

Differentiable Collision Avoidance Using Collision Primitives

Simon Zimmermann¹, Matthias Busenham¹, Simon Huber¹, Roi Poranne², Stelian Coros¹

Abstract—A central aspect of robotic motion planning is collision avoidance, where a multitude of different approaches are currently in use. Optimization-based motion planning is one method, that often heavily relies on distance computations between robots and obstacles. These computations can easily become a bottleneck, as they do not scale well with the complexity of the robots or the environment. To improve performance, many different methods suggested to use collision primitives, i.e. simple shapes that approximate the more complex rigid bodies, and that are simpler to compute distances to and from. However, each pair of primitives requires its own specialized code, and certain pairs are known to suffer from numerical issues. In this paper, we propose an easy-to-use, unified treatment of a wide variety of primitives. We formulate distance computation as a minimization problem, which we solve iteratively. We show how to take derivatives of this minimization problem, allowing it to be seamlessly integrated into a trajectory optimization method. We demonstrate that the resulting method can be used to plan smooth and collision-free paths based on a variety of single- and multi-robot scenarios with different obstacles.

I. INTRODUCTION

Collision avoidance is an integral part of robotic motion planning. Cluttered environments like construction sites, where many potential collisions may occur, are burdened with great computational load. Planning paths for multiple robots requires dynamic and flexible collision avoidance, making the problem more difficult. Furthermore, where human collaborators are involved, robots need to plan ahead while treating the humans as unpredictable, moving obstacles. Indeed, the increasing complexity of tasks that robots are expected to perform certainly requires an equal increase in the efficiency of motion planning algorithms.

Our goal in this paper is to derive a simple, yet robust and *customizable* approach for collision avoiding trajectory optimization. A common practice with this approach is to utilize distance functions between obstacles and define motion planning as a constrained optimization problem. In that sense, the distance functions are used to penalize proximity of the robot to obstacles. Computing the true distance to an obstacle, however, can be computationally demanding. Instead, previous approaches opted to use approximations in the form of collision primitives. That is, they replace the complex geometry of robots and obstacles by simpler shapes, such as spheres, that are easy to compute distances to. However, the accuracy of the approximation depends on the number of collision primitives used. Indeed, an oblong

shape such as the links in a robot’s arm might only be faithfully represented by a multiplicity of spheres, depending on their length and widths. Follow up work suggested using alternative primitives, such as ellipsoids, capsules, boxes, and their combinations, that can better fit to different geometries, and thus reduce the number of primitives necessary [1]. However, this typically requires specialized code that computes distances for every type of primitive pair, which could be cumbersome to maintain, especially when using gradient-based methods. In addition, some of the computations involved are known to be numerically sensitive [2], which can cause numerical issues throughout the planning process, especially when derivatives are required.

Our approach is to formulate distance computation as a low-level, *differentiable* optimization problem, where we explicitly handle numerical issues for all primitive pairs by adding a simple regularization term. We combine this approach with a straight-forward parameterization of different collision primitives, resulting in a light-weight and easy-to-understand implementation that can be found online¹. We show how to take derivatives of the distance computation by leveraging sensitivity analysis. Finally, we integrate this approach into a high-level trajectory optimization problem for collision-free multi-robot motion planning with both fixed and moving obstacles. We evaluate the efficacy of our method based on a variety of simulated and real-world experiments, involving single- as well as multi-robot scenarios.

II. RELATED WORK

The literature on path planning and collision avoidance is vast, with many conceptually different approaches currently in use, each with its own benefits and drawbacks. Therefore, we only highlight a few selected papers, and focus our attention on collision-free path planning via smooth trajectory optimization. In these type of formulations, collision avoidance is typically incorporated in form of inequality constraints, similarly to our approach.

Some of these related methods are agnostic to the representation of the obstacles, and rely on collision *detection* only to define an optimization problem [3]. A straightforward approach is to describe the robot and obstacles as *polyhedrons*, which leads to disjunctive linear inequalities that can be addressed in a mixed-integer problem [4]. Using Farkas’ lemma, it is also possible to convert these disjunctive inequalities into standard inequalities [5]. A more recent approach presented in [6] can also handle more general ellipsoidal obstacles. It employs duality theory to formulate

¹The authors are with the Department of Computer Science, ETH, Zurich, Switzerland. simon.zimmermann@inf.ethz.ch; busenham@student.ethz.ch; simon.huber@inf.ethz.ch; roi.poranne@inf.ethz.ch; scoros@gmail.com

