



Sampling-based optimal kinodynamic planning with motion primitives

Basak Sakcak¹ · Luca Bascetta¹ · Gianni Ferretti¹ · Maria Prandini¹

Received: 16 January 2018 / Accepted: 2 January 2019 / Published online: 14 January 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2019

Abstract

This paper proposes a novel sampling-based motion planner, which integrates in Rapidly exploring Random Tree star (RRT*) a database of pre-computed motion primitives to alleviate its computational load and allow for motion planning in a dynamic or partially known environment. The database is built by considering a set of initial and final state pairs in some grid space, and determining for each pair an optimal trajectory that is compatible with the system dynamics and constraints, while minimizing a cost. Nodes are progressively added to the tree of feasible trajectories in the RRT* algorithm by extracting at random a sample in the gridded state space and selecting the best obstacle-free motion primitive in the database that joins it to an existing node. The tree is rewired if some nodes can be reached from the new sampled state through an obstacle-free motion primitive with lower cost. The computationally more intensive part of motion planning is thus moved to the preliminary offline phase of the database construction at the price of some performance degradation due to gridding. Grid resolution can be tuned so as to compromise between (sub)optimality and size of the database. The planner is shown to be asymptotically optimal as the grid resolution goes to zero and the number of sampled states grows to infinity.

Keywords Optimal sampling-based planning · Kinodynamic planning · Motion primitives

1 Introduction

Motion planning is one of the fundamental problems in robotics, and consists of guiding the robot from an initial state to a final one along a collision-free path.

For many robots, focusing only on kinematics could result in collision-free paths that are impossible to be executed by the actual system. In particular, for systems that are differentially constrained, such a decoupled approach makes it difficult to turn a collision-free path into a feasible trajectory. In order to tackle this problem, Donald et al. (1993) proposed the idea of *kinodynamic planning*, which combines the search for a collision-free path with the underlying dynamics of the system, so that the resulting trajectory would be feasible.

For most robotic applications, the solution to the planning problem should not only be feasible and collision-free, but also satisfy some properties such as, e.g., reaching the

goal in minimum time, minimizing the energy consumption, and maximizing safety. These additional requirements have shifted the focus from simply designing collision-free and feasible trajectories to finding optimal ones.

This paper deals with *optimal kinodynamic motion planning* for systems with complex dynamics and subject to constraints.

1.1 Literature review

A significant amount of work in the robotics community has then been dedicated to the problem of kinodynamic planning so as to determine a trajectory that fulfills the differential constraints arising from the dynamics of the robot. Solving this problem is complex, in general, since it requires a search in the state space of the robot, which often implies a higher-dimensional search space compared to a pure kinematic planning.

Most of the motion planners can be classified under the two categories of *exact* and *sampling-based* methods, LaValle (2006). The former looks for a solution in the continuous state space, while the latter samples this space redefining

✉ Basak Sakcak
Basak.Sakcak@polimi.it

¹ Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Piazza L. Da Vinci 32, 20133 Milan, Italy

it as a graph where nodes are connected via edges representing local trajectories between sampled states.

Exact methods are said to be *complete*, since they terminate in finite time with a solution, if one exists, and return the nonexistence otherwise. However, with the exact approaches even the simplest problem is PSPACE hard (Reif 1979). Exact methods require that the obstacles are represented explicitly in the state space, dramatically increasing the problem complexity. However, they can provide practical solutions for problems that are characterized by a low dimensional state space, or for which a low dimensional obstacle representation can be adopted. Obstacles introduce non convex constraints in the admissible state space, and make the problem of computing an optimal trajectory hard. In particular, most of the algorithms, that are gradient based, can only find a solution in the same homotopy class of the initial guess (LaValle 2006). There has been a significant progress to address this issue and, in particular, the ideas of dividing the global optimal trajectory planning problem into simpler subproblems and of using numerical optimization to compute locally optimal trajectories have been explored in, e.g., (Park et al. 2015; Kuderer et al. 2014).

Sampling-based approaches emerged to handle systems with high dimensional state spaces, and they became the most popular approaches in the planning literature, representing the practical way to tackle the problem. The basic idea is to sample states (nodes) in the continuous state space and connect these nodes with trajectories in the collision-free space, building a roadmap in the form of a graph or a tree of feasible trajectories. These algorithms avoid an explicit representation of obstacles by using a collision check module that allows to determine the feasibility of a tentative trajectory. They are not complete, but they satisfy the *probabilistic completeness* property, i.e., they return a solution with a probability converging to one as the number of samples grows to infinity, if such a solution exists.

Probabilistic Road Map (PRM), introduced by Kavraki et al. (1996), and Rapidly exploring Random Trees (RRT), introduced by LaValle and Kuffner (2001), were the first popular sampling-based planners. PRM first creates a graph in the free configuration space by randomly sampling nodes and connecting them to the already existing ones in the graph using a local planner. The graph can then be used to answer multiple queries, where in each query a start node and a goal node are added to the graph and a path connecting the two nodes is looked for. RRT, on the other hand, incrementally builds a tree starting from a given node, returning a solution as soon as the tree reaches the goal region, hence providing a fast on-line implementation. In all the different formulations of sampling-based planners, a steering function is required to design a trajectory (edge in the tree terminology) connecting two nodes of the tree.

Considering the quality of the solution, an important progress has been made with the introduction of RRT* (Rapidly exploring Random Tree star) and PRM* (Probabilistic Road Map star), which have been proven to be *asymptotically optimal*, i.e., the probability of finding an optimal solution, if there exists one, converges to 1 as the tree cardinality grows to infinity, (Karaman and Frazzoli 2011). The main idea of these algorithms is to ensure that each node is connected to the graph optimally, possibly rewiring the graph by testing connections with pre-existing nodes that are in a suitably defined neighborhood. The same strategy applies to kinodynamic planning (Karaman and Frazzoli 2010) as well, with the additional difficulty that when optimality is required, implementing the steering function involves solving a *two point boundary value problem* (TPBVP), which is computationally challenging especially when dealing with complex dynamics, such as for non-holonomic robots, in presence of actuation constraints. In the context of kinodynamic planning RRT* and PRM* cannot be considered in the same way. In fact, PRM* is limited to symmetric costs and to those systems for which the cost associated to a TPBVP is conserved when the boundary pairs are swapped. Note that, nonholonomic systems do not belong to this class. Considering instead RRT*, it must be noticed that for various dynamical systems, such as non-holonomic vehicles, the presence of kinodynamic constraints makes the constrained optimization problem that the steering function has to solve extremely complex.

To deal with this computational complexity, some effort has been made towards developing effective steering functions for different types of dynamical systems. Webb and van den Berg (2013) have obtained the closed-form analytical solution for a minimum time minimum energy optimization problem for systems with linear dynamics, and extended it to non-linear dynamics using first-order Taylor approximation. Other works (Perez et al. 2012; Goretkin et al. 2013) have focused on approximating the solution for systems with linearizable dynamics, by locally linearizing the system and applying linear quadratic regulation (LQR).

Some recent attempts have been made towards optimality without formulating and solving a TPBVP, as well. For example, algorithms like Stable Sparse RRT* (SST*) (Li et al. 2016) have proved asymptotic optimality given access only to a forward propagation model. The idea is to iteratively sample a set of controls and final times instead of explicitly solving the BVP. Similarly, a variant of RRT* (hwan Jeon et al. 2011) uses a shooting method without a steering function to improve the solution by pruning branches from the tree. If a sampled node has a lower cost compared to another one that is close by and that shares the same parent, the pre-existing node is pruned from the tree and its branches are connected to the newly sampled node, or they are pruned completely if they are not collision-free. This approach generates feasible

but inherently suboptimal solutions. Other works on extending RRT* to handle kinodynamic constraints include limiting the volume in the state space from which nodes are selected by tailoring it to the considered dynamical system in order to improve computational effectiveness (Karaman and Frazzoli 2013).

Nevertheless, solving the TPBVP for an arbitrary nonlinear system remains challenging and typically calls for numerical solvers. The algorithms that account for a nonlinear optimization tool, like for example ACADO toolkit (Houska et al. 2011), GPOPS-II (Patterson and Rao 2014), etc., commonly use *Sequential Quadratic Programming* (SQP) (Boggs and Tolle 1995) for solving the TPBVPs numerically, and embed it as a subroutine in the sampling based planning framework such as in RRT* (Stoneman and Lampariello 2014) or in Batch-Informed-Trees star (BIT*) (Gammell et al. 2014).

