* graph search. From (17), we notice that we extend the feasible set of the optimization problem compared to the collision avoidance planning problem [21]. Extending the feasible set forces our method to traverse more nodes on the graph compared to the collision avoidance method. Even though our method can generate a less conservative result with less cost compared to the collision avoidance method, the computational time of our can be larger compared to collision avoidance.

To improve computational efficiency and reduce the computational time of our method, we can replace the post-impact motion primitive generation technique introduced in Sec. VI-D in A* graph search with a more efficient variant that is inspired by jump point search [47]. Specifically, we notice that when the robot needs to add a new waypoint between the collision point \mathbf{p}_e^- of the motion primitive and the goal \mathbf{p}_{goal} ($\xi = 2$), we can modify $\mathbf{p}_e^+ = \mathbf{p}_{add}$. Performing this modification will help us eliminate traversing multiple nodes with the same \mathbf{p}_{add} . This way, the number of nodes we are traversing can reduce, thus reducing computational time. Even though applying this technique can be at expense of optimality of the solution, solving the planning problem with less computational time can be more important in practice.

If colliding with an obstacle (as shown in Fig. 6), we modify \mathbf{p}_e^+ and duration τ as $\mathbf{p}_e^+ \leftarrow \mathbf{p}_{add}$ and $\tau \leftarrow \tau + \tau_{add}$ with $\tau_{add} = \frac{\|\mathbf{p}_{add,y} - \mathbf{p}_e^-\|}{\|v_e^x\|}$. The cost will be updated with new

$\tau \leftarrow \tau + \tau_{add}$. When we go through edges with $\xi = 2$, we split the trajectory of this given edge with two segments, given the start and the end waypoints as \mathbf{p}_0 and \mathbf{p}_e^- for the first segment, \mathbf{p}_e^- and \mathbf{p}_e^+ for the second segment. The time duration of the first segment is $\tau - \tau_{add}$ and the time duration of the second segment is τ_{add} . We set the $\xi = 0$ and $\zeta = 0$ with respect to the waypoint \mathbf{p}_e^+ .

G. Trajectory Refinement

Following the aforementioned approach results in a collision-free trajectory including specific times needed to reach each waypoint. This is then fed as a prior to create smooth trajectories in higher dimensions. The refined trajectory $\mathbf{s}_d^*(t)$ is derived from solving a gradient-based trajectory generation problem similar to the one in Sec. V-A with given initial and end states $\mathbf{s}_{d,0}$ and $\mathbf{s}_{d,goal}$ and intermediate waypoints $\mathbf{p}_k, k \in \{0, 1, \dots, N\}$. The time for each trajectory segment τ_k is also given from the prior trajectory. All \mathbf{p}_k are stored in `waypoint_list` (see Alg. 2).

We apply a two-step optimization strategy similar to [42] which can be summarized as follows: 1) First, optimize the collision cost of the path generated from waypoints only. Positions of intermediate waypoints on the initial path are left as free variables, and will be pushed away from the obstacles. 2) Second, revise the time scaling of the trajectory according to current waypoints' positions. Then optimize the objective with additional smoothness and dynamical penalty terms.

The output trajectory comprises $N_c + 1$ continuous polynomial trajectory segments. The differential variable of the waypoint with $\xi_k \geq 1$ is fixed end variable in the collision-inclusive method, which indicates $\mathbf{s}_{d,ed,i} = \mathbf{s}_{d,k}^-, k \in \{0, 1, \dots, N\}, i \in \{1, \dots, N_c\}, \forall \xi_k \geq 1$. The differential variable of the next fixed initial state is $\mathbf{s}_{d,st,i+1} = \mathbf{s}_{d,k}^+, k \in \{0, 1, \dots, N\}, i \in \{1, \dots, N_c\}, \forall \xi_k \geq 1$. The collision state waypoints $\xi_k \geq 1$ are generated from a grid map with augmented obstacles. We need to adjust those waypoints before trajectory generation by relocating them so that the distance to the closest block is $d_c \leq r_{rob}$; this way we ensure that planned collisions occur. The output of the search-based algorithm may have two colliding and reflecting states that are close-by (Fig. 7). In this case, if \mathbf{p}_i is visible to both \mathbf{p}_{i-2} and \mathbf{p}_{i+1} , we can delete \mathbf{p}_{i-1} to reduce redundant collisions. We also disregard obstacles that the robot planned to collide on ($\xi_k \geq 1$) when computing the potential field for trajectory generation in the second step for computational expediency. The trajectory after refinement is n -th order continuous.

It is important to note that even though the refinement step produces a smoother trajectory, the refined course might be dynamically infeasible; we need to perform time scaling as in [22] to reduce the maximum dynamics of the refined trajectory. The refined trajectory might collide with the obstacle in the trajectory segment that is checked to be collision free according to Sec. VI-D. In such cases, our DRR strategy ensures robustness and safety.




Fig. 7: Collision state merging.

VII. EXPERIMENTAL RESULTS

We validate the effectiveness of our unified framework for collision-inclusive motion planning and control by presenting several benchmark testing results in simulation and via physical experimentation with our robot. 1) First, we test the deformation controller on the robot to evaluate its performance and generate a post-collision model which is required in simulation. 2) Second, we test the local DRR trajectory generation component experimentally with our robot. 3) Then, we test our global planning method in a double corridor environment and compare it with state-of-the-art search-based collision avoidance and sampling-based collision-inclusive methods. 4) After that we test the overall planning strategy in simulation in unknown, partially-observable environments. 5) Lastly, we evaluate the overall method with our impact-resilient robot experimentally.

A. Experimental Setup and Implementation Details

Testing the deformation controller (Sec. VII-B) and DRR (Sec. VII-C) experimentally takes place in a 2.0×2.5 m area with a rectangular pillar serving as a static polygon-shaped obstacle. The overall method is tested experimentally (Sec. VII-F) in a 2.5×3.5 m area with a long rectangular pillar in the middle to form a U-shaped corridor environment.

We use the two active omni-directional impact-resilient wheeled robots we built in-house (Fig. 1). The main chassis is connected to a deflection ‘ring’ via 4 or 8 arms that feature a visco-elastic prismatic joint each. Each arm has embedded Hall effect sensors to measure the length of the arm and detect collisions along each of their direction when the deformation exceeds a certain threshold. In physical experiments, odometry feedback is provided by a 12-camera VICON motion capture system. An onboard Intel NUC mini PC (2.3 GHz i7 CPU; 16 GB RAM) processes odometry data and sends control commands to the robot at 10 Hz. The robot is equipped with a single-beam LiDAR (RPLidar A2) with 8 m range to detect the obstacles in the environment.