²Department of Computer Science, University of Haifa, Haifa, Israel

¹Code written in C++: <https://github.com/SimiPro/DCA>

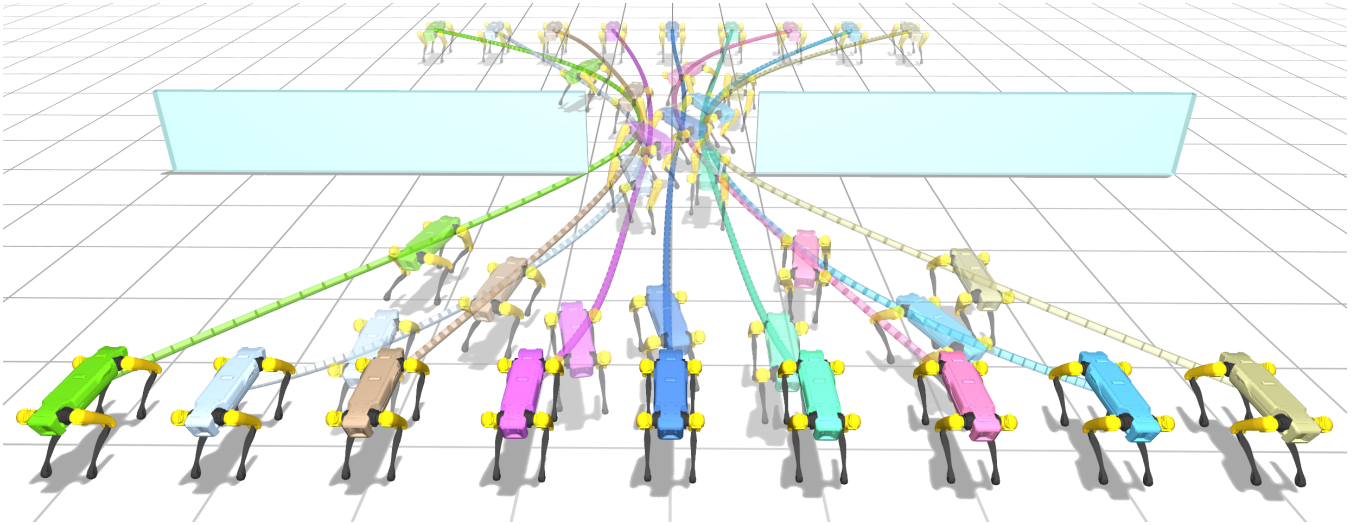


Fig. 1. A pack of Spot robots are finding their way through a gap without colliding. Several poses are overlaid.

simple, differentiable constraints. The resulting optimization problem is then solved using IPOPT [7], a generic interior point solver.

Another common approach we briefly mention uses *potential fields*. First presented in [8], the idea is to create an artificial potential that attracts the robot to the goal and a potential that *repels* the robot away from obstacles. A few examples can be found in [9]–[14]. A related approach uses distance fields instead of potential fields. It is used to constrain the distance between a robot and an obstacle to be greater than a certain safety margin. Notable examples are CHOMP [15] and STOMP [16], both of which sample a distance field on a regular grid as a preprocessing step. While this approach is fairly efficient, it is limited to static obstacles and a single robot scenarios. In addition, the memory requirements might be prohibitive.

We take an approach that, in some sense, lies between the above-mentioned methodologies. We utilize *collision primitives*, i.e. simple shapes that allow simple and efficient distance computations, to approximate the robots and obstacles in the scene. This has been a common strategy in past as well as in recent work. Particularly spheres are applied as collision primitives in countless papers due to their ease-of-use [17], [18]. Ellipsoids and capsules are also commonly found [1], [19]–[21], as well as boxes [22]. Our approach unifies all of these primitives in an easy-to-use manner.

An alternative methodology is presented by Schulman et al [23]. There, robot bodies and obstacles are represented by convex shapes, using a support mapping representation. The distance between two shapes is formalized using signed distance functions, which are computed using the Gilbert–Johnson–Keerthi (GJK) distance algorithm. Collision avoidance is then incorporated into a high-level sequential convex optimization problem, where they approximate the corresponding derivatives by linearizing the signed distance function. Our approach presents a different strategy: We start with basic collision primitives like spheres and

extend the concept to more complex shapes. With this approach, we can easily handle degenerate cases that plague the GJK algorithm and result in numerical issues, which occur surprisingly often in practice.

III. METHOD

To simplify the exposition, we first describe our approach using a simple example of two rigid bodies. The method is easily generalized to multiple robots and obstacles, both stationary or mobile, which we discuss at the end of this section.

A. Collision-Free Motion Planning

Our goal is to plan a smooth, collision-free trajectory for two free-floating rigid bodies denoted by \mathcal{B}^A and \mathcal{B}^B , respectively. We formulate this as a *time-discretized* trajectory optimization problem. To this end, let \mathbf{x}_i^A and \mathbf{x}_i^B be the corresponding states of \mathcal{B}^A and \mathcal{B}^B , comprising of a rotation and a translation in world coordinates at a specific trajectory step i . We represent the entire state by stacking both states into one vector $\mathbf{x}_i = (\mathbf{x}_i^A, \mathbf{x}_i^B)$. Then, $\mathbf{x} := (\mathbf{x}_1, \dots, \mathbf{x}_N)$ represents the entire state trajectory that we want to optimize, consisting of a total number of N steps. Furthermore, let $\mathcal{D}^{AB}(\mathbf{x}_i)$ be the squared distance between the two rigid bodies at trajectory step i , defined as the *shortest* squared distance between any pair of points on the two rigid bodies in their respective states. We include a safety margin with each of the individual rigid bodies, which we denote by r^A and r^B , respectively. Collision avoidance is formulated as an inequality constraint, which forces the shortest distance between the rigid bodies to be greater than the sum of their safety margins. With these definitions, we write the trajectory optimization problem as

$$\min_{\mathbf{x}} \mathcal{O}(\mathbf{x}) \quad (1a)$$

$$\text{s.t. } \mathcal{D}^{AB}(\mathbf{x}_i) \geq r^A + r^B \quad \forall i \in 1, \dots, N \quad (1b)$$

with objective \mathcal{O} , which typically consists of two terms. The first term expresses a goal objective, which matches the state of the rigid bodies at a given trajectory step i to a predefined target $\bar{\mathbf{x}}_i$. The second term is a regularization term that encourages smooth motions by penalizing high accelerations throughout the trajectory. We discretize the acceleration as $\ddot{\mathbf{x}}_i \approx \frac{\mathbf{x}_i - 2\mathbf{x}_{i-1} + \mathbf{x}_{i-2}}{h^2}$, where h is the step duration. The objective (1a) is then

$$\mathcal{O}(\mathbf{x}) = \sum_{i \in I} \|\mathbf{x}_i - \bar{\mathbf{x}}_i\|^2 + w_S \sum_{i=1}^N \|\ddot{\mathbf{x}}_i\|^2, \quad (2)$$

where I is the set of fixed target states $\bar{\mathbf{x}}_i$, and w_S denotes the regularization weight.