A different class of algorithms, aiming at optimality, is based on graph search and adopt a gridding approach. The main idea is to discretize the state space, building a grid, and compute a graph. The motion planning problem is then recast into finding the best sequence of motions by traversing this graph with an optimal search algorithm like A* (Pearl 1984). This graph is often represented by a state lattice, a set of states distributed in a regular pattern, where the connections between states are provided by feasible/optimal trajectories (Pivtoraiko et al. 2009). Likhachev and Ferguson (2009) improved the idea of state lattice by using a multi-resolution lattice such that the portion of the graph that is close to either the initial or the goal state has a higher resolution than the other parts. These approaches have been successfully applied to several robotic systems and found to be effective for dynamic environments (Likhachev and Ferguson 2009; Dolgov et al. 2010). However, these algorithms are *resolution optimal*, such that the optimality is guaranteed up to the grid resolution. Furthermore, their computational effectiveness is highly related to the resolution: the finer is the grid, the higher the branching factor, and thus the computational time and the required memory to execute a graph traversal algorithm.

1.2 Contribution of the paper

The main contribution of this paper is proposing an algorithm, called RRT*_MotionPrimitives, which extends RRT* by introducing a database of pre-computed motion primitives in order to avoid the online solution of a constrained TPBVP for the edge computation. The database is composed of a set of trajectories, each one connecting an initial state to a final one in a suitably defined grid. By sampling in the gridded state space, the implementation of the steering function adopted for growing and rewiring the RRT* tree reduces to the search of a motion primitive in a pre-computed Look Up Table (LUT).

The proposed approach is applicable to any dynamical system described by differential equations and subject to analytical constraints, for which edge design can be formulated as a TPBVP. Notably, when a model of the robot is not available, the database can be derived directly from experimental trajectories.

The main difference of the proposed approach, with respect to existing algorithms that use a database of pre-computed trajectories, e.g., search based approaches as in (Pivtoraiko et al. 2009), is that it leverages on a dynamic tree whose size depends only on the number of samples, but not on the number of motion primitives that affects the accuracy in the approximation of the robot kinematic and dynamic characteristics. As a consequence, memory consumption to store the tree and computation time to determine a solution on a given tree can be bounded selecting an appropriate maximum number of samples, without introducing undesired constraints in the robot action space.

Search based approaches, instead, have to strongly limit the action space, keeping the number of motion primitives low, as the branching factor of the graph, i.e., the number of edges generated expanding each node, depends on the size of the database.

The effectiveness of RRT*_MotionPrimitives has been validated in simulation, showing that the time required to build the tree is greatly reduced with the introduction of a LUT. This represents a promising result for online applications, especially in dynamic environments where the planner has to generate a new trajectory in response to changes in the obstacle-free state space.

An analysis of the probabilistic completeness and optimality properties of RRT*_MotionPrimitives is also provided. This involves a two-step procedure where we assess how close the proposed sample-based solution gets to the optimal one in the gridded state space as the number of samples grows to infinity, and how it gets close to the optimal solution in the continuous state space as the gridding gets finer and finer.

1.3 Paper structure

The paper is organized as follows. Section 2 introduces a formal description of the problem. In Sect. 3 the proposed RRT*_MotionPrimitives algorithm is explained in detail. An analysis of its probabilistic completeness and optimality properties is presented in Sect. 4. A numerical validation of RRT*_MotionPrimitives is provided in Sect. 5. Finally, some concluding remarks are drawn in Sect. 6.

2 Problem statement

In this work, dynamical systems with state vector $\mathbf{q} \in \mathbb{R}^d$ and control input $\mathbf{u} \in \mathbb{R}^m$, governed by

$$\dot{\mathbf{q}}(t) = f(\mathbf{q}(t), \mathbf{u}(t)) \quad (1)$$

where f is continuously differentiable as a function of both arguments, are considered. The control input $\mathbf{u}(t)$ is subject to actuation constraints, and the admissible control space is denoted as $U \subset \mathbb{R}^m$. The state $\mathbf{q}(t)$ is constrained in the set $Q \subset \mathbb{R}^d$, and initialized with $\mathbf{q}(0) = \mathbf{q}_0 \in Q$. Both U and Q are assumed to be compact. An open subset Q_{goal} of Q represents the goal region that the state has to reach.

A trajectory of system (1) is defined by the tuple $\mathbf{z} = (\mathbf{q}(\cdot), \mathbf{u}(\cdot), \tau)$, where τ is the duration of the trajectory, and $\mathbf{q}(\cdot) : [0, \tau] \rightarrow Q$ and $\mathbf{u}(\cdot) : [0, \tau] \rightarrow U$ define the state and control input evolution along the time interval $[0, \tau]$, satisfying the differential equation (1) for $t \in [0, \tau]$, the initial condition $\mathbf{q}(0) = \mathbf{q}_0$, and the final condition $\mathbf{q}(\tau) \in Q_{goal}$.

Obstacles are represented via an open subset Q_{obs} of Q . The free space is then defined as $Q_{free} := Q \setminus Q_{obs}$, and the assumption $\mathbf{q}_0 \in Q_{free}$ is enforced.

A trajectory $\mathbf{z} = (\mathbf{q}(\cdot), \mathbf{u}(\cdot), \tau)$ of system (1) is said to be *collision free*, if it avoids collisions with obstacles, i.e., $\mathbf{q}(t) \in Q_{free}, t \in [0, \tau]$.

The set of collision free trajectories is denoted as Z_{free} .

An *optimal kinodynamic motion planning problem* can then be formalized as finding a feasible and collision free trajectory $\mathbf{z}^* = (\mathbf{q}^*(\cdot), \mathbf{u}^*(\cdot), \tau^*) \in Z_{free}$ that is *optimal* according to a cost criterion $J(\mathbf{z}) : Z_{free} \rightarrow \mathbb{R}_{\geq 0}$ that is expressed as

$$J(\mathbf{z}) = \int_0^\tau g(\mathbf{q}(t), \mathbf{u}(t)) dt$$

where $g : Q \times U \rightarrow \mathbb{R}_{\geq 0}$ is an instantaneous cost function. We assume that trajectories joining two different states have a non-zero cost. This is for instance the case in minimum time optimization where $g(\mathbf{q}, \mathbf{u}) = 1$ and the trajectory duration τ is one of the optimization variables of the problem.

Remark 1 (*translation invariance property*) In the context of motion planning, system (1) represents the robot equations of motion and, consequently, its state vector \mathbf{q} includes the robot position with respect to a given absolute reference frame. In the following, a translation invariance property is supposed to hold. This means that, if obstacles are neglected and a pair of initial and final states and the associated optimal trajectory $\mathbf{z}^* = (\mathbf{q}^*(\cdot), \mathbf{u}^*(\cdot), \tau^*)$ are considered, by shifting the origin of the coordinate system and rewriting all relevant quantities—including system dynamics (1), and initial

and final states—in the new coordinate system and applying input $\mathbf{u}^*(\cdot)$, the optimal robot trajectory is obtained, which is \mathbf{z}^* rewritten in the new coordinates.

3 RRT* with motion primitives

The approach here proposed is based on previous works on search-based (Pivtoraiko et al. 2009; Likhachev and Ferguson 2009), and sampling-based (Karaman and Frazzoli 2011, 2010) methods, and combines them in a novel way.

In particular, it relies on a uniform discretization of the state space, and on the computation of a finite set of motion primitives by solving a constrained optimization problem with boundary conditions on the grid points of a smaller (uniform) grid. The motion primitives are then embedded in the RRT* algorithm, where they are used to connect the randomly generated nodes to the tree, thus eliminating the need of solving online challenging and time consuming TPBVPs.

3.1 Database of motion primitives

The database of motion primitives is built by gridding the continuous state space in order to obtain a finite set of boundary conditions (initial and final states), and by solving offline a constrained boundary value optimization problem for each pair. The resulting set of optimal trajectories is then used repeatedly online, implementing a procedure that, when an edge connecting two nodes is requested by the planner, simply picks up from the database a suitable trajectory.

Given a state tuple $\mathbf{q} \in Q$, $\mathbf{q} = [\boldsymbol{\pi}^T, \dots]^T$, including the robot position $\boldsymbol{\pi}$, motion primitives are computed for each pair of initial and final states $(\bar{\mathbf{q}}_0, \bar{\mathbf{q}}_f)$ with $\bar{\mathbf{q}}_0 = [\bar{\boldsymbol{\pi}}_0^T, \dots]^T$. Given a boundary value pair $(\bar{\mathbf{q}}_0, \bar{\mathbf{q}}_f)$, a motion primitive is computed by solving the following optimization problem

$$\begin{aligned} & \underset{\mathbf{u}(\cdot), \tau}{\text{minimize}} && \int_0^\tau g(\mathbf{q}(t), \mathbf{u}(t)) dt \\ & \text{subject to} && \dot{\mathbf{q}}(t) = f(\mathbf{q}(t), \mathbf{u}(t)) \\ & && \mathbf{u}(t) \in U, t \in [0, \tau] \\ & && \mathbf{q}(t) \in Q_{free}, t \in [0, \tau] \\ & && \mathbf{q}(0) = \bar{\mathbf{q}}_0, \mathbf{q}(\tau) = \bar{\mathbf{q}}_f \end{aligned} \quad (2)$$

Finally, the database is generated by considering distinct pairs of initial and final states $(\bar{\mathbf{q}}_0^i, \bar{\mathbf{q}}_f^i), i = 1, \dots, N$, and computing the corresponding set of motion primitives $\mathcal{Z} = \{\mathbf{z}_i^* = (\mathbf{q}_i^*(t), \mathbf{u}_i^*(t), \tau_i^*), i = 1, \dots, N\}$ with the associated costs $\mathcal{C} = \{C_i^*, i = 1, \dots, N\}$.