The robot may flip when colliding with velocities over a bound. To identify a theoretical collision velocity bound to avoid flipping, we use an energy conservation argument. Assume the kinetic energy before collision transfers into elastic potential energy of the arm, and the gravitational potential energy of the robot with small flipping angle counters the negative work input from the controller, i.e.

$$E_{k,t}(v_{max}) = E_{ep}(l_e) - E_{ep}(l_s) + E_{gp}(\sigma_{max}) + m a_{in,max}(l_s - l_e)$$

Then,

$$v_{max} = \left\{ \frac{k[(l_e - l_0)^2 - (l_s - l_0)^2]}{2m} + g(\rho - l_s + l_e) \sin \sigma_{max} + a_{in,max}(l_s - l_e) \right\}^{\frac{1}{2}}.$$

The robot’s radius is 0.3 m. The difference between the initial and neutral position of each arm is $l_s = 30.0$ mm, the maximum load length is $l_e = 15.0$ mm, and the neutral length is $l_0 = 41.5$ mm. The spring coefficient is $k = 2.31$ N/mm. We set the largest flip angle $\sigma_{max} = 3^\circ$. For the 4-arm robot, the maximum acceleration input is $a_{in,max} = 5.0$ m/s², and its mass is 6.0 kg. Then, we compute an upper theoretical velocity bound of $v_{max} \approx 0.7$ m/s.¹¹

Simulated comparison against other methods (Sec. VII-D) takes place in a double-corridor environment, whereas simulated benchmark testing of our method when noise is added takes place in the same double-corridor environment but with added isolated obstacles added as well (Sec. VII-E). We further consider a similar environment that features non-convex obstacles (Sec. VII-E).

We use a rigid cylinder body to emulate the robot. A numerical model is generated from the experiments for the deformation recovery controller to determine the output velocity after collision. The output velocity is generated by adding uniform random noise to the reference velocity ${}^c v_T$. Then, we use a ray-casting algorithm to emulate the LiDAR (we consider the range of the LiDAR can cover all visible operating space). We implement simulation benchmarks in a python environment. All simulations run on a workstation with Intel Core Xeon-E2146G CPU.

B. Experimental Testing of the Deformation Controller

To examine the deformation controller’s effect in local trajectory generation, we command the robot to collide with an obstacle and then apply the proposed deformation recovery controller. We perform 10 trials of various input-output velocity combinations [20, Table 2]. Collision detection is very accurate; only 9 out of 249 were not detected.

Results suggest that the deformation controller generates a negative velocity to make the robot detach from the obstacle after collision. Actual output velocity \bar{v}_{out} is determined by the actual input velocity \bar{v}_{in} and the set output value $v_{out,set}$ though the latter may not be reached in practice. That is because feedback linearization is not robust to system parameter uncertainties that occur in practice. We observe that the velocity along the normal to collision direction is closer to the set velocity than the velocity along the tangent direction. This is because most of the uncertainties in system parameters enter as unmodeled friction dynamics along the tangent direction. Further, the sensor is more accurate when the input velocity is along the normal direction; the average value of deformation detected in this case is 29% larger.

¹¹The 8-arm robot features motors with higher torque and different gear ratio that increase $a_{in,max}$ and despite the mass increase to 8 kg, the same upper theoretical velocity bound remains valid.

C. Experimental Testing of the DRR Strategy

We test our DRR strategy with a trajectory generated based on using the online safe trajectory generation method in [48] with time scaling as in [22] without collision checking. We compare the strategies in two cases: 1) when the previous path does not intersect with the collision surface; and 2) when the previous path intersects with the collision surface.

Case 1 tests the condition in Fig. 5a, i.e. no waypoint is added as per Alg. 2. Case 2 tests Fig. 5, i.e. a waypoint is added to the list. In case 2, we run RRT* to generate a collision free path and perform path simplification to remove nodes without affecting the path’s collision safety [20, Fig. 4]. The path simplification technique removes intermediate waypoints between two waypoints if a line segment between those two does not intersect with the obstacle. Then, we use the trajectory generation strategy in [48]. We perform 10 trials for each case. Instances of DRR and all experimental trajectories are shown in Fig. 8 and Fig. 9, respectively.



Fig. 8: Composite images (taken every 2 s) of a sample experiment with DRR. The robot must go from start to goal passing through all waypoints, including intermediate ones created post-collision.

Even though we design a collision-free desired trajectory with the strategy in [48], the robot may still collide with the environment given for instance unmodeled dynamics such as drift. In case 1 there are 3 out of 10 trials that the robot in fact collides with the obstacle applying trajectory generation [48] that aims to avoid collisions. Table II shows statistics on mean arrival times, path lengths and control energy.

TABLE II: Comparison of trajectory generation strategy in [48] (Collision-avoidance) and DRR (Collision-inclusive) strategies.

	Strategy in [48]		DRR (our method)	
	Case 1	Case 2	Case 1	Case 2
\bar{T}_{end} [s]	7.7	8.71	6.16	9.17
$STD(T_{end})$	0	2.20	0.22	0.31
\bar{s} [m]	3.153	3.448	2.977	4.38
$STD(s)$	0.117	0.495	0.150	0.451
\bar{E}_c [m^2/s^3]	56.83	80.62	58.42	255.96
$STD(E_c)$	29.34	54.38	32.92	133.54

In case 1 for DRR, mean arrival times \bar{T}_{end} and path lengths \bar{s} decrease by 25% and 6%, while the control energy increases by 2.8% on average. However, the error in the end point increases by 25%. In case 2, mean arrival times and path lengths increase by 5.2% and 27%, and control energy increases by 258%. This is because the output velocity of DRR is not flat since the robot decelerates and then accelerates during boundary following. The path generated by the boundary following is not the shortest. However, since the path between the collision point and the new inserted

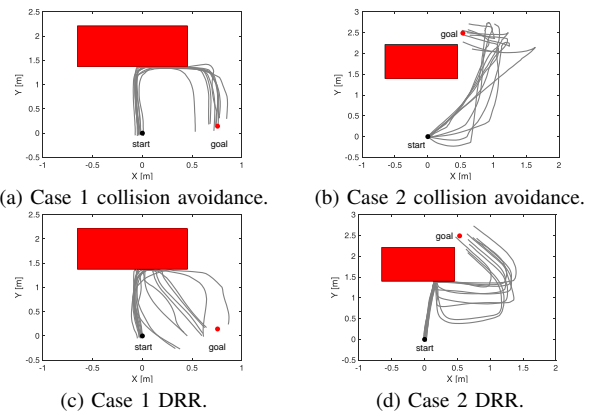


Fig. 9: Experimental trajectories generated from DRR and collision avoidance [48] when the preplanned path intersects or does not intersect with the obstacle. (In all cases we conduct 10 trials.)

waypoint is close to the obstacle surface, the existence of the obstacle decreases the control error in free space. The error in the end point decreases by 12%. These results show the tradeoff between online reactive execution (whereby collision checking is skipped) and collision avoidance.