B. Collision Primitives

The problem stated in Eq. (1) is an idealized one, as \mathcal{D}^{AB} is set to compute the *true* distance between the two bodies \mathcal{B}^A and \mathcal{B}^B . However, this is often not practical since the computation of this distance – and its derivatives – can be expensive for general shapes. Instead, a common practice is to approximate rigid bodies using simple collision primitives, which are endowed with simpler distance computation routines. The simplest primitives used (for trajectory optimization) are perhaps spheres. They can be seen as points with a safety margin corresponding to the radius, and therefore require only point-to-point distances. Furthermore, when a rigid body has a shape that can not be approximated by one sphere sufficiently well, several smaller spheres can easily be used together to obtain a better approximation. The drawback, however, is that some shapes require *too* many spheres to be approximated well, offsetting the computational benefit of using them in the first place. In these cases, more elaborate primitive shapes such as ellipsoids, capsules, or boxes can be used at the cost of slightly increased computation times, as mentioned in Sec. II. Therefore, creating an efficient approximation of an object using collision primitives requires balancing the complexity of the primitives with their numbers. To simplify the process, we propose a unified formulation for a set of different collision primitives. This formulation is based on the simple observation that common primitives can be described *parametrically* using a point \mathbf{p} and a varying number of vectors \mathbf{v}_l , each scaled by a parameter t_l . The number of vectors and scale parameters depends on the type of primitive. For example, a primitive with no scaled vector simply results in a point \mathbf{p} , which, combined with its radius as a safety margin, results in a sphere collision primitive. In the same manner, a point and one scaled vector $\mathbf{p} + t_1 \mathbf{v}_1$, $0 \leq t_1 \leq 1$ describes a line segment, which can be used as a capsule primitive by choosing its radius as a safety margin again.

We now extend this formulation for the general case: For a rigid body \mathcal{B} with state \mathbf{x}_i , let \mathbf{t}_i be the set of parameters stacked into one vector that describe all points lying on the corresponding collision primitive. In the following, we neglect the subscript i for brevity, and consequently use \mathbf{x} to

represent the state of a single trajectory step. Let $\mathbf{P}(\mathbf{x}, \mathbf{t})$ be a function that returns a specific point on the primitive as a function of \mathbf{x} and \mathbf{t} . In this work, we select a set of collision primitives that can all be described by

$$\mathbf{P}(\mathbf{x}, \mathbf{t}) = \mathbf{p}(\mathbf{x}) + \sum_{l=1}^L t_l \mathbf{v}_l(\mathbf{x}), \quad 0 \leq t_l \leq 1, \quad \forall l \quad (3)$$

where $\mathbf{p}(\mathbf{x})$ is a point on the primitive, and $\mathbf{v}_l(\mathbf{x})$ denote a varying number L of vectors, depending on the type of primitive. All primitive shapes used in this work are visualized in Table I. We find that this set of simple primitives is sufficient for our application. Nevertheless, we note that more constraints can be added to create different shapes, e.g. $\sum_{l=1}^L t_l \leq 1$ for a simplex. Additionally, any convex collision primitive that can be described parametrically in a continuously differentiable manner with respect to both \mathbf{x} and \mathbf{t} can be used as well.

C. Shortest Distance Computation

We return to the two rigid bodies \mathcal{B}^A and \mathcal{B}^B , and assume their collision primitives are given by $\mathbf{P}^A(\mathbf{x}^A, \mathbf{t}^A)$ and $\mathbf{P}^B(\mathbf{x}^B, \mathbf{t}^B)$, where we stack $\mathbf{t} = (\mathbf{t}^A, \mathbf{t}^B)$. The squared shortest distance between the two primitives can be written as

$$\mathcal{D}^{AB}(\mathbf{x}) = \min_{0 \leq \mathbf{t} \leq 1} \mathcal{D}^{AB}(\mathbf{x}, \mathbf{t}) \quad (4)$$

with

$$\mathcal{D}^{AB}(\mathbf{x}, \mathbf{t}) = \|\mathbf{P}^A(\mathbf{x}^A, \mathbf{t}^A) - \mathbf{P}^B(\mathbf{x}^B, \mathbf{t}^B)\|^2. \quad (5)$$

This problem can be solved analytically for the cases we discuss in this paper. However, this must be done in a case-by-case manner, and additionally can result in some numerical issues, especially when computing derivatives. For example, for two capsule primitives parameterized as line segments, these issues occur when the two lines are close to being parallel, as the closest points between them go to infinity [2]. Similar issues arise for other pairs of primitives as well, and we find that they frequently occur in practice. Instead of treating each individual case, we propose a more generic approach: We solve Eq. (4) using iterative numerical optimization. This also allows us to include regularization terms, and to make the collision constraints *soft*, which in turn makes collision avoidance more robust. Concretely, we add a regularization term