Note that, thanks to the translation invariance property introduced in Remark 1, the size of the database can be kept small while covering a wide range of the space where the robot is moving. Indeed, one can, without loss of generality, set the initial position $\bar{\boldsymbol{\pi}}_0$ to $\bar{\boldsymbol{\pi}}_0 = 0$ when building the

database, and recover the optimal trajectory for an arbitrary initial position π_0 by simply centring the coordinate system in π_0 .

Example 1 Consider, as an example, a planning problem for a unicycle robot characterised by a 3D search space (x, y, θ) , including the position (x, y) and the orientation θ , and by a 2D actuation space (v, ω) , constituted by the linear velocity v and the angular velocity ω .

Motion primitives are computed solving the following TPBVP

$$\begin{aligned} & \underset{v(\cdot), \omega(\cdot), \tau}{\text{minimize}} && \int_0^\tau \left(1 + 0.5v(t)^2 + 0.5\omega(t)^2 \right) dt \\ & \text{subject to} && \dot{x}(t) = v(t) \cos(\theta(t)) \\ & && \dot{y}(t) = v(t) \sin(\theta(t)) \\ & && \dot{\theta}(t) = \omega(t) \\ & && v(t) \in [0, 2], t \in [0, \tau] \\ & && \omega(t) \in [-2, 2], t \in [0, \tau] \\ & && x(0) = \bar{x}_0, y(0) = \bar{y}_0, \theta(0) = \bar{\theta}_0 \\ & && x(\tau) = \bar{x}_f, y(\tau) = \bar{y}_f, \theta(\tau) = \bar{\theta}_f \end{aligned}$$

for different initial and final poses.

Figure 1 shows a subset of these motion primitives, characterized by trajectories starting from $\bar{x}_0 = \bar{y}_0 = \bar{\theta}_0 = 0$.




Fig. 1 A subset of the motion primitives computed for a 3D search space (x, y, θ) . Red dots correspond to the initial and final positions, and black lines represent the resulting trajectories for different final orientations (θ) (Color figure online)




Fig. 2 “Reference trajectory” corresponding to (q_0, q_f) (red solid line), and symmetric trajectories defined by the pair of boundary conditions $(q_0, q_f^1), (q_0, q_f^2), (q_0, q_f^3)$ (Color figure online)

As can be seen in Fig. 1, with the dynamical system and cost function considered in this example, motion primitives, corresponding to boundary conditions that are symmetric with respect to the x -axis, are symmetric. A further analysis reveals that the same property holds for the y -axis as well, and that symmetric primitives are characterized by the same cost.

Remark 2 When for the considered dynamical system and cost criterion stronger invariance properties hold, like the axis symmetry in Example 1, the size of the database can be further reduced by storing only a few “reference trajectories”, and generating all the others using the invariance transformation.

Figure 2 shows an example, referred to the TPBVP considered in Example 1, where the trajectories represented by the pairs $(q_0, q_f^1), (q_0, q_f^2), (q_0, q_f^3)$ are characterised by the same cost and can be easily mapped to a “reference trajectory” corresponding to the boundary value pair (q_0, q_f) . In this case, the size of the database can be further reduced storing only the “reference trajectory”.

3.2 Search space design

In order to take advantage of the pre-computed database of motion primitives in sampling-based planning, the search space of the planner has to be defined appropriately so that every time the planner needs to connect two nodes, the corresponding optimal trajectory can be found in the database. To guarantee that this is indeed the case, the search space of the planner is uniformly gridded as the region where motion primitives are built. The translation invariance property¹ can then be exploited, as any optimal trajectory which connects a

¹ This approach can be easily extended in case stronger invariance properties hold.

pair of initial and final states (in the discretized search space) can be computed by first shifting the initial and final states so that the initial robot position corresponds to the zero position, then picking a suitable motion primitive in the database, and finally shifting the motion primitive so as to get back to the original reference coordinate system. Translation invariance, jointly with uniform gridding, allow a reduced number of motion primitives to cover the entire (discretized) search space. Note that, as the resolution of the database and the uniform grid size of the search space are coupled, we often use these two terms interchangeably.

Determining optimal state space discretization depends on the specific application, and is out of the scope of this work. We shall assume here that the state space grid should include the initial state \mathbf{q}_0 , at least a grid point in the goal region, and few points in the free space Q_{free} . Moreover, in order to find a solution that reaches the goal region, the algorithm should be able to search within all homotopy classes that are feasible given the robot footprint. In other words, one should be able to represent in the grid space all sets of trajectories in the continuous state space that can be obtained applying a smooth transformation and lead to the goal region. Missing a homotopy class could highly deteriorate performance in terms of achieved cost.

In Sect. 5 some numerical examples are provided, in which different grids are used for solving the same case study and results are compared.

Algorithm 1: RRT*_MotionPrimitives

```

1  $Q_T \leftarrow \{\mathbf{q}_0\}, E_T \leftarrow \emptyset, n \leftarrow 1$ 
2 while  $n \leq N$  do
3    $\mathbf{q}_{rand} \leftarrow \text{SAMPLE}(Q_{free})$ 
4    $Q_{near} \leftarrow \text{NEAR\_NODES}(\mathbf{q}_{rand})$ 
5    $\mathbf{q}_{best} \leftarrow \text{EXTEND}(Q_{near}, \mathbf{q}_{rand})$ 
6   if  $\mathbf{q}_{rand} \notin Q_T \wedge \mathbf{q}_{best} \neq \emptyset$  then
7      $Q_T \leftarrow Q_T \cup \{\mathbf{q}_{rand}\}$ 
8      $E_T \leftarrow E_T \cup \{(\mathbf{q}_{best}, \mathbf{q}_{rand})\}$ 
9      $n \leftarrow n + 1$ 
10     $E_T \leftarrow \text{REWIRE}(Q_T, E_T, \mathbf{q}_{rand}, Q_{near})$ 
11  else if  $\mathbf{q}_{rand} \in Q_T$  then
12     $\mathbf{q}_{prev} \leftarrow \text{PARENT}(\mathbf{q}_{rand})$ 
13    if  $\mathbf{q}_{best} \neq \mathbf{q}_{prev}$  then
14       $E_T \leftarrow (E_T \setminus \{(\mathbf{q}_{prev}, \mathbf{q}_{rand})\}) \cup \{(\mathbf{q}_{best}, \mathbf{q}_{rand})\}$ 
15       $E_T \leftarrow \text{REWIRE}(Q_T, E_T, \mathbf{q}_{rand}, Q_{near})$ 
16 return  $Q_T, E_T$ 

```

3.3 Motion planning

This section introduces RRT*_MotionPrimitives (Algorithm 1), the proposed variant of RRT* integrating the database of motion primitives for the computation of a collision-free optimal trajectory (cf. Sect. 2).

As in RRT*, RRT*_MotionPrimitives is based on the construction of a random tree $T = (Q_T, E_T)$ where $Q_T \subset Q$ is the set of nodes, and E_T is the set of edges. Nodes are states and edges are optimal trajectories, each one connecting a pair of origin and destination nodes and solving the TPBVP in (2).

The tree T is expanded for a maximum number of iterations N , defined by the user, starting from $Q_T = \{\mathbf{q}_0\}$, where $\mathbf{q}_0 \in Q_{free}$ is the initial state, and $E_T = \emptyset$, as described in Algorithm 1.

Every node $\mathbf{q} \in Q_T$ is connected to \mathbf{q}_0 via a single sequence of intermediate nodes $\mathbf{q}_i \in Q_T, i = 1, \dots, n-1, n \leq N$, and associated edges $e_i = e_{\mathbf{q}_i, \mathbf{q}_{i+1}} \in E_T, i = 0, 1, \dots, n-1$, with $\mathbf{q}_n = \mathbf{q}$.

One can then associate to this sequence a cost $C(\rightarrow \mathbf{q}_n)$ given by

$$C(\rightarrow \mathbf{q}_n) = \sum_{i=0}^{n-1} C(e_i)$$

where $C(e_i)$ denotes the cost associated with edge $e_i \in E_T$ and computed as in (2).

Algorithm 2: NEAR_NODES

```

1  $Q_{near} \leftarrow \emptyset$ 
2  $\pi_{rand} \leftarrow \text{GetPosition}(\mathbf{q}_{rand})$ 
3 for  $\forall \mathbf{q} \in Q_T$  do
4   if  $\mathbf{q} \in \text{BoundingBox}(\pi_{rand})$  then
5     if  $C(e_{\mathbf{q}_{rand}, \mathbf{q}}) \leq l(n) \vee C(e_{\mathbf{q}, \mathbf{q}_{rand}}) \leq l(n)$  then
6        $Q_{near} \leftarrow Q_{near} \cup \{\mathbf{q}\}$ 
7 return  $Q_{near}$ 

```

Tree growing is based on four main steps—random sampling, finding near nodes, extending the tree, and rewiring—that are described in the following.