D. Simulated Tests of the Collision-inclusive Global Planner

To test our proposed framework in simulation, we first benchmark it in a double corridor environment to test our search-based collision-inclusive global planner. We compare our method for global planning against two methods: 1) a search-based collision-avoidance motion planning algorithm [21], and 2) an RRT*-based (sampling-based) collision-inclusive planning algorithm adapted from [4] to ensure fair comparisons.¹² In all tests, the dynamic limits are set as $a_{max} = 5.0$ m/s². The holonomic robot only translates but does not rotate during the process.¹³ We set the upper bound of the robot velocity $v_{max} = 2.0$ m/s. The cost function in all methods considers $\rho_t = 1.0$. The overall size of the map is 70×70 m. The position resolution of the grid map in the benchmark is 1.0 m, and the position resolution of the velocity map in the benchmark is 0.1 m/s. The time interval for each motion primitive is set as $\tau = 5.0$ s (selected via an ablation study the results of which are shown in Table III) and the resolution of acceleration $r = 1.0$ m/s² also selected via an ablation study the results of which are shown in Table IV). We set $\lambda_s = 0.5$, $\lambda_o = 1.0$ and $\lambda_v = 10.0$.¹⁴ Parameter ρ_c in the cost function is one of the most critical ones as it determines how much collisions are being penalized. We select $\rho_c = 1.0$, with the ablation studies to determine this value being demonstrated qualitatively in Fig. 10 and expanded in more depth in Table V.

Results from testing the global planner are shown in Fig. 10; both collision avoidance as per [21] (panel a) and collision-inclusive (our method, panel b-f) results are

¹² No open-source python code is available for either [21] and [4], so we implemented both ourselves to the best possible extent.

¹³ Constant orientation is maintained via a separate stabilizing controller.

¹⁴ The values were selected empirically to improve trajectory refinement.

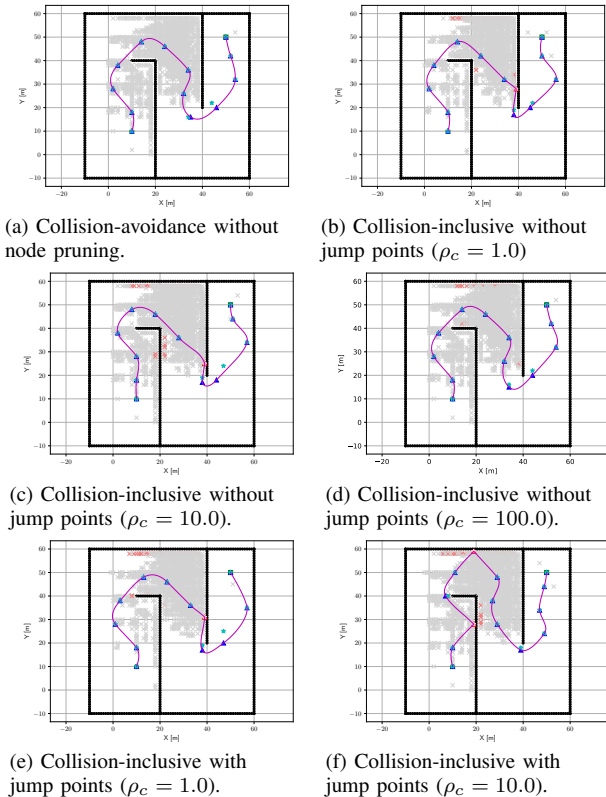


Fig. 10: Search-based collision avoidance global planner [21] (top panels) and our proposed search-based collision-inclusive global planner for three different ρ_c values. Higher ρ_c values (bottom panels) penalize collisions more, thus recovering behaviors that resemble collision avoidance. The collision-inclusive with jump points case when $\rho_c = 100.0$ is very similar to the case without jump points in panel (d) in the sense of number of closed nodes N_p (c.f. Table V, and hence not shown here for brevity).

highlighted. We demonstrate our method’s results with and without implementing jump points. We also consider three cases for varying values of parameter ρ_c which affects how much collisions are being penalized in the cost function: 1) $\rho_c = 1.0$ (panels b and e) which corresponds to minimal penalty; 2) $\rho_c = 10.0$ (panels c and f) which corresponds to a medium penalty; and 3) $\rho_c = 100.0$ (panels d) which corresponds to a severe penalty. It can be readily verified that both cases of $\rho_c = \{1.0, 10.0\}$ can lead to paths that contain collisions, although in some cases (especially when jump points are considered) a higher ρ_c value of 10.0 may make the output trajectory unnecessarily complex and suboptimal (panel f). As such, if collisions are to be considered, setting $\rho_c = 1.0$ should be preferred. At very high ρ_c values (of 100.0), we observe that our method can recover collision avoidance behaviors (c.f. panels a and d). This highlights our global planner’s ability to switch between collision-inclusive and collision avoidance planning on-demand by only updating the value of a single parameter.

We also demonstrate the utility of formulating the motion primitives as discussed in Sec. VI-B as compared to directly pruning dynamically infeasible primitives. Our proposed method can feature primitives of longer duration τ , which in fact helps increase the efficiency of exploring the space.

TABLE III: Comparison of global planners’ performance between our method and that to prune primitives.

$r = 1.0$ [m/s ²]	τ [s]	Comp. Time[s]	N_p	Traj. Time[s]	Ctrl. Cost[m ² /s ³]
method [21] without pruning	5.0	31.86	1427	71.5	12.01
Our method without pruning $\rho_c = 1.0$	5.0	34.60	1340	69.1	13.25
method [21] with pruning	0.5	18629.12	102836	91.1	201.98
Our method with pruning $\rho_c = 1.0$	0.5	15168.97	89450	90.9	200.72

Table III provides comparative numbers for both collision avoidance [21] and our collision-inclusive method. Results verify that our selected primitives generation method can explore the space with significantly less computational time when compared to the approach of pruning infeasible states.