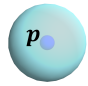
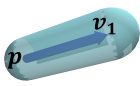
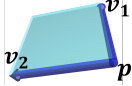
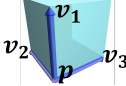
$$\mathcal{R}(\mathbf{t}) = \|\mathbf{t} - 0.5\|^2, \quad (6)$$

which regularizes the \mathbf{t} 's such that the resulting points on the primitives are closer to their center. This simple term effectively avoids the numerical issues described above. In addition, we replace the box constraints $0 \leq \mathbf{t} \leq 1$ in Eq. (4) by soft barrier constraints. Specifically, we define two unilateral quadratic barrier functions as

$$S_l^+(t) = \begin{cases} 0 & t \leq l \\ (t-l)^2 & t > l \end{cases}, \quad S_l^-(t) = \begin{cases} (t-l)^2 & t \leq l \\ 0 & t > l \end{cases}. \quad (7)$$

TABLE I

Collision primitives used in this work. All can be parameterized using Eq. (3). For simplicity, we temporarily drop all dependencies from \mathbf{x} .

Primitive	Sphere	Capsule	Rectangle	Box
$\mathbf{P}(\mathbf{x}, \mathbf{t}) =$	\mathbf{p}	$\mathbf{p} + t_1 \mathbf{v}_1$	$\mathbf{p} + t_1 \mathbf{v}_1 + t_2 \mathbf{v}_2$	$\mathbf{p} + t_1 \mathbf{v}_1 + t_2 \mathbf{v}_2 + t_3 \mathbf{v}_3$
				

Thus, we rewrite the optimization problem (4) as

$$\mathcal{U}^{AB}(\mathbf{x}) = \min_{\mathbf{t}} \mathcal{U}^{AB}(\mathbf{x}, \mathbf{t}) \quad (8)$$

with

$$\mathcal{U}^{AB}(\mathbf{x}, \mathbf{t}) = \mathcal{D}^{AB}(\mathbf{x}, \mathbf{t}) + w_R \mathcal{R}(\mathbf{t}) + w_C (\mathcal{S}_1^+(\mathbf{t}) + \mathcal{S}_0^-(\mathbf{t})), \quad (9)$$

where w_R, w_C are the regularization and penalty weights, which we choose to be $w_R = 0.01$ and $w_C = 10$. We solve this unconstrained optimization problem using Newton's method.

D. Computing Derivatives

To solve the motion planning problem (1), we want to be able to compute the derivatives of the distance function \mathcal{D} w.r.t. to \mathbf{x} . Following the derivation above, \mathcal{D} is not only a function of \mathbf{x} , but also of \mathbf{t} , which in turn has to be seen as a function of \mathbf{x} as well. Consequently, $\mathcal{D}(\mathbf{x}) := \mathcal{D}(\mathbf{x}, \mathbf{t}(\mathbf{x}))$, and computing the gradient requires the use of the chain rule

$$\frac{d\mathcal{D}^{AB}}{d\mathbf{x}} = \frac{\partial \mathcal{D}^{AB}}{\partial \mathbf{x}} + \frac{\partial \mathcal{D}^{AB}}{\partial \mathbf{t}} \frac{d\mathbf{t}}{d\mathbf{x}}. \quad (10)$$

The Jacobian $\frac{d\mathbf{t}}{d\mathbf{x}}$ is known as the *sensitivity matrix*, as it describes how \mathbf{t} changes with respect to changes in \mathbf{x} . Since \mathbf{t} is computed numerically by solving the optimization problem (8), there is no direct analytical expression to compute $\frac{d\mathbf{t}}{d\mathbf{x}}$. However, as described in [24], we can readily compute this Jacobian by leveraging the implicit function theorem. For our particular application, it can be applied when the gradient of the optimization problem (8) is zero, i.e. $\frac{\partial \mathcal{U}^{AB}}{\partial \mathbf{t}} = 0$. This is the case when the optimization has been solved. Then, by using the resulting solution for \mathbf{t} , the sensitivity matrix can be computed analytically as

$$\frac{d\mathbf{t}}{d\mathbf{x}} = -\frac{\partial^2 \mathcal{U}^{AB}}{\partial \mathbf{t}^2}^{-1} \frac{\partial^2 \mathcal{U}^{AB}}{\partial \mathbf{t} \partial \mathbf{x}}. \quad (11)$$

We refer the reader to [24] for a full derivation, and simply state the Hessian for our particular use-case. Just as in [24], we use an approximation of the true Hessian, as it avoids the costly computation the higher-order terms. This results in the expression

$$\frac{d^2 \mathcal{D}^{AB}}{d\mathbf{x}^2} \approx \left(\frac{d\mathbf{t}}{d\mathbf{x}}^T \frac{\partial^2 \mathcal{D}^{AB}}{\partial \mathbf{t}^2} + 2 \frac{\partial^2 \mathcal{D}^{AB}}{\partial \mathbf{x} \partial \mathbf{t}} \right) \frac{d\mathbf{t}}{d\mathbf{x}} + \frac{\partial^2 \mathcal{D}^{AB}}{\partial \mathbf{x}^2}, \quad (12)$$

where we can re-use the Jacobian from the gradient computation.