Random sampling

A random state \mathbf{q}_{rand} is sampled from the free state space Q_{free} according to a uniform distribution by SAMPLE(Q_{free}). Unlike the original RRT* algorithm, however, the node is not sampled from the continuous state space, but from its discretization according to a uniform grid. For this reason, there is also a non zero probability that the same state \mathbf{q}_{rand} is sampled again in the next iterations of the algorithm.

Near nodes (Algorithm 2)

In RRT* a random state \mathbf{q}_{rand} can be connected only to a node that is within the set of its near nodes.

For Euclidean cost metrics, the set of near nodes is defined as a d -dimensional ball centered at \mathbf{q}_{rand} of radius

$$\gamma_{ball} = \gamma_{RRT^*} (\log(n)/n)^{1/d}$$

where n is the tree cardinality at the current iteration of the algorithm and γ_{RRT^*} is a suitable constant selected as

$$\gamma_{RRT^*} > 2 (1 + 1/d)^{1/d} (\mu(Q_{free})/\zeta_d)^{1/d}$$

$\mu(Q_{free})$ and ζ_d denoting the volume of the free configuration space and of the unit ball, respectively, in a d -dimensional Euclidean space.

For non-Euclidean cost metrics, the distance between two states is represented by the optimal cost of the trajectory that connects them, and near nodes are selected from a set of reachable states, Q_{reach} , defined as the set of states that can be reached from \mathbf{q}_{rand} or that can reach \mathbf{q}_{rand} with a cost that satisfies some threshold value. More specifically,

$$Q_{reach} = \{\mathbf{q} \in Q : C(e_{\mathbf{q}_{rand}, \mathbf{q}}) \leq l(n) \vee C(e_{\mathbf{q}, \mathbf{q}_{rand}}) \leq l(n)\} \tag{3}$$

where $e_{\mathbf{q}_i, \mathbf{q}_j}$ denotes the edge from \mathbf{q}_i to \mathbf{q}_j , and $l(n)$ is a cost threshold that decreases over the iterations of the algorithm as $l(n) = \gamma_l (\log n/n)$ such that a ball of volume $\gamma^d (\log n/n)$ is contained within Q_{reach} , where γ_l and γ are suitable constants (see Karaman and Frazzoli 2010, for further details).

In `RRT*_MotionPrimitives`, however, the set Q_{reach} has to be further constrained to ensure that there is a pre-computed trajectory in the database for each connection in the set of near nodes, that is thus redefined as follows

$$Q_{near} = Q_{reach} \cap \text{BoundingBox}(\pi_{rand})$$

where $\text{BoundingBox}(\pi_{rand})$ denotes the box of grid points in the state space Q adopted for the database construction, with the origin of the reference coordinate system shifted from $\tilde{\pi}_0 = 0$ to π_{rand} , which is the robot position associated to state \mathbf{q}_{rand} and obtained by using `GetPosition`. If Q_{near} occurs to be an empty set, the algorithm continues to the next iteration selecting a new \mathbf{q}_{rand} .

Extend (Algorithm 3)

The tree is extended to include \mathbf{q}_{rand} by selecting the node $\mathbf{q}_{best} \in Q_T$ such that the edge $e_{\mathbf{q}_{best}, \mathbf{q}_{rand}}$ connects \mathbf{q}_{rand} with a minimum cost collision free trajectory.

\mathbf{q}_{best} is determined as follows

$$\mathbf{q}_{best} = \underset{\mathbf{q} \in Q_{feasible}}{\operatorname{argmin}} C(\rightarrow \mathbf{q}) + C(e_{\mathbf{q}, \mathbf{q}_{rand}})$$

where $Q_{feasible} \subseteq Q_{near}$ is the set of nodes \mathbf{q} that belong to Q_{near} and such that the trajectory connecting \mathbf{q} to \mathbf{q}_{rand} is collision free, i.e.,

$$Q_{feasible} = \{\mathbf{q} \in Q_{near} | \text{CollisionFree}(e_{\mathbf{q}, \mathbf{q}_{rand}}) = 1\}$$

where $\text{CollisionFree} : E_T \rightarrow \{0, 1\}$ is a function that returns 1 for an edge that is collision free, 0 otherwise.

The selection of an edge from the database, connecting \mathbf{q} to \mathbf{q}_{rand} , is a peculiarity of `RRT*_MotionPrimitives`, and is performed by the `FindTrajectory` function as follows:

1. a translation is applied to the pair of initial and final states $(\mathbf{q}, \mathbf{q}_{rand})$ (Fig. 3a), obtaining the normalized pair $(\tilde{\mathbf{q}}, \tilde{\mathbf{q}}_{rand})$, such that the resulting $\tilde{\mathbf{q}}$ has the position $\tilde{\pi}$ corresponding to the null vector (Fig. 3b);
2. a query is executed on the database to look for the trajectory $\mathbf{z}_i \in Z$ and the cost $C(e_{\tilde{\mathbf{q}}, \tilde{\mathbf{q}}_{rand}}) \in \mathbb{R}$;
3. the inverse of the previous translation is applied in order to recover the trajectory connecting the actual pair of boundary values $(\mathbf{q}, \mathbf{q}_{rand})$, determining the edge $e_{\mathbf{q}, \mathbf{q}_{rand}}$ (Fig. 3c).

At the end of this procedure, if \mathbf{q}_{rand} is not already in the tree, then it is added to the tree together with the minimum cost edge, i.e., Q_T is replaced by $\{\mathbf{q}_{rand}\} \cup Q_T$ and E_T by $\{e_{\mathbf{q}_{best}, \mathbf{q}_{rand}}\} \cup E_T$ (see steps 7 and 8 in Algorithm 1). If \mathbf{q}_{rand} is already in the tree and if the computed \mathbf{q}_{best} is different from the current parent node, \mathbf{q}_{prev} , of \mathbf{q}_{rand} , given by `PARENT`(\mathbf{q}_{rand}), then the previous edge, $e_{\mathbf{q}_{prev}, \mathbf{q}_{rand}}$, is replaced by the new edge, $e_{\mathbf{q}_{best}, \mathbf{q}_{rand}}$ (see step 14 in Algorithm 1).

Algorithm 3: EXTEND

```

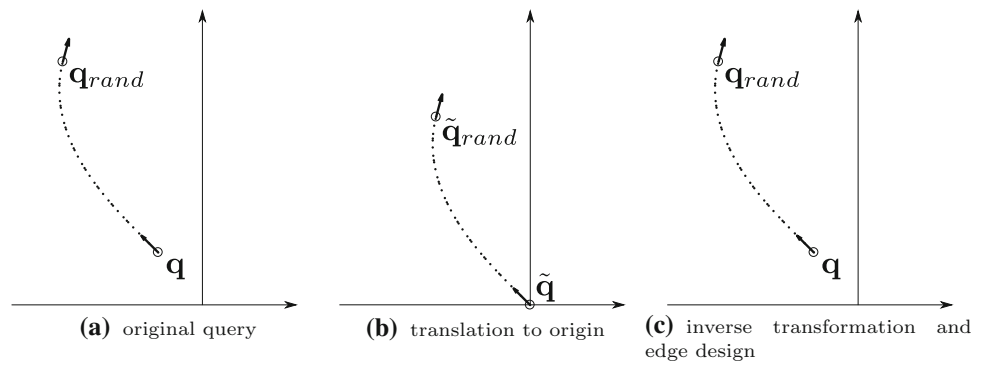
1  $\mathbf{q}_{best} \leftarrow \emptyset, e_{best} \leftarrow \emptyset, c_{best} \leftarrow \infty$ 
2 for  $\mathbf{q} \in Q_{near}$  do
3    $z, C(e_{\mathbf{q}, \mathbf{q}_{rand}}) \leftarrow \text{FindTrajectory}(\mathbf{q}, \mathbf{q}_{rand})$ 
4   if  $C(e_{\mathbf{q}, \mathbf{q}_{rand}}) < c_{best}$  then
5     if  $\text{CollisionFree}(e_{\mathbf{q}, \mathbf{q}_{rand}})$  then
6        $\mathbf{q}_{best} \leftarrow \mathbf{q}$ 
7        $e_{best} \leftarrow z$ 
8        $c_{best} \leftarrow C(e_{\mathbf{q}, \mathbf{q}_{rand}})$ 
9 return  $\mathbf{q}_{best}$ 

```

Rewiring (Algorithm 4)

In order to ensure that all node pairs are connected by an optimal sequence of edges, every time a new node \mathbf{q}_{rand} is added to the tree, a check is performed to verify if an already existing node can be reached from this newly added node with a smaller cost.