Table IV contains the results from the ablation study on the resolution parameter r . We found that in the environment with simple obstacles (as in Fig. 14a shown later), a high resolution of $r = 0.5$ leads to computational times for both collision-avoidance and collision-inclusive methods that are much higher since the graph is denser. The computational time of the collision-inclusive method is higher than the collision-avoidance method since we modify those primitives that intersect with the obstacles instead of pruning them altogether. With lower resolution $r = 1.0$ or $r = 2.0$, the computational time of both collision-avoidance and collision-inclusive methods rapidly decreases. When $r = 1.0$ collision-avoidance and our collision-inclusive method have comparable computational performance. However, as we further increase the resolution ($r = 2.0$), the computational time of our collision-inclusive method gets much lower than the collision-avoidance method, which appears to be affected less by this change. These observations suggest that with lower resolution, collision-inclusive primitives can explore the space with higher efficiency. However, the trajectory time and control cost are higher than applying higher resolution. Similar observations can be made when testing in a more complicated environment (Fig. 14b). Therefore, taken computational time, trajectory time and control cost into consideration, we select $r = 1.0$ for which both collision-avoidance and collision-inclusive method have better results.

Further, we conduct a more extensive analysis to evaluate the effect of different values of parameter ρ_c in more detail. Table V contains more detailed results and also presents comparisons against the sampling-based (RRT*) method in [4], which was adapted herein to feature a trapezoidal velocity pattern to connect any two nodes in the tree to better match our search-based global planning method and enable fair comparisons. Due to the non-deterministic nature of this method, we perform 10 trials and report statistics. In all other cases (that are deterministic), we perform a single trial.

Both our collision-inclusive method and collision avoidance in [21] can generate kinodynamically-feasible trajectories. When the initial velocity is $v_0 = [0 \ 0]^T$ m/s and $\rho_c = 1.0$, our method without jump points tends to gen-

TABLE IV: Comparison of global planners' performance for different values of parameter r

non-convex env. Fig. 14a	r [m/s ²]	Comp. Time[s]	N_p	Traj. Time[s]	Ctrl. Cost[m ² /s ³]
method [21]	0.5	480.19	6752	63.1	11.86
Our method without jump point $\rho_c = 1.0$	0.5	873.78	9311	66.3	14.09
method [21]	1.0	31.86	1427	71.5	12.01
Our method without jump point $\rho_c = 1.0$	1.0	34.60	1340	69.1	13.25
method [21]	2.0	34.68	1657	74.6	17.48
Our method without jump point $\rho_c = 1.0$	2.0	9.70	557	74.7	12.57
<hr/>					
non-convex env. Fig. 14b	r [m/s ²]	Comp. Time[s]	N_p	Traj. Time[s]	Ctrl. Cost[m ² /s ³]
method [21]	0.5	564.95	8385	74.9	9.58
Our method without jump point $\rho_c = 1.0$	0.5	1995.16	15908	81.3	11.94
method [21]	1.0	55.11	2388	79.9	12.43
Our method without jump point $\rho_c = 1.0$	1.0	52.69	1799	76.6	20.15
method [21]	2.0	48.33	2151	89.0	20.84
Our method without jump point $\rho_c = 1.0$	2.0	24.39	1726	79.6	12.33

erate a path with the shortest duration compared to both the collision-inclusive planner with jump points and the collision-avoidance planner. However, the control cost for doing so is slightly higher. The computational time of the collision-inclusive planner with jump points is the second-lowest among all the methods. Comparing these results with those obtained by the RRT* method in [4], we find the RRT*-based approach is time consuming since node rewiring requires significant computational time (about 88% of total time). Also, results are not deterministic compared to the search-based method. Thus, we deduce that the search-based collision-inclusive method with jump points can be the global planner in our unified collision-inclusive motion planning and control framework with parameters selected in this section.

E. Simulated Tests of our Unified Collision-inclusive Method

We first test our unified collision-inclusive motion planning and control strategy in a double corridor environment with online sensing, and compare its performance against that of a collision avoidance framework similar to [24]. In the collision avoidance framework, the global planner is the search-based method in [21]; we also make the optimistic assumption of treating the unknown space as free. The local trajectory generation method is based on gradients [42] and time duration adjustment [22]. We design a backup safety trajectory to ensure the robot will stop at the frontier. Then, we test both methods in a double corridor environment populated with circular isolated obstacles of increasing density. In all

TABLE V: Comparison of global planners' performance for different values of parameter ρ_c .

$v_0 = [0 \ 0]^T$ [m/s]	Comp. Time[s]	N_p	Traj. Time[s]	Ctrl. Cost[m ² /s ³]	Succ. Rate[%]	
search based method [21]	29.02	1427	71.5	12.01	100.0	
Our method no jump point $\rho_c = 1.0$	34.60	1340	69.1	13.25	100.0	
Our method with jump point $\rho_c = 1.0$	30.75	1388	73.3	10.101	100.0	
Our method no jump point $\rho_c = 10.0$	48.28	1701	74.1	10.74	100.0	
Our method with jump point $\rho_c = 10.0$	44.32	1747	76.7	12.60	100.0	
Our method no jump point $\rho_c = 100.0$	44.93	1588	71.5	12.01	100.0	
Our method with jump point $\rho_c = 100.0$	41.74	1602	71.5	12.01	100.0	
sampling based method [4]	mean	124.94	211	101.0	9.86	70.0
	std	55.63	38.2	9.86	2.91	
	min	93.72	183	84.2	5.77	
	max	259.40	298	114.8	15.27	
<hr/>						
$v_0 = [2 \ 2]^T$	Comp. Time[s]	N_p	Traj. Time[s]	Ctrl. Cost[m ² /s ³]	Succ. Rate[%]	
search based method [21]	22.41	1113	71.7	14.02	100.0	
Our method no jump point $\rho_c = 1.0$	24.45	1012	81.0	15.03	100.0	
Our method with jump point $\rho_c = 1.0$	23.78	1077	78.9	8.20	100.0	
Our method no jump point $\rho_c = 10.0$	37.30	1388	74.9	13.28	100.0	
Our method with jump point $\rho_c = 10.0$	36.67	1417	74.9	13.28	100.0	
Our method no jump point $\rho_c = 100.0$	30.97	1145	71.7	14.02	100.0	
Our method with jump point $\rho_c = 100.0$	30.38	1145	71.7	14.02	100.0	
sampling based method [4]	mean	88.50	169	117.0	14.18	70.0
	std	43.88	55.7	10.58	5.14	
	min	17.09	63	103.6	8.28	
	max	142.83	235	134.9	21.83	

cases, each method is run for 10 times with the same initial configuration and parameter settings.

We test with and without additive estimation noise in the global planner. Position estimation noise is zero-mean truncated Gaussian with variance of 0.3 and bounds of ± 0.9 . Velocity estimation noise is zero-mean truncated Gaussian with variance of 0.1 and bounds of ± 0.3 . Comparative results are presented in Figs. 11 and 12. Output trajectories of our method (with added noise) are shown in Fig. 13.