E. Extension to Robot Motion Planning

As a final step, we extend our formulation from two rigid bodies to the case of robotic motion planning. To this end, we model a robot as a kinematic arm with floating base: Let $\mathbf{x}_i \in \mathbb{R}^{n+6}$ denote the state of a robot at trajectory step $i \in [1, \dots, N]$, where N and n denote the total number of trajectory steps and joint angles, respectively. The state \mathbf{x}_i consists of the position and orientation of the base, as well as the joint angles of the arm. Same as before, we describe the entire robot trajectory by stacking all states \mathbf{x}_i into one vector $\mathbf{x} := (\mathbf{x}_1, \dots, \mathbf{x}_N)$. Using these definitions, we can now extend the optimization problem presented in Eq. (1). Since we are now handling kinematic arms, it is more convenient to include end-effector targets into our formulation: Let \mathbf{z}_i be the target pose in global coordinates for the robot's end-effector at a predefined trajectory step i . Let $\mathcal{K}(\mathbf{x}_i, \mathbf{l})$ be the forward kinematics function that transforms the local coordinates \mathbf{l} of the end-effector into its global pose at state \mathbf{x}_i . We can then formulate an inverse kinematics (IK) objective as

$$\|\mathcal{K}(\mathbf{x}_i, \mathbf{l}) - \mathbf{z}_i\|^2 \quad (13)$$

for each end-effector and trajectory step we wish to set a target for. Furthermore, physical limitations of an individual robot like joint, velocity and acceleration limits need to be considered for motion planning. This can be included into the formulation in form of box constraints

$$\mathbf{b}_l \leq \mathbf{h}(\mathbf{x}_i) \leq \mathbf{b}_u, \quad \forall i \quad (14)$$

where \mathbf{b}_l and \mathbf{b}_u denote the lower and upper bounds, respectively, and $\mathbf{h}(\mathbf{x}_i)$ returns the corresponding value at trajectory step i . Velocities and accelerations are approximated using finite differences.

In order to apply collision avoidance, a robot can be approximated by different collision primitives in the same manner as in the single rigid body case. Examples for two different robots are given in Fig. 2. This can be done conveniently by expressing a specific collision primitive in local coordinates of the body it approximates. Since the robot's bodies are connected via a kinematic chain, the formulation as presented in (1b) needs to be slightly adjusted, as the world coordinates of the collision primitives have to be computed using the forward kinematics function \mathcal{K} . Therefore, the inequality constraints are written as

$$\mathcal{D}^{ab}(\mathcal{K}(\mathbf{x}_i, \mathbf{c}^a), \mathcal{K}(\mathbf{x}_i, \mathbf{c}^b)) \geq \mathbf{r}^a + \mathbf{r}^b, \quad \forall i, \forall a, b \in \mathbb{C}, \quad (15)$$




Fig. 2. The YuMi (right) and Kinova (left) robot approximated by collision primitives. The YuMi uses three boxes for the body and five spheres and two capsules per arm. The Kinova uses three capsules and four spheres.

where \mathbf{c}^a and \mathbf{c}^b denote the local coordinates of collision primitive a and b , respectively, and \mathbb{C} includes all collision primitives used in the scene. These constraints formulate self-collision avoidance of an individual robot for all trajectory steps. In case primitives are used to approximate static obstacles, the dependency on \mathcal{K} can be dropped, as the state of the obstacle can be written directly in world coordinates. This formulation can easily be extended to plan motions for multiple robots. To this end, we simply stack the states of all robots in the scene into one large optimization vector \mathbf{x} . Then, the inequality constraints (15) not only cover self-collision avoidance, but also avoidance between different robots. We solve this multi-robot motion planning problem using Newton’s method. To this end, we convert all inequality constraints into soft constraints using barrier functions as presented in Eq. (8). Then, we can compute the gradient and Hessian of the total objective $\mathcal{O}(\mathbf{x})$ in order to apply Newton steps until convergence. An overview of the overall solving strategy is given in Algorithm 1. We note that in practice, we only compute the exact distance and its derivatives between individual primitive pairs if they are in close proximity to another. To this end, we first roughly estimate the distances between each pair according to the distance of their center points.

IV. RESULTS

We evaluate the efficacy of our method on a variety of simulated and real-world experiments. These comprise of path planning tasks for different robotic platforms in both single- and multi-robot scenarios. Our method outputs collision-free, smooth motion trajectories for every robot in the scene. It avoids self-collision for a single robot, collisions between two different robots, and collisions with obstacles. We encourage the reader to watch the accompanying video, where all conducted experiments are shown. An extended version of this video can be found online². More quantitative

²Accompanying video: <https://youtu.be/et0bu--wuy4>




Fig. 3. Three instances from the *package packing* demonstration, visualized both in simulation and in real-world: the Kinova robot places a package into a slightly larger box.

information about the individual experiments can be found in Table II. An Intel Core i7-7709K 4.2Ghz PC has been used to record all measurements.

A. Experiments

Interactive Avoidance. We demonstrate our collision avoidance method in an interactive, real-world setting using a dual-armed YuMi robot [25]. Two users move around obstacles that tracked by a motion capture system and approximated by collision spheres in the planning framework. The robot successfully avoids the obstacles while ensuring that its two arms do not collide with each other. Motion planning in this case is run in a receding horizon fashion. The motion optimization objective for the robot contains a regularization term that encourages the robot to return to its rest pose when it is undisturbed by the obstacles.

Spot Pack. This experiment demonstrates kinematic path planning for the bases of multiple legged robots. We use the simulation model of several Boston Dynamics’ Spot robots [26], which are modeled as floating bases with pose and velocity constraints that resemble those of the physical Spot. The accompanying video and Fig. 1 show different scenarios where the robots are tasked with switching positions while avoiding each other as well as some large world obstacles. We note that motion planning is only conducted for the robot’s bases, but leg motions are kinematically computed in post-processing for visualization purposes.