Fig. 3 Steps involved in the coordinate transformation within the FindTrajectory procedure



Algorithm 4: REWIRE

```

1 for  $q \in Q_{near}$  do
2    $z, C(e_{q_{rand},q}) \leftarrow \text{FindTrajectory}(q_{rand}, q)$ 
3   if  $C(\rightarrow q_{rand}) + C(e_{q_{rand},q}) < C(\rightarrow q)$  then
4     if  $\text{CollisionFree}(e_{q,q_{rand}})$  then
5        $q_{parent} \leftarrow \text{Parent}(q)$ 
6        $e_{prev} \leftarrow e_{q_{parent},q}$ 
7        $e \leftarrow e_{q_{rand},q}$ 
8        $E_T = \{E_T \setminus e_{prev}\} \cup \{e\}$ 
9 return  $E_T$ 

```

Therefore, $\forall q \in Q_{near}$ if $e_{q_{rand},q}$ is collision free, and the following conditions hold

$$C(e_{q_{rand},q}) \leq l(n)$$

$$C(\rightarrow q_{rand}) + C(e_{q_{rand},q}) \leq C(\rightarrow q)$$

the tree is rewired, i.e.,

$$E_T \leftarrow \{E_T \setminus e_{prev}\} \cup \{e_{q_{rand},q}\}$$

where e_{prev} is the previous edge connecting the node q to the tree.

Termination and best sequence selection

After the maximum number of iterations is reached the procedure to build the tree terminates. The best trajectory is selected as the node sequence reaching the goal region with the minimum cumulative cost C . Note that, using a discretized search space limits the number of nodes that can be sampled, once all of them have been sampled the tree cardinality does not increase any more, but the algorithm can still continue updating the edges to ensure that each node is connected with the best possible parent node.

4 Completeness and optimality analysis

In this section, probabilistic completeness of the proposed planning algorithm and optimality of the solution are discussed. Furthermore, some results to assess how close the solution obtained using a discretized state space and motion primitives is to the optimal trajectory computed considering a continuous state space are provided.

Let Q^Δ define the set of grid points that represent the discretized state space² and similarly $Q_{free}^\Delta := Q^\Delta \cap Q_{free}$ represents the free discrete state space. Assuming that the discretization step size is chosen properly, then, the collection of all grid points $q \in Q_{free}^\Delta$ that can be reached from q_0 by concatenating a sequence of motion primitives in Q_{free} is a non empty set. We shall denote this set as V_{free}^Δ and its cardinality as N^Δ . Note that the end points of the concatenated motion primitives are grid points in Q_{free}^Δ and hence they belong to V_{free}^Δ .

Let $\mathcal{G}_{free}^\Delta = (V_{free}^\Delta, E_{free}^\Delta)$ be a graph where the set of nodes is given by V_{free}^Δ defined before and the set of edges E_{free}^Δ is the collection of all the (possibly translated) motion primitives iteratively built as follows: starting from q_0 consider all the (translated) motion primitives that lie in Q_{free} and connect q_0 to all possible grid points in Q_{free}^Δ , and, then, continue with the same strategy for all of the newly reached grid points iteratively until it is not possible to further expand the graph.

Finally, Q_{goal}^Δ denotes the set of those grid points of V_{free}^Δ that belong to Q_{goal} . The motion planning problem using the grid representation admits a solution if Q_{goal}^Δ is not empty since this means that there exists a way of reaching a state in Q_{goal} starting from q_0 with the available motion primitives. In the following derivations we assume that $Q_{goal}^\Delta \neq \emptyset$.

Similarly to RRT*, RRT*_MotionPrimitives generates a tree T based on the random samples extracted from Q_{free}^Δ . However, unlike RRT*, the nodes and edges added to

² In this section, a Δ superscript is used to denote all variables that are associated with the grid state space, so as to distinguish them from their continuous state space counterpart.

the tree belong respectively to V_{free}^Δ and E_{free}^Δ , so that the obtained tree T is a sub-graph of $\mathcal{G}_{free}^\Delta$, i.e., $T \subset \mathcal{G}_{free}^\Delta$. The subscript i is used to denote the generated tree and the cost of the lowest cost trajectory represented in that tree after i -th iterations, i.e., T_i and c_i , respectively.

Since $Q_{goal}^\Delta \neq \emptyset$, then, there exists at least an *optimal branch* of $\mathcal{G}_{free}^\Delta$ composed of the ordered sequence of nodes

$$S^* := \{\mathbf{q}_0^*, \mathbf{q}_1^*, \dots, \mathbf{q}_k^*\},$$

which represents a *resolution optimal Δ -trajectory*, a minimum cost trajectory that starts at the initial state \mathbf{q}_0 and ends in the goal region, i.e., $\mathbf{q}_k^* \in Q_{goal}^\Delta$, such that $\mathbf{q}_j^* \in V_{free}^\Delta$ and $e_{\mathbf{q}_{j-1}^*, \mathbf{q}_j^*} \in E_{free}^\Delta$, for any $j = 1, \dots, k$. Let $c^{*\Delta}$ denote the cost of this optimal trajectory, i.e., $C(\rightarrow \mathbf{q}_k^*) = c^{*\Delta}$, which is named *resolution optimal Δ -cost*.

The goal of `RRT*_MotionPrimitives` can then be reformulated as that of generating a tree that contains an optimal branch S^* to reach Q_{goal}^Δ that is represented in $\mathcal{G}_{free}^\Delta$. Note that, one could in principle build G_{free}^Δ and apply an exhaustive search on it. However, this can still be an issue due to the combinatorial nature of the problem, in particular due to the branching caused by the dimensionality of the state space and the number of nodes contained in the graph.

In this section the quality of the solution obtained by `RRT*_MotionPrimitives` is analyzed by addressing the following questions:

1. resolution optimality: if there exists a resolution optimal Δ -trajectory S^* in $\mathcal{G}_{free}^\Delta$, then, is it possible to obtain such a trajectory?
2. asymptotic optimality: how close is the resolution optimal Δ -cost to the cost of the optimal trajectory, as the grid resolution increases and the grid converges to the continuous state space?

Theorem 1 *As the number of iterations goes to infinity the cost of the trajectory returned by `RRT*_MotionPrimitives` converges to the resolution optimal Δ -cost with a probability equal to 1, i.e.,*

$$\mathbb{P}\left(\left\{\lim_{i \rightarrow \infty} c_i = c^{*\Delta}\right\}\right) = 1.$$

Moreover, the expected number of iterations required to converge to c^{Δ} is upper bounded by $k|Q_{free}^\Delta|$, where k is the length of an optimal branch of $\mathcal{G}_{free}^\Delta$.*

Proof `RRT*_MotionPrimitives` returns the resolution optimal solution if it discovers an optimal branch, $S^* = \{\mathbf{q}_0^*, \mathbf{q}_1^*, \dots, \mathbf{q}_k^*\}$, defined on $\mathcal{G}_{free}^\Delta$. Assuming that the constant defining the cost threshold for selecting nearby nodes,

γ_l , is selected large enough (e.g. γ_l can be chosen so that Q_{reach} defined in (3) contains the `BoundingBox` (π_{rand}) for $n = N^\Delta$), as soon as $\{\mathbf{q}_0^*, \mathbf{q}_1^*, \dots, \mathbf{q}_{j-1}^*\}$ is a branch in the tree, $T = (Q_T, E_T)$, then the algorithm adds the edge $e_{\mathbf{q}_{j-1}^*, \mathbf{q}_j^*}$ to the tree in one of the following two ways:

1. any time \mathbf{q}_j^* is sampled it is connected to \mathbf{q}_{j-1}^* by the `EXTEND` procedure. In fact, as $e_{\mathbf{q}_{j-1}^*, \mathbf{q}_j^*}$ is part of the optimal sequence, there is no better way of reaching \mathbf{q}_j^* other than $e_{\mathbf{q}_{j-1}^*, \mathbf{q}_j^*}$.
2. if \mathbf{q}_j^* is sampled before \mathbf{q}_{j-1}^* and connected to some node \mathbf{q}_{prev} , such that $\mathbf{q}_j^* \in Q_T$ and $e_{\mathbf{q}_{prev}, \mathbf{q}_j^*} \in E_T$, the edge $e_{\mathbf{q}_{j-1}^*, \mathbf{q}_j^*}$ is selected any time \mathbf{q}_{j-1}^* is sampled, thanks to the `REWIRE` procedure.