With reference to Fig. 11, when replanning every 5 sec, our method generates shorter paths with a lower trajectory time on average. When the obstacle density is low (9.3%), our method generates trajectories with higher control energy; however, when the obstacle density increases (13.5% & 20.7%) our method requires lower control energy since the

robot can utilize the obstacles to change its heading.¹⁵ When replanning every 10 sec, our method consistently generates paths with lower length on average. Our method also has lower trajectory times and control energy.

With $\rho_c = 100$, $T_{rep} = 5.0$ s, our algorithm causes oscillations around the obstacle by avoiding collisions which increase path length and trajectory time. However, its ability to use collisions makes its output trajectory better than the collision-avoidance strategy in terms of control energy, trajectory and path length. With $T_{rep} = 10.0$ s the path length, trajectory time and control energy of the collision-inclusive trajectory is higher since there is no safety maneuver making the robot stop before the frontier; hence it will have to turn sharply and possibly oscillate when replanning. The computational time of the collision-inclusive planning is higher since it visits more nodes in the graph.

Further, success rates of the proposed collision-inclusive method and collision avoidance are shown in Fig. 12. Our method has higher success rates as it addresses over-conservativeness in collision avoidance to ensure safety.

We then test our proposed unified collision-inclusive motion planning and control strategy in environments with non-convex obstacles (Table VI and Fig. 14). With reference to the environment shown in Fig. 14a, and using a replanning time interval $T_{rep} = 5.0$ s, our strategy can reach the goal with lower computational time, trajectory time and path length but higher control energy since the robot trajectories oscillate at the corner of the obstacle after collision-recovery and detouring. Similar patterns are observed for $T_{rep} = 10.0$ s. In the case shown in Fig. 14b, and with $T_{rep} = 5.0$ s, our strategy is better than collision avoidance in terms of computational time, trajectory time, path length and control energy. When $T_{rep} = 10.0$ s, the control energy of our method increases because of oscillations around corners.

TABLE VI: Results by testing with non-convex obstacles.

non-convex env. Fig 14a		Comp. Time[s]	Traj. Time[s]	Ctrl. Cost[m ² /s ³]	Path Len.[m]
Our method	mean	35.90	81.34	34.24	144.29
$T_{rep} = 5.0$	std	0.953	0.398	0.434	0.231
Collision-avoidance $T_{rep} = 5.0$		52.55	85.0	25.814	160.67
Our method	mean	16.71	80.43	21.15	146.51
$T_{rep} = 10.0$	std	0.243	0.857	0.312	0.121
Collision-avoidance $T_{rep} = 10.0$		18.62	90.0	15.11	167.25
non-convex env. Fig 14b		Comp. Time[s]	Traj. Time[s]	Ctrl. Cost[m ² /s ³]	Path Len.[m]
Our method	mean	9.09	80.48	19.72	146.02
$T_{rep} = 5.0$	std	0.210	0.623	2.205	2.543
Collision-avoidance $T_{rep} = 5.0$		32.62	93.4	27.38	173.60
Our method	mean	6.98	85.33	32.54	153.17
$T_{rep} = 10.0$	std	0.199	0.857	0.556	0.395
Collision-avoidance $T_{rep} = 10.0$		14.45	94.4	19.50	166.94

¹⁵ In collision avoidance, and with $T_{rep} = 5.0$ s, the robot can get trapped oscillating in a area to avoid collisions and cannot reach the goal; however, adding some random behavior may help the robot break the tie.

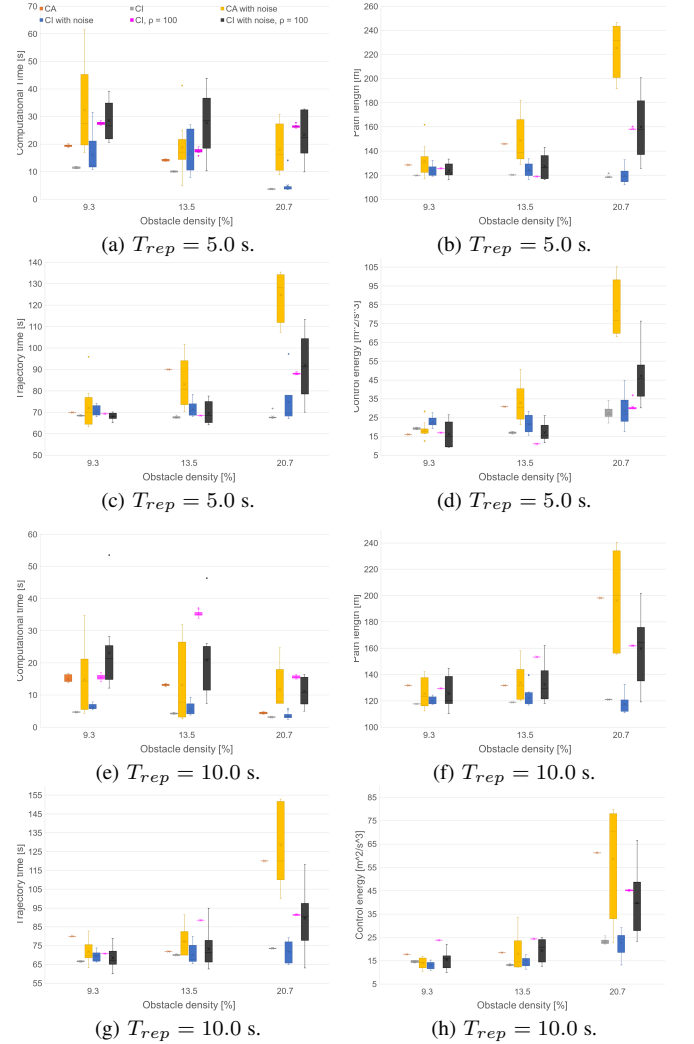


Fig. 11: Computational time, path length, trajectory time and control energy for collision-avoidance (CA) without (in red) and with (in yellow) noise, collision-inclusive (CI) without (in grey) and with (in blue) noise with $\rho_c = 1.0$ and collision-inclusive (CI) without (in magenta) and with (in black) noise with $\rho_c = 100.0$ methods, for two cases of replanning time $T_{rep} = 5.0$ s and $T_{rep} = 10.0$ s.

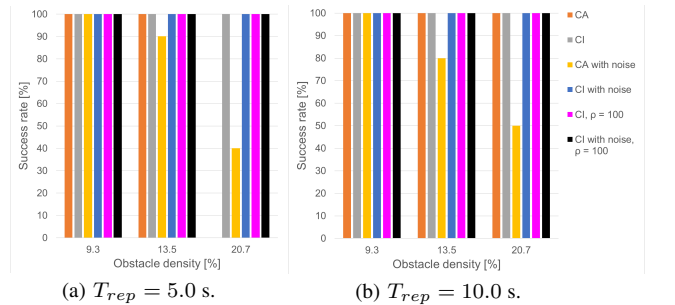


Fig. 12: Success rates of both collision-avoidance and collision-inclusive frameworks (same notation as in Fig. 11).