Package Packing. A statically mounted robot arm picks up a package, and places it in a slightly larger, empty box. The empty box is approximated by four rectangular collision primitives, one for each closed side. This experiment is also conducted in the real world by passing the nominal trajectories computed in simulation onto the physical robot in an open-loop manner. We use a Kinova Gen3 7Dof robot [27] equipped with a Sake EZGripper [28] to do so. Fig. 3 show different frames of this experiment.

Table Reach. A Kinova arm mounted on a Spot robot (see [29]) is tasked with retrieving an object that lies under a table. The table is modeled by five box primitives, one for each leg, and one for the plate. The grasp maneuver is shown in simulation in the video, and is depicted in Fig. 4.

Algorithm 1: Collision-free multi-robot kinematic trajectory optimization using Newton’s method

Input: Total objective $\mathcal{O}(\mathbf{x})$ (including collision avoidance constraints as soft constraints), initial \mathbf{x}

Output: Optimal motion trajectories \mathbf{x}^*

while convergence criterion not reached **do**

 Compute \mathbf{t} by solving optimization problem (8) for each primitive pair in close proximity using Newton’s method

 Compute gradient $\frac{d\mathcal{O}}{d\mathbf{x}}$ and Hessian $\frac{d^2\mathcal{O}}{d\mathbf{x}^2}$ (includes the use of Eq. (10) and (12) by following the chain rule)

 Compute search direction $\Delta\mathbf{x}$ by solving linear system $\frac{d^2\mathcal{O}}{d\mathbf{x}^2}\Delta\mathbf{x} = -\frac{d\mathcal{O}}{d\mathbf{x}}$

 Run backtracking line search on α in $\mathbf{x} := \mathbf{x} + \alpha\Delta\mathbf{x}$

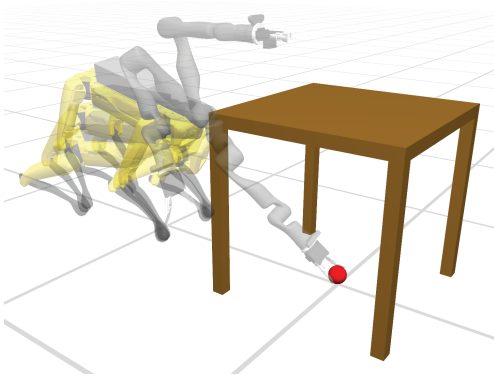


Fig. 4. The combined platform consisting of Spot and Kinova are retrieving a ball from under a table. Several poses of the motion trajectory are overlaid. The motion both start and end at the most upright pose.




Fig. 5. Four different frames of the *house assembly* experiment: The RFL robots cooperatively assemble the structure of a small wooden house.

Gap Handover. The YuMi robot is used to hand over a package from one gripper to the other through a gap in a wall. The wall is modeled with four rectangle collision primitives. The video shows the experiment both in simulation and performed on the physical robot.

House Assembly. This demonstration employs a kinematic model of the Robotic Fabrication Lab (RFL) setup at ETH Zurich [30], which includes four large robotic arms connected by bridges. These robot arms cooperatively assemble a simplified house by planning collision-free paths around each other as well as the structure that has already been built. While a robot is holding a building block in its gripper, the corresponding collision primitive becomes part of the robot, and otherwise it is treated as an obstacle. The assembly is broken down into individual pick and place tasks, where the assembly order and task assignment is done by the user. Fig. 5 shows different scenes from the assembly process.

B. Primitive Configuration

Generally, using more primitives has the greater potential to better approximate the shape of a body. However, more primitives lead to a higher computational cost. We illustrate this using a simple example shown in Fig. 6. There, we solve a trajectory optimization problem involving two cube-shaped obstacles. We approximate one cube using a varying number of spheres and capsules to compare the timings. We show that using more primitives reduces the approximation error, given here by the Hausdorff distance, but increases the computation time per iteration. We can see that for the same Hausdorff distance (i.e. around 0.045 and 0.18), fewer capsules than




Fig. 6. A free-floating rigid body and a world obstacle, both in form of a box, are approximated by different number of primitives. We compare the Hausdorff distance with the average run-time to compute Newton steps. Approximating both boxes with multiple spheres or capsules decreases the Hausdorff distance, but has a quadratic effect on the computation time. For this example, it is highly beneficial to use a box primitive.

spheres are needed, and despite the slight increase in distance computation time per capsule, the overall time is significantly shorter. Naturally, using the cube itself as the primitive is the ideal choice, and this is also evident in the figure. We additionally show in Table III the times required to compute the shortest distance between different primitive pairs using our method. Again, making the primitive higher-dimensional introduces additional Newton steps, which also contributes to higher computation times, but the overall effect for the path planning problem is a reduction in time.

C. Comparison to Other Motion Planners

We conducted several experiments to compare our method to two other optimization-based motion planners: CHOMP [15] and STOMP [16]. We used the default parameter settings

TABLE II

Quantitative results for the different experiments. The state dimension is the sum of all robot state sizes in the corresponding scene. The last two columns show the average computation time per solver iteration and the total computation time (by multiplying the values of the two columns before), respectively.

Experiment	# Robots	# Primitives	State Dim.	# Traj. Steps	# Iterations	\varnothing Time / Iteration	Total Time
Interactive Avoidance	1 (dual arm)	19	20	1	9	0.003 s	0.027 s
Spot Pack	8 - 10	8 - 12	48 - 60	100	446 - 1577	0.217 - 0.488 s	97 - 770 s
Package Packing	1	5	13	160	32	0.218 s	7 s
Table Reach	1	12	13	120	215	0.783 s	168 s
Gap Handover	1 (dual arm)	6	20	160	184	0.373 s	69 s
House Assembly	4	16 - 33	48	30	15 - 68	0.004 - 0.080 s	0.06 - 5.44 s

TABLE III

Number of Newton steps and average time per step for the distance computation between each pair of primitives. Time measurements are averaged over 10'000 runs and expressed in microseconds.