This property allows to model the process of determining the sequence S^* as an *absorbing Markov chain* initialized at \mathbf{q}_0^* with \mathbf{q}_k^* as absorbing state and all intermediate \mathbf{q}_j^* , $j = 1, 2, \dots, k - 1$ that are transient states. There is a positive probability \mathcal{P}_j of advancing in the sequence, i.e., moving from \mathbf{q}_{j-1}^* to \mathbf{q}_j^* , and a probability $1 - \mathcal{P}_j$ of staying at the same state. Considering that at each iteration a new grid point is sampled from Q_{free}^Δ independently and according to a uniform distribution, each one has a probability $\frac{1}{|Q_{free}^\Delta|}$ of being extracted. Therefore, there is a probability $\mathcal{P}_j = \frac{1}{|Q_{free}^\Delta|}$ of advancing in the *Markov chain* and, since all states are transient states apart from \mathbf{q}_k^* which is the absorbing state, the probability that the process is absorbed by \mathbf{q}_k^* tends to 1 as i tends to infinity (Bertsekas and Tsitsiklis 2002). Moreover, there is a finite number of expected iterations, $k/\mathcal{P}_j = k|Q_{free}^\Delta|$, before the process is absorbed, i.e., before `RRT*_MotionPrimitives` returns the resolution optimal solution. \square

Clearly, the number of expected iterations increases with the depth of the solution, k , and the number of states represented in the grid. However, one advantage of `RRT*_MotionPrimitives`, as `RRT*` is the possibility of obtaining a solution rapidly and possibly improving its quality within the allowed computing time.

Corollary 1 *`RRT*_MotionPrimitives` is probabilistically resolution complete, as the number of iterations goes to infinity the algorithm will return a solution to the motion planning problem, if there exists one in $\mathcal{G}_{free}^\Delta$, with a probability 1.*

The remaining of this section deals with the relation between the resolution optimal Δ -trajectory returned by `RRT*_MotionPrimitives` as the number of iterations grows to infinity and the truly optimal trajectory in the continuous state space. To this purpose, we shall focus with the case when there are no actuation constraints and enforce

the following assumptions regarding the properties of the dynamical system (1) and the existence of a solution.

Assumption 1 The following properties hold for the dynamical system in (1)

- the system is small-time locally attainable (STLA)³;
- function $f(\cdot)$, representing the system dynamics, is Lipschitz continuous with Lipschitz constant \mathcal{K}_f ;
- function $C(\cdot)$, assigning a cost to an edge, satisfies the following Lipschitz-like continuity condition with Lipschitz constant \mathcal{K}_c :

$$|C(e_{\mathbf{q}_0, \mathbf{q}_1}) - C(e_{\tilde{\mathbf{q}}_0, \tilde{\mathbf{q}}_1})| \leq \mathcal{K}_c \left\| \begin{matrix} \mathbf{q}_0 \\ \mathbf{q}_1 \end{matrix} \right\| - \left\| \begin{matrix} \tilde{\mathbf{q}}_0 \\ \tilde{\mathbf{q}}_1 \end{matrix} \right\|$$

for each pair of edges $e_{\mathbf{q}_0, \mathbf{q}_1}$ and $e_{\tilde{\mathbf{q}}_0, \tilde{\mathbf{q}}_1}$.

Derivations in the rest of this section apply straightforwardly to state spaces that are Euclidean, and can be generalized to state spaces that are manifolds if the following assumption holds.

Assumption 2 The state space manifold of system (1) with d state variables is a subspace of the d -dimensional Euclidean space, \mathbb{R}^d , therefore can be locally treated as \mathbb{R}^d .

With a slight abuse of the previously introduced notation, in the rest of this section we use the term “trajectory” for the state space component of the tuple \mathbf{z} defined in Sect. 2. In order to compare the trajectory returned by `RRT*MotionPrimitives` and the optimal trajectory in the continuous state space, firstly, trajectories whose points are all away from obstacles by a certain distance are considered. For this reason, the definition of obstacle clearance of a trajectory, i.e., the minimum distance between obstacles and points belonging to the trajectory, has to be introduced.

Definition 1 (ϵ -obstacle clearance) Given a trajectory $\sigma(t)$, $t \in [0, T]$, if the ball $B_\epsilon(\sigma(t))$ of radius ϵ and centred at $\sigma(t)$ is strictly inside Q_{free} , for any $t \in [0, T]$, then, the ϵ -obstacle clearance property holds for σ .

Definition 2 (ϵ -free space) Let $\sigma(t) : [0, T] \rightarrow Q_{free}$ be a trajectory which has ϵ -obstacle clearance, the ϵ -free space along σ is given by

$$Q_\sigma^\epsilon := \bigcup_{t \in [0, T]} B_\epsilon(\sigma(t))$$

³ A system is STLA from a state $\mathbf{q} \in Q$ if $\forall T > 0$ the reachable set of states from q in time $0 < t \leq T$, $\mathcal{R}(\mathbf{q}, \leq T)$ contains a d -dimensional subset of \mathcal{N} , where \mathcal{N} denotes the set of neighborhood states in terms of Euclidean distance, (Choset 2005).




Fig. 4 An example of a trajectory σ_2 that is ϵ -similar to a trajectory σ_1

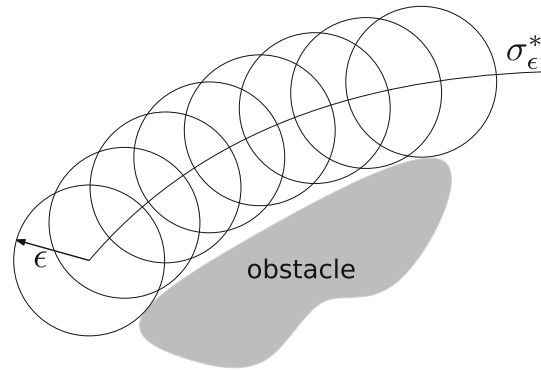


Fig. 5 σ_ϵ^* is the optimal trajectory with at least ϵ -obstacle and dynamic clearance

Then, a set of trajectories that are called to be ϵ -similar to σ can be introduced.

Definition 3 (ϵ -similarity) Any trajectory $\tilde{\sigma}(t) : [0, \tilde{T}] \rightarrow Q_{free}$ is said to be ϵ -similar to σ if it lies in the Q_σ^ϵ free-space, i.e., if

$$\tilde{\sigma}(t) \in Q_\sigma^\epsilon, t \in [0, \tilde{T}].$$

Figure 4 shows an example of a trajectory that is ϵ -similar to another one.

Note that having an ϵ -free space along a trajectory is not sufficient to guarantee the existence of ϵ -similar trajectories (ϵ -dynamic clearance), as this existence depends also on the properties of the dynamical system (1). The following definition relates the ϵ -similarity with the ϵ -free space through the system dynamics.

Definition 4 (ϵ -dynamic clearance) Given a trajectory $\sigma(t) : [0, T] \rightarrow Q_{free}$ which has ϵ obstacle clearance, if for any pair of time instants t_1, t_2 , such that $0 \leq t_1 < t_2 \leq T$, there exists a set of states inside a ball of radius $\alpha\epsilon$, with $0 < \alpha \leq 1$, centered at $\sigma(t_2)$ that are reachable from $\sigma(t_1)$ according to dynamics (1) without leaving the ϵ -free space around $\sigma(t)$, then σ has ϵ -dynamic clearance.

Let Σ_ϵ denote all the trajectories that solve the motion planning problem and have at least ϵ -obstacle and dynamic clearance. Let c_ϵ^* denote the minimum cost over all Σ_ϵ , which corresponds to the ϵ -optimal trajectory, σ_ϵ^* (Fig. 5). Note

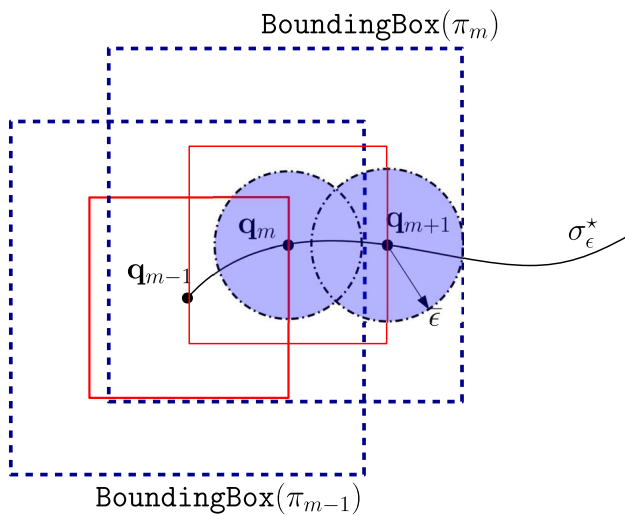


Fig. 6 Consecutive samples taken along the ϵ -optimal trajectory to design a $\{\mathcal{B}_\delta\}$ sequence. The corresponding $\text{BB}(\pi_m)$ are shown in red (Color figure online)

that, as ϵ tends to zero σ_ϵ^* converges to the truly optimal trajectory in the continuous state space. Due to the discretized nature of $\text{RRT}^*_\text{MotionPrimitives}$ it is not possible to converge to the optimal trajectory in the continuous state space, however in the following it is proven that the graph, $\mathcal{G}_{free}^\Delta$, contains an ϵ -similar trajectory, σ_ϵ^Δ , to the optimal trajectory σ_ϵ^* with a particular ϵ -clearance. Consequently, this result will be used to show that as the resolution of the grid increases, and as ϵ converges to zero, the resolution optimal Δ -cost will converge to the truly optimal cost, i.e., c^* .