F. Experimental Validation of our Framework

Finally, we validate our proposed framework experimentally, and also test it against the collision avoidance strategy in Sec. VII-E, in a single corridor environment similar to [3] (Fig. 15). Each method is repeated for 10 times using the same parameter settings. Output trajectories are depicted




Fig. 13: Simulated trajectories of the unified collision-inclusive motion planning and control framework with online sensing and noise added to the input of the global planner in the double corridor for environment with increasing obstacle density.

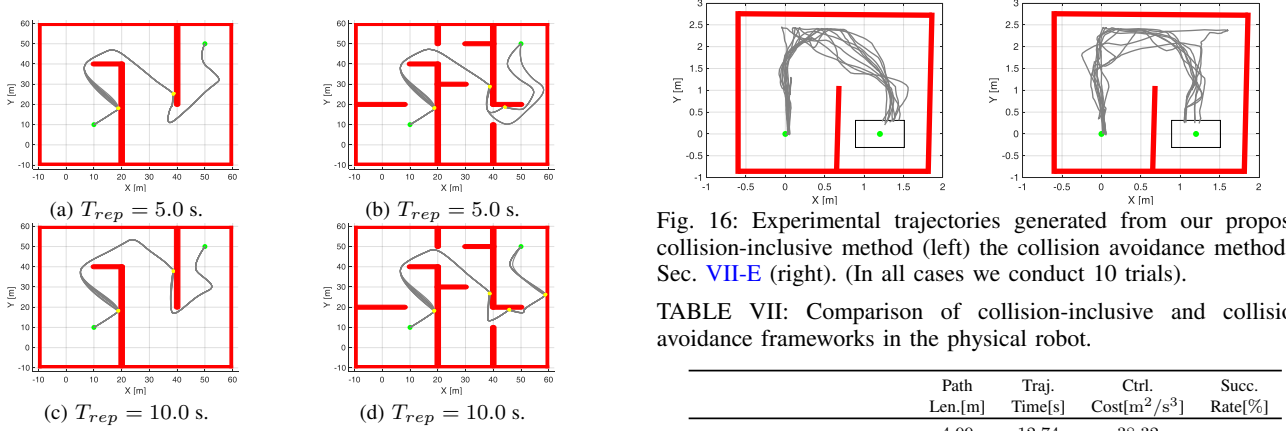


Fig. 14: Simulated trajectories generated from our method in environment with increasing obstacle density.



Fig. 15: Composite images of a sample experiment with our unified collision-inclusive motion planning and control framework (left) and of a collision avoidance sample experiment (right). Snapshots shown every 2.5 s. (See supplementary video file for more details.)

in Fig. 16, while detailed numerical results are given in Table VII. By implementing our proposed collision-inclusive planning method, the robot can reach the goal area with higher success rates since unmodeled dynamics in physical testing make the robot collide with the obstacle even if the reference trajectory generated from collision avoidance is designed to be collision-free. Further, by utilizing collisions, the robot can reach the goal faster while requiring less control energy by trading off the average path length.

VIII. DISCUSSION AND CONCLUSIONS

A. Summary of Contributions and Main Findings

In this article, we proposed a unified collision-inclusive motion planning and control framework applied for nav-




Fig. 16: Experimental trajectories generated from our proposed collision-inclusive method (left) the collision avoidance method in Sec. VII-E (right). (In all cases we conduct 10 trials.)

TABLE VII: Comparison of collision-inclusive and collision-avoidance frameworks in the physical robot.

		Path Len.[m]	Traj. Time[s]	Ctrl. Cost[m ² /s ³]	Succ. Rate[%]
Collision-avoidance	mean	4.99	12.74	38.32	90.0
	std	1.80	2.61	16.75	
Collision-inclusive	mean	5.36	11.98	32.45	100.0
	std	0.33	1.19	10.05	

igation in unknown environment. A global search-based method is devised to generate a path which contains explicit information about collisions. The effect of the collisions is explored in the global planner. The local planner is enhanced by a lower-level deformation recovery control and trajectory replanning strategy, which enables the robot to detect and recover from collisions and move toward the goal. The deformation controller is designed based on robot dynamics, which herein is a holonomic omni-directional wheeled robot.

The planning system was evaluated extensively through several benchmark comparisons in simulation as well as via physical experimental testing. The conducted ablation study demonstrated the utility of certain key design choices made in this work (e.g., not pruning primitives altogether), and evaluated the effect of key parameters (e.g., how much collisions are to be penalized via parameter ρ_c). The proposed collision-inclusive planning method is implemented in simulation first and then integrated with state estimation, mapping and control into our custom-made robot platform to check the feasibility of the method in physical world experiments. Results show that the proposed method is robust and can generate fast and safe trajectories compared to collision-

avoidance methods. Overall, this work pushes forward the state-of-the-art in collision-inclusive motion planning and control, and provides a competitive alternative to traditional collision avoidance methods for a class of impact-resilient mobile robots operating in partially-observable environments populated with isolated (non-)convex obstacles.

B. Discussion of Key Selections in our Framework

Application to Other Robots in 2D and 3D: We considered the family of omni-directional wheeled robots (Fig. 1). Yet, we anticipate that our proposed framework can apply to other impact-resilient robots in 2D (e.g., wheeled [11] or aerial [13], [14], [18], [37], [49] robots) provided that they can adjust their position and redirect post impact by using the collision to save energy. The omni-directional wheeled robot employed here is one example along those types of robots. The higher-level part of the framework can readily apply in 3D for such systems; same holds for the overall methodology as in whole. However, the proposed lower-level collision recovery would need to be adjusted to consider the 3D dynamics for post-impact stabilization [14].

Use of Motion Primitives: Besides the use of motion primitives (as herein) other methods are possible. For instance, direct control of the kinematic model (1) of the specific robot considered herein, or use of fixed motion patterns (e.g., as in [38]–[40]) can be viable alternatives. However, use of motion primitives at the higher-level provides a unified way to make the proposed framework applicable to all the aforementioned types of robots and extendable from 2D to 3D, and hence it was preferred to over simpler approaches that would have worked specifically for the omni-directional robot we tested with herein but would be hard to scale to other types of robots. Furthermore, use of primitives allows for more flexibility which is critical to help determine where the robot should collide with the environment to help it redirect toward the globally-planned goal; this is achieved by directly using information on the velocity as per (12).