	Sphere	Capsule	Rectangle	Box
Sphere	0 / 0.0	-	-	-
Capsule	1-2 / 0.85	1-3 / 1.54	-	-
Rectangle	1-7 / 1.65	1-5 / 2.30	1-11 / 3.34	-
Box	1-5 / 2.25	1-11 / 3.13	2-9 / 3.86	2-14 / 4.60

including initialization procedures set in their individual ROS implementations (see details online³). One of the experimental setups involved a Kinova arm [27] tasked with finding a collision-free path around an obstacle modeled by either a sphere, a capsule (a cylinder in ROS), or a box. We ran the experiment multiple times, where we gradually increased the size of the individual obstacles to make the path planning problem more elaborate. All methods usually succeeded in finding a feasible path around the obstacle. An overview of the individual measurements is given in Table IV. We note that the overall computation time is dependent on the convergence criteria of the individual planners. When using the default settings, the computation times of all methods were usually comparable. However, we noticed a difference in the quality of the solution. Especially for larger obstacles, some of the planners seemed to struggle to find a smooth path. Our method often provided a smoother and more direct path, both in terms of joint angles (see Fig. 7) as well as for the robot's end-effector (see Table IV).

V. DISCUSSION

To summarize, our approach provides a unified, straightforward framework that can be applied to various collision primitives and safely handles numerical issues that can arise when computing distances and its derivatives. Therefore, our distance computation scheme can seamlessly be integrated into path planners that profit from these properties. Specifically, for optimization-based path planners, the treatment of numerical issues contributes to a good convergence behavior as well as to the satisfaction of the collision avoidance constraints. In terms of limitations, our overall trajectory optimization framework suffers from the same drawbacks

TABLE IV

Performance comparison to other planners. The table shows both the total computation time (in seconds) as well as the distance travelled by the robot's end-effector (in meters) for the final solution. All measurements have been conducted for a small (S), medium (M) and large (L) obstacle.

		CHOMP		STOMP		Ours	
		Time	Dist.	Time	Dist.	Time	Dist.
Sphere	S	1.550	0.621	0.214	0.615	0.629	0.657
	M	1.709	0.946	0.586	0.951	1.483	0.715
	L	2.216	1.115	3.992	1.427	2.143	0.874
Capsule / Cylinder	S	1.951	0.916	0.828	1.175	1.715	0.688
	M	2.109	1.165	1.499	1.319	1.561	0.743
	L	2.457	1.421	3.510	1.491	3.703	0.907
Box	S	1.478	0.621	0.159	0.615	0.545	0.640
	M	1.603	0.739	0.196	0.988	1.026	0.744
	L	2.297	1.329	3.480	1.647	2.538	1.061

as other gradient-based methods: it can only find a local minimum. Consequently, in very cluttered environments and tight spaces, our method might fail to find a collision-free trajectory without a proper initialization. One solution would be to find an initial trajectory using a sampling-based method and use it as a starting point for the optimization. However, due to the relatively quick iteration time and predictable behaviour, we believe it to be simpler for a user to provide input and guide the process toward a feasible solution. Additionally, we experimented with an automatic continuation method that shows promising results. With these methods, infeasible cases can be resolved by starting the optimization in an obstacle-free environment, such that an initial, collision-free trajectory can be generated. Then, obstacles are gradually introduced into the scene while keeping the optimization process running. This is a reliable strategy to find feasible solutions for all the scenarios that we investigated in this paper. However, further analysis and experimentation in this regard will be part of future investigations, with focus on enabling our method to find suitable motion plans more independently. Another avenue to explore is the automatic generation of primitives that would cover the objects. Currently, this is done manually based on previous experience. However, we believe that finding an optimal configuration of primitives that achieves the best performance would be of great value.

³https://ros-planning.github.io/moveit_tutorials/index.html




Fig. 7. Resulting joint velocities for different motion planners for the experiment described in section IV-C using the largest box obstacle. The number of trajectory steps was chosen by the individual planners, and we chose 100 steps for our method for all experiments. Our method typically produced smoother velocities and found a more direct path.