Theorem 2 Let $\bar{\epsilon} > 0$ be smaller than half of the shortest side of the bounding box and such that the system (1) admits an ϵ -optimal trajectory for any $\epsilon \leq \bar{\epsilon}$, then,

- for a sufficiently fine gridding, $\mathcal{G}_{free}^\Delta$ contains an ϵ -similar trajectory to σ_ϵ^* , $\forall \epsilon \leq \bar{\epsilon}$,
- as the grid resolution increases, the resolution optimal Δ -cost converges to c^* .

Proof Let us fix ϵ , with $0 < \epsilon \leq \bar{\epsilon}$, and sample a set of states, $\{\mathbf{q}_m : m = 0, 1, \dots, M\}$, along the ϵ -optimal trajectory, σ_ϵ^* , starting from the initial state and ending in the final state, in such a way that a ball of radius $\bar{\epsilon}$ centered at sample \mathbf{q}_m would be contained in the $\text{BoundingBox}(\pi_{m-1})$ centered at the preceding sample \mathbf{q}_{m-1} and touching its boundary for $m = 1, \dots, M - 1$ (Fig. 6), till the final state is included within $\text{BoundingBox}(\pi_{M-1})$.

As σ_ϵ^* has ϵ -dynamic clearance, it is true that for any sample \mathbf{q}_m along the trajectory there exists a set of states within a ball of radius $\alpha\epsilon$ centered at \mathbf{q}_{m+1} that is reachable without leaving $Q_{\sigma_\epsilon^*}^\epsilon$, which is the ϵ -free space along σ_ϵ^* . From

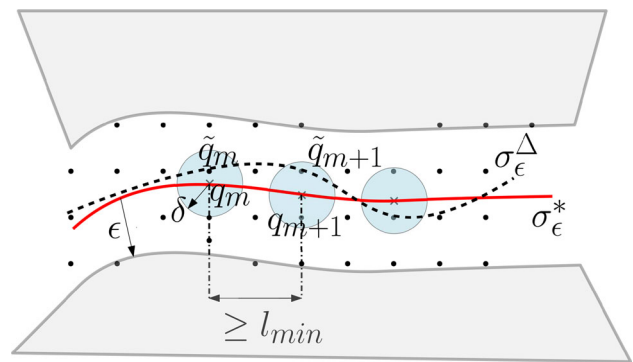


Fig. 7 Construction of the ball sequence $\{\mathcal{B}_\delta(\mathbf{q}_m)\}$. Tiling σ_ϵ^* with balls of radius δ centered at consecutive samples, there exist trajectories that are ϵ -similar to σ_ϵ^* represented on $\mathcal{G}_{free}^\Delta$. The cost of the trajectory connecting two consecutive samples is larger than or equal to l_{min}

Assumption 1, in particular from the Lipschitz continuity of the system dynamics, it follows that a sequence of non-overlapping balls, $\{\mathcal{B}_\delta(\mathbf{q}_m)\}$, centered at the samples along σ_ϵ^* and characterized by radius δ where

$$\delta = \frac{\alpha\epsilon}{2\mathcal{K}_f} \tag{4}$$

can be determined such that any state within $\mathcal{B}_\delta(\mathbf{q}_m)$ can reach any state in $\mathcal{B}_\delta(\mathbf{q}_{m+1})$ without leaving⁴ $Q_{\sigma_\epsilon^*}^\epsilon$ (see, Khalil 1996; Karaman and Frazzoli 2010). Assuming that the discretization is fine enough so that none of the $\mathcal{B}_\delta(\mathbf{q}_m)$ is empty, there exists a sequence of nodes $\{\tilde{\mathbf{q}}_m : m = 0, 1, \dots, M\}$ represented in $\mathcal{G}_{free}^\Delta$, such that $\tilde{\mathbf{q}}_m \in \mathcal{B}_\delta(\mathbf{q}_m), m = 0, 1, \dots, M, \tilde{\mathbf{q}}_0 = \mathbf{q}_0, \tilde{\mathbf{q}}_M$ belongs to the goal region and the corresponding trajectory σ_ϵ^Δ in $\mathcal{G}_{free}^\Delta$ is ϵ -similar to σ_ϵ^* (Fig. 7).

The cost of the optimal trajectory can be defined as the sum of the costs of the trajectories connecting consecutive $\mathcal{B}_\delta(\mathbf{q}_m)$ balls, i.e.,

$$c_\epsilon^* = \sum_{m=1}^M C(e_{\mathbf{q}_{m-1}, \mathbf{q}_m}), \tag{5}$$

and similarly the cost of the σ_ϵ^Δ is

$$c_\epsilon^\Delta = \sum_{m=1}^M C(e_{\tilde{\mathbf{q}}_{m-1}, \tilde{\mathbf{q}}_m}). \tag{6}$$

Subtracting (5) from (6), the difference can be written as

$$c_\epsilon^\Delta - c_\epsilon^* = \sum_{m=1}^M (C(e_{\tilde{\mathbf{q}}_{m-1}, \tilde{\mathbf{q}}_m}) - C(e_{\mathbf{q}_{m-1}, \mathbf{q}_m})).$$

⁴ It is assumed that $\mathcal{K}_f \geq 1$.

By the Lipschitz continuity of the cost function given in Assumption 1 it follows that

$$c_\epsilon^\Delta - c_\epsilon^* \leq \sum_{m=1}^M \mathcal{K}_c \delta,$$

where \mathcal{K}_c is the Lipschitz constant for the cost function. Therefore,

$$c_\epsilon^\Delta \leq c_\epsilon^* + \mathcal{K}_c M \delta. \quad (7)$$

We next derive an upper bound on the number of trajectory segments, M , using the lower bound on the cost of the trajectory connecting two consecutive samples, \mathbf{q}_{m-1} and \mathbf{q}_m . To this purpose, let us define a smaller bounding box centered at \mathbf{q}_m for each sample, which we shall denote as $\text{BB}(\boldsymbol{\pi}_m)$, such that the frontier of $\text{BB}(\boldsymbol{\pi}_m)$ intersects with σ_ϵ^* at \mathbf{q}_m (see Fig. 6) and all its sides are obtained by decreasing those of $\text{BoundingBox}(\boldsymbol{\pi}_m)$ of $2\bar{\epsilon}$ evenly. Then, we can define a minimum cost l_{min} among all the trajectories that reach the frontier of $\text{BB}(\boldsymbol{\pi}_0)$ from $\boldsymbol{\pi}_0$. Note that $l_{min} > 0$ since the length of the shortest side of $\text{BB}(\boldsymbol{\pi}_0)$ is larger than zero and we assumed non-zero cost for trajectories joining two different states. Then, the cost $C(e_{\mathbf{q}_{m-1}, \mathbf{q}_m})$ of the trajectory connecting \mathbf{q}_{m-1} and \mathbf{q}_m , for $m = 1, \dots, M - 1$, satisfies the following inequality

$$l_{min} \leq C(e_{\mathbf{q}_{m-1}, \mathbf{q}_m}). \quad (8)$$

Note that we have to treat separately the last trajectory segment connecting \mathbf{q}_{M-1} to \mathbf{q}_M since the existence of an $\bar{\epsilon}$ -ball centered at \mathbf{q}_M and contained within $\text{BoundingBox}(\boldsymbol{\pi}_{M-1})$ can not be guaranteed as \mathbf{q}_M is the final state which is fixed. Therefore, (8) may not hold for the last trajectory segment. We can then rewrite Eq. (5) as follows

$$c_\epsilon^* = \sum_{m=1}^{M-1} C(e_{\mathbf{q}_{m-1}, \mathbf{q}_m}) + C(e_{\mathbf{q}_{M-1}, \mathbf{q}_M})$$

and lower bound the cost of σ_ϵ^* as

$$c_\epsilon^* \geq (M - 1)l_{min} + 0.$$

Considering that the number of trajectory segments can be upper bounded by

$$M - 1 \leq \frac{c_\epsilon^*}{l_{min}},$$

from (7) it follows that

$$c_\epsilon^\Delta \leq \left(1 + \frac{\mathcal{K}_c \delta}{l_{min}}\right) c_\epsilon^* + \mathcal{K}_c \delta.$$

Substituting the definition of δ given in (4), we get

$$c_\epsilon^\Delta \leq \left(1 + \frac{\mathcal{K}_c \alpha \epsilon}{2\mathcal{K}_f l_{min}}\right) c_\epsilon^* + \frac{\mathcal{K}_c \alpha \epsilon}{2\mathcal{K}_f},$$

which shows that the cost of the ϵ -similar trajectory for a specific gridding is upper bounded by the cost of the ϵ -optimal trajectory and the ϵ -clearance. Now the cost c_ϵ^* of the optimal trajectory that can be obtained given a particular discretization satisfies

$$c_\epsilon^* \leq c_\epsilon^\Delta \leq \left(1 + \frac{\mathcal{K}_c \alpha \epsilon}{2\mathcal{K}_f l_{min}}\right) c_\epsilon^* + \frac{\mathcal{K}_c \alpha \epsilon}{2\mathcal{K}_f} \quad (9)$$

which provides an upper bound on the *resolution optimal* Δ -cost, c_ϵ^* .