Choosing a Search- or Sampling-based Global Planner: We showed that it is possible to derive collision-inclusive planning frameworks with the global planner being either search-based or sampling-based. Each has its own strengths and weaknesses, and as a matter of fact, our results are consistent with observations made in collision-avoidance methods. Consistent with collision avoidance, a user can choose which approach to select (search-based over sampling-based global planner) according to their application needs; our proposed framework can accommodate both. We highlight here that the sampling-based global planner can be further optimized by biasing search toward free space to increase computational efficiency (e.g., [50]). Integration of the sampling-based planner into our overall real-time framework would require further adaptations of the collision-inclusive RRT* planner to make it online (faster nearest neighbor search, minimal cost path generation, and optimized rewiring methods). Similar to collision-avoidance online RRT* methods (e.g., [51], [52]), a collision-inclusive

anytime planning algorithm is required to extend the RRT* method for planning collision-inclusive trajectories online.

C. Directions for Future Work

The framework developed herein lays the basis toward a general method for collision-inclusive motion planning and control, and creates multiple opportunities for future research along these lines. These include extension to other robots and to systems with higher-order dynamics, evaluation of direct controllers against motion primitives (as well as different parameterizations of the latter), and integration of sampling-based planners into the overall framework.

Further, at its current form, our method does not consider the perception model of the robot in online planning; extension of the proposed algorithm to consider the perception problem based on the collision-inclusive method is another interesting direction of future research. Lastly, we have shown that it is possible to handle navigation in environments populated with isolated non-convex environments; however, study of navigation in more cluttered environments (e.g., maze-like) is a direction of research enabled by this work.

ACKNOWLEDGEMENT

The authors wish to thank Hanzhe Teng for his help implementing the mapping package used in this work.

REFERENCES

- [1] Z. Lu and K. Karydis, "Optimal steering of stochastic mobile robots that undergo collisions with their environment," in *IEEE International Conference on Robotics and Biomimetics*, 2019, pp. 668–675.
- [2] Z. Lu, Z. Liu, G. J. Correa, and K. Karydis, "Motion planning for collision-resilient mobile robots in obstacle-cluttered unknown environments with risk reward trade-offs," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 7064–7070.
- [3] M. Mote, M. Egerstedt, E. Feron, A. Bylard, and M. Pavone, "Collision-inclusive trajectory optimization for free-flying spacecraft," *Journal of Guidance, Control, and Dynamics*, pp. 1–12, 2020.
- [4] J. Zha and M. W. Mueller, "Exploiting collisions for sampling-based multicopter motion planning," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 7943–7949.
- [5] T. Lew, T. Emmei, D. D. Fan, T. Bartlett, A. Santamaria-Navarro, R. Thakker, and A.-a. Agha-mohammadi, "Contact inertial odometry: collisions are your friends," in *The International Symposium of Robotics Research*. Springer, 2019, pp. 938–958.
- [6] T. Schmickl, R. Thenius, C. Moeslinger, G. Radspieler, S. Kernbach, M. Szymanski, and K. Crailsheim, "Get in touch: cooperative decision making based on robot-to-robot collisions," *Autonomous Agents and Multi-Agent Systems*, vol. 18, no. 1, pp. 133–155, 2009.
- [7] K. Karydis, D. Zarrouk, I. Poulakakis, R. S. Fearing, and H. G. Tanner, "Planning with the star (s)," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 3033–3038.
- [8] D. W. Haldane, M. M. Plecnik, J. K. Yim, and R. S. Fearing, "Robotic vertical jumping agility via series-elastic power modulation," *Science Robotics*, vol. 1, no. 1, 2016.
- [9] Y. Mulgaonkar, A. Makineni, L. Guerrero-Bonilla, and V. Kumar, "Robust aerial robot swarms without collision avoidance," *IEEE Robotics and Automation Letters*, vol. 3, no. 1, pp. 596–603, 2017.
- [10] S. Mayya, P. Pierpaoli, G. Nair, and M. Egerstedt, "Localization in densely packed swarms using interrobot collisions as a sensing modality," *IEEE Transactions on Robotics*, vol. 35, no. 1, pp. 21–34, 2018.
- [11] A. Stager and H. G. Tanner, "Composition of local potential functions with reflection," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 5558–5564.