REFERENCES

- [1] D. Rakita, H. Shi, B. Mutlu, and M. Gleicher, "Collisionnik: A per-instant pose optimization method for generating robot motions with environment collision avoidance," *CoRR*, vol. abs/2102.13187, 2021. [Online]. Available: <https://arxiv.org/abs/2102.13187>
- [2] D. Eberly, "Robust computation of distance between line segments," 2015.
- [3] J. Alonso-Mora, T. Naegeli, R. Siegwart, and P. Beardsley, "Collision avoidance for aerial vehicles in multi-agent scenarios," *Auton Robot*, vol. 39, no. 1, pp. 101–121, Jun. 2015. [Online]. Available: <https://doi.org/10.1007/s10514-015-9429-0>
- [4] L. Blackmore and B. Williams, "Optimal manipulator path planning with obstacles using disjunctive programming," in *2006 American Control Conference*. IEEE, 2006, pp. 3–pp.
- [5] M. Gerdtts, R. Henrion, D. Hömberg, and C. Landry, "Path planning and collision avoidance for robots," *Numerical Algebra, Control & Optimization*, vol. 2, no. 3, p. 437, 2012.
- [6] X. Zhang, A. Liniger, and F. Borrelli, "Optimization-Based Collision Avoidance," *IEEE Transactions on Control Systems Technology*, pp. 1–12, 2020, conference Name: IEEE Transactions on Control Systems Technology.
- [7] A. Wächter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006.
- [8] O. Khatib, "Real-time obstacle avoidance for manipulators and mobile robots," in *1985 IEEE International Conference on Robotics and Automation Proceedings*, vol. 2, Mar. 1985, pp. 500–505.
- [9] S. S. Ge and Y. J. Cui, "New potential functions for mobile robot path planning," *IEEE Transactions on Robotics and Automation*, vol. 16, no. 5, pp. 615–620, Oct. 2000, conference Name: IEEE Transactions on Robotics and Automation.
- [10] G. Yngve and G. Turk, "Robust creation of implicit surfaces from polygonal meshes," *IEEE Trans. Visual. Comput. Graphics*, vol. 8, no. 4, pp. 346–359, Oct. 2002. [Online]. Available: <http://ieeexplore.ieee.org/document/1044520/>
- [11] S. Tao and J. Tan, "Path Planning with Obstacle Avoidance Based on Normalized R-Functions," Oct. 2018, iSSN: 1687-9600 Pages: e5868915 Publisher: Hindawi Volume: 2018. [Online]. Available: <https://www.hindawi.com/journals/jtr/2018/5868915/>
- [12] S. N. Gai, R. Sun, S. J. Chen, and S. Ji, "6-DOF Robotic Obstacle Avoidance Path Planning Based on Artificial Potential Field Method," in *2019 16th International Conference on Ubiquitous Robots (UR)*, Jun. 2019, pp. 165–168, iSSN: 2325-033X.
- [13] M. Safeea, R. Béarée, and P. Neto, "Collision Avoidance of Redundant Robotic Manipulators Using Newton's Method," *J Intell Robot Syst*, vol. 99, no. 3-4, pp. 673–681, Sep. 2020. [Online]. Available: <http://link.springer.com/10.1007/s10846-020-01159-3>
- [14] C. Scoccia, G. Palmieri, M. C. Palpacelli, and M. Callegari, "A Collision Avoidance Strategy for Redundant Manipulators in Dynamically Variable Environments: On-Line Perturbations of Off-Line Generated Trajectories," *Machines*, vol. 9, no. 2, p. 30, Feb. 2021, number: 2 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: <https://www.mdpi.com/2075-1702/9/2/30>
- [15] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, "Chomp: Gradient optimization techniques for efficient motion planning," in *2009 IEEE International Conference on Robotics and Automation*, 2009, pp. 489–494.
- [16] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "Stomp: Stochastic trajectory optimization for motion planning," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 4569–4574.
- [17] S. Duenser, J. M. Bern, R. Poranne, and S. Coros, "Interactive robotic manipulation of elastic objects," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 3476–3481.
- [18] M. Gaertner, M. Bjelonic, F. Farshidian, and M. Hutter, "Collision-free MPC for legged robots in static and dynamic scenes," *CoRR*, vol. abs/2103.13987, 2021. [Online]. Available: <https://arxiv.org/abs/2103.13987>
- [19] C. Dube, "Self collision avoidance for humanoids using circular and elliptical capsule bounding volumes," in *2013 Africon*, 2013, pp. 1–6.
- [20] B. Brito, B. Floor, L. Ferranti, and J. Alonso-Mora, "Model predictive contouring control for collision avoidance in unstructured dynamic environments," *IEEE Robotics and Automation Letters*, vol. PP, pp. 1–1, 07 2019.
- [21] A. El Khoury, F. Lamiroux, and M. Taïx, "Optimal motion planning for humanoid robots," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 3136–3141.
- [22] M. Toussaint, "Logic-geometric programming: An optimization-based approach to combined task and motion planning," in *Proceedings of the 24th International Conference on Artificial Intelligence*, ser. IJCAI'15. AAAI Press, 2015, p. 1930–1936.
- [23] J. Schulman, Y. Duan, J. Ho, A. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *The International Journal of Robotics Research*, vol. 33, no. 9, pp. 1251–1270, Aug. 2014, publisher: SAGE Publications Ltd STM. [Online]. Available: <https://doi.org/10.1177/0278364914528132>
- [24] S. Zimmermann, R. Poranne, and S. Coros, "Optimal control via second order sensitivity analysis," *arXiv e-prints*, p. arXiv:1905.08534, May 2019.
- [25] "YuMi | IRB 14000 | Kollaborativer Roboter." [Online]. Available: <https://new.abb.com/products/robotics/de/kollaborative-roboter/yumi>
- [26] "Spot@ | Boston Dynamics." [Online]. Available: <https://www.bostondynamics.com/spot>
- [27] "Kinova Robotics: Kinova Gen3." [Online]. Available: <https://www.kinovarobotics.com/en/products/gen3-robot>
- [28] "Sake Robotics: EZGripper Gen. 2." [Online]. Available: <https://sakerobotics.com/>
- [29] S. Zimmermann, R. Poranne, and S. Coros, "Go Fetch! - Dynamic Grasps Using Boston Dynamics Spot With External Robotic Arm," in *IEEE International Conference on Robotics and Automation, ICRA*, 2021.
- [30] "RFL | Robotic Fabrication Laboratory, ITA Insitute of Technology in Architecture, ETH Zurich." [Online]. Available: <https://ita.arch.ethz.ch/archteclab/rfl.html>