- [12] N. Khedekar, F. Mascarich, C. Papachristos, T. Dang, and K. Alexis, "Contact-based navigation path planning for aerial robots," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 4161–4167.
- [13] Y. Mulgaonkar, W. Liu, D. Thakur, K. Daniilidis, C. J. Taylor, and V. Kumar, "The tiercel: A novel autonomous micro aerial vehicle that can map the environment by flying into obstacles," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 7448–7454.
- [14] Z. Liu and K. Karydis, "Toward impact-resilient quadrotor design, collision characterization and recovery control to sustain flight after collisions," in *IEEE International Conference on Robotics and Automation*, 2021, pp. 183–189.
- [15] M. Hoy, A. S. Matveev, and A. V. Savkin, "Algorithms for collision-free navigation of mobile robots in complex cluttered environments: a survey," *Robotica*, vol. 33, no. 3, pp. 463–497, 2015.
- [16] S. Campbell, W. Naeem, and G. W. Irwin, "A review on improving the autonomy of unmanned surface vehicles through intelligent collision avoidance manoeuvres," *Annual Reviews in Control*, vol. 36, no. 2, pp. 267–283, 2012.
- [17] H. Oleynikova, Z. Taylor, R. Siegwart, and J. Nieto, "Safe local exploration for replanning in cluttered unknown environments for microaerial vehicles," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1474–1481, 2018.
- [18] P. De Petris, H. Nguyen, T. Dang, F. Mascarich, and K. Alexis, "Collision-tolerant autonomous navigation through manhole-sized confined environments," in *IEEE International Symposium on Safety, Security, and Rescue Robotics*, 2020, pp. 84–89.
- [19] A. Stager and H. G. Tanner, "Stochastic behavior of robots that navigate by interacting with their environment," in *IEEE 55th Conference on Decision and Control*, 2016, pp. 6871–6876.
- [20] Z. Lu, Z. Liu, and K. Karydis, "Deformation recovery control and post-impact trajectory replanning for collision-resilient mobile robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2021, pp. 2030–2037.
- [21] S. Liu, N. Atanasov, K. Mohta, and V. Kumar, "Search-based motion planning for quadrotors using linear quadratic minimum time control," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 2872–2879.
- [22] S. Liu, M. Watterson, K. Mohta, K. Sun, S. Bhattacharya, C. J. Taylor, and V. Kumar, "Planning dynamically feasible trajectories for quadrotors using safe flight corridors in 3-d complex environments," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1688–1695, 2017.
- [23] B. Zhou, J. Pan, F. Gao, and S. Shen, "Raptor: Robust and perception-aware trajectory replanning for quadrotor fast flight," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1992–2009, 2021.
- [24] J. Tordesillas, B. T. Lopez, M. Everett, and J. P. How, "Faster: Fast and safe trajectory planner for navigation in unknown environments," *IEEE Transactions on Robotics*, vol. 38, no. 2, pp. 922–938, 2021.
- [25] R. Deits and R. Tedrake, "Efficient mixed-integer planning for uavs in cluttered environments," in *IEEE International Conference on Robotics and Automation*, 2015, pp. 42–49.
- [26] F. Gao, W. Wu, Y. Lin, and S. Shen, "Online safe trajectory generation for quadrotors using fast marching method and bernstein basis polynomial," in *IEEE International Conference on Robotics and Automation*, 2018, pp. 344–351.
- [27] J. Tordesillas, B. T. Lopez, J. Carter, J. Ware, and J. P. How, "Real-time planning with multi-fidelity models for agile flights in unknown environments," in *IEEE International Conference on Robotics and Automation*, 2019, pp. 725–731.
- [28] D. Liu, M. Cong, and Y. Du, "Episodic memory-based robotic planning under uncertainty," *IEEE Transactions on Industrial Electronics*, vol. 64, no. 2, pp. 1762–1772, 2016.
- [29] B. T. Lopez and J. P. How, "Aggressive collision avoidance with limited field-of-view sensing," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 1358–1365.
- [30] K. Mohta, M. Watterson, Y. Mulgaonkar, S. Liu, C. Qu, A. Makineni, K. Saulnier, K. Sun, A. Zhu, J. Delmerico, K. Karydis, N. Atanasov, G. Loiano, D. Scaramuzza, K. Daniilidis, C. J. Taylor, and V. Kumar, "Fast, autonomous flight in gps-denied and cluttered environments," *Journal of Field Robotics*, vol. 35, no. 1, pp. 101–120, 2018.
- [31] C. Richter and N. Roy, "Learning to plan for visibility in navigation of unknown environments," in *International Symposium on Experimental Robotics*. Springer, 2016, pp. 387–398.
- [32] E. Heiden, K. Hausman, G. S. Sukhatme, and A.-a. Agha-mohammadi, "Planning high-speed safe trajectories in confidence-rich maps," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 2880–2886.
- [33] A. Briod, P. Kornatowski, J.-C. Zufferey, and D. Floreano, "A collision-resilient flying robot," *Journal of Field Robotics*, vol. 31, no. 4, pp. 496–509, 2014.
- [34] D. W. Haldane, C. S. Casarez, J. T. Karras, J. Lee, C. Li, A. O. Pullin, E. W. Schaler, D. Yun, H. Ota, A. Javey, and R. S. Fearing, "Integrated manufacture of exoskeletons and sensing structures for folded millirobots," *Journal of Mechanisms and Robotics*, vol. 7, no. 2, p. 021011, 2015.
- [35] T. Li, Z. Zou, G. Mao, X. Yang, Y. Liang, C. Li, S. Qu, Z. Suo, and W. Yang, "Agile and resilient insect-scale robot," *Soft Robotics*, vol. 6, no. 1, pp. 133–141, 2019.
- [36] A. Battiston, I. Sharf, and M. Nahon, "Attitude estimation for collision recovery of a quadcopter unmanned aerial vehicle," *The International Journal of Robotics Research*, vol. 38, no. 10-11, pp. 1286–1306, 2019.
- [37] A. Briod, P. Kornatowski, A. Klapotcz, A. Garnier, M. Pagnamenta, J.-C. Zufferey, and D. Floreano, "Contact-based navigation for an autonomous flying robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2013, pp. 3987–3992.
- [38] A. Q. Nilles, Y. Ren, I. Becerra, and S. M. LaValle, "A visibility-based approach to computing non-deterministic bouncing strategies," *The International Journal of Robotics Research*, vol. 40, no. 10-11, pp. 1196–1211, 2021.
- [39] T. Alam, L. Bobadilla, and D. A. Shell, "Minimalist robot navigation and coverage using a dynamical system approach," in *IEEE International Conference on Robotic Computing*, 2017, pp. 249–256.
- [40] J. S. Lewis and J. M. O’Kane, "Planning for provably reliable navigation using an unreliable, nearly sensorless robot," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1342–1357, 2013.
- [41] P. E. Sandin, *Robot mechanisms and mechanical devices illustrated*. McGraw-Hill, 2003.
- [42] F. Gao, Y. Lin, and S. Shen, "Gradient-based online safe trajectory generation for quadrotor flight in complex environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2017, pp. 3681–3688.
- [43] M. Likhachev, G. J. Gordon, and S. Thrun, "Ara*: Anytime a* with provable bounds on sub-optimality," *Advances in Neural Information Processing Systems*, vol. 16, pp. 767–774, 2003.
- [44] B. Zhou, F. Gao, L. Wang, C. Liu, and S. Shen, "Robust and efficient quadrotor trajectory generation for fast autonomous flight," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3529–3536, 2019.
- [45] F. L. Lewis, D. Vrabie, and V. L. Syrmos, *Optimal control*. John Wiley & Sons, 2012.
- [46] E. Verriest and F. Lewis, "On the linear quadratic minimum-time problem," *IEEE Transactions on Automatic Control*, vol. 36, no. 7, pp. 859–863, 1991.
- [47] D. Harabor and A. Grastien, "Online graph pruning for pathfinding on grid maps," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 25, no. 1, 2011.
- [48] C. Richter, A. Bry, and N. Roy, "Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments," in *Robotics Research*. Springer, 2016, pp. 649–666.
- [49] J. Zha, X. Wu, J. Kroeger, N. Perez, and M. W. Mueller, "A collision-resilient aerial vehicle with icosahedron tensegrity structure," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020, pp. 1407–1412.
- [50] R. Sandström, D. Uwacu, J. Denny, and N. M. Amato, "Topology-guided roadmap construction with dynamic region sampling," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6161–6168, 2020.
- [51] J. D. Hernández, E. Vidal, G. Vallicrosa, E. Galceran, and M. Carreras, "Online path planning for autonomous underwater vehicles in unknown environments," in *IEEE International Conference on Robotics and Automation*, 2015, pp. 1152–1157.
- [52] S. Karaman, M. R. Walter, A. Perez, E. Frazzoli, and S. Teller, "Any-time motion planning using the rrt," in *IEEE International Conference on Robotics and Automation*, 2011, pp. 1478–1483.