As the resolution of the discretization increases and as the grid converges to the continuous state space, ϵ can converge to zero. Then, since \mathcal{K}_f and \mathcal{K}_c are constants and l_{min} is fixed (it depends on the upper bound $\bar{\epsilon}$ on ϵ), from (9) it follows that as ϵ goes to zero the *resolution optimal* Δ -cost converges to the cost c^* of the optimal trajectory, i.e.,

$$\lim_{\epsilon \rightarrow 0} c_\epsilon^* = c^*. \quad (10)$$

□

Theorem 2 states that when the number of nodes in $\mathcal{G}_{free}^\Delta$ goes to infinity, i.e., when the uniform gridding converges to the continuous state space, the resolution optimal Δ -cost converges to the cost c^* of the optimal trajectory in the continuous state space without gridding. By Theorem 1, it then follows that, when the uniform gridding converges to the continuous state space, as the number of iterations goes to infinity the cost of the trajectory returned by `RRT*_MotionPrimitives` converges to the cost of the optimal trajectory. Furthermore, Theorem 2 establishes an upper bound on the *resolution optimal* Δ -cost as a function of the cost of the ϵ -optimal trajectory and the ϵ -clearance such that as the grid resolution increases the upper bound decreases (see Eq. (9)). However, as shown in Theorem 1, as the number of discrete states increase, `RRT*_MotionPrimitives` will take more iterations to return the *resolution optimal* Δ -trajectory. Depending on the problem at hand, one should make the best compromise between the computing time (and size of the database) and the performance in terms of cost.

5 Numerical example

In this section a numerical example is presented to show the effectiveness of the proposed algorithm.

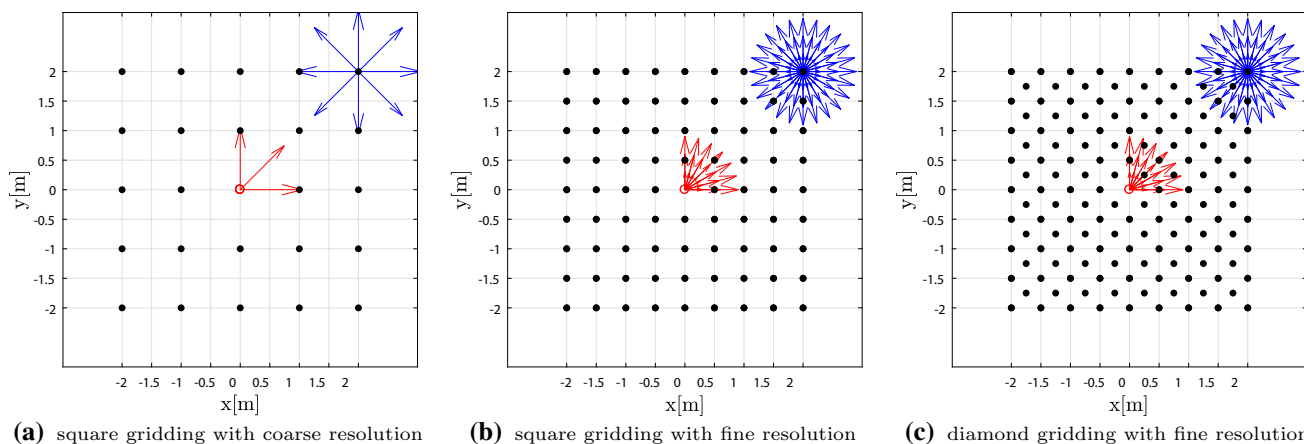


Fig. 8 The three grids used to generate the database. The red circle is the initial position, the blue dots the final ones. Red arrows represent the initial headings and velocities, blue arrows are the final ones. Different arrow sizes correspond to different velocities (Color figure online)

A 4D state-space (x, y, θ, v) representing a unicycle like robot moving on a planar surface is considered. The robot is described by the following equations

$$\begin{cases} \dot{x}(t) = v(t) \cos \theta(t) \\ \dot{y}(t) = v(t) \sin \theta(t) \\ \dot{\theta}(t) = w(t) \\ \dot{v}(t) = a(t) \end{cases} \quad (11)$$

where (x, y) is the position of the robot and θ the orientation with respect to a global reference frame, v and w are the linear and angular velocity, respectively. The control input is represented by $\mathbf{u} = [w, a]^T$, where w and a are the angular velocity and linear acceleration, respectively.

The motion primitives in the database are computed for each pair of initial and final state, $\mathbf{q}_0 = [x_0, y_0, \theta_0, v_0]$ and $\mathbf{q}_f = [x_f, y_f, \theta_f, v_f]$, solving the TPBVP in (2) for the differential equations given in (11) and the cost function

$$J(\mathbf{u}, \tau) = \int_0^\tau [1 + \mathbf{u}(t)^T R \mathbf{u}(t)] dt$$

that minimizes the total time of the trajectory τ , penalizing the total actuation effort with a weight $R = 0.5I_2$. The control variables a and w are bounded as $a \in [-3, 3] \text{ m/s}^2$, $w \in [-5, 5] \text{ rad/s}$.

TPBVPs are solved using MATLAB toolbox GPOPS (Patterson and Rao 2014), a nonlinear optimization tool based on the Gauss pseudo-spectral collocation method.

Three databases based on different grids have been considered. For all of them the initial state is characterised by the same position $(x_0, y_0) = (0, 0)$.

The first database is based on a coarse resolution uniform square grid (Fig. 8a), where $(x_f, y_f) \in [-2, 2] \times [-2, 2] \setminus \{(0, 0)\}$ and each square cell has a size of one




Fig. 9 Linear velocity and actuation profiles for the optimal trajectories in Figs. 10c (pink line), 11c (blue line) and 12c (green line). Red lines show the velocity and actuation limits (Color figure online)

meter. The initial orientation θ_0 is selected among three values $\{0, \pi/4, \pi/2\}$ rad, the final orientation θ_f can take 8 equally spaced values in the range $[0, 2\pi)$ rad. For the initial and final velocities, v_0 and v_f , a minimum and a maximum velocity of 1 m/s and 4 m/s is considered together with the zero velocity.

The second database is based on a fine resolution uniform square grid (Fig. 8b), where $(x_f, y_f) \in [-2, 2] \times [-2, 2] \setminus$

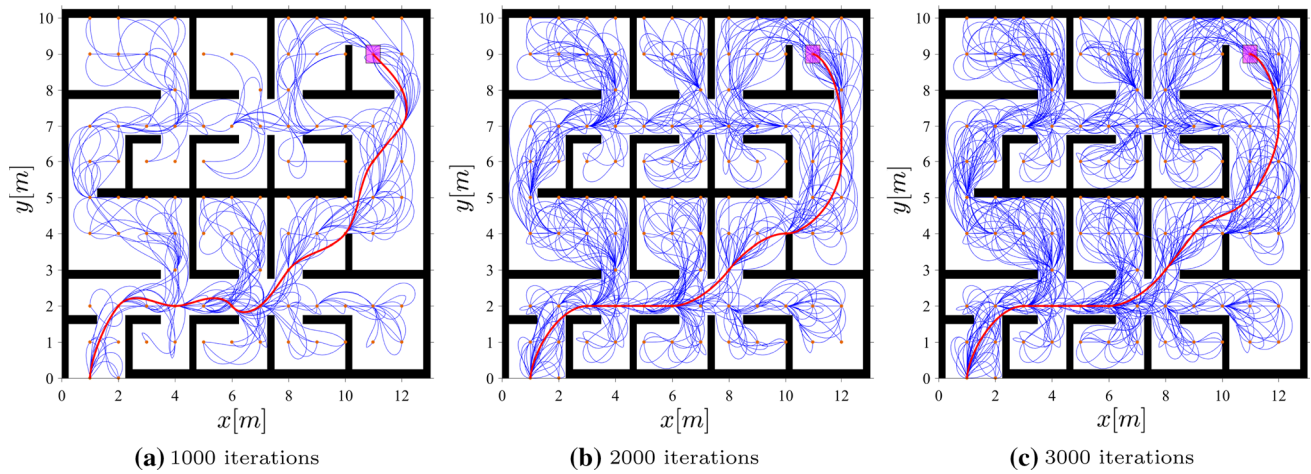


Fig. 10 Trajectories generated with an indoor map and a coarse resolution uniform square gridding for various number of iterations. Magenta square is the goal region, and the optimal trajectories are represented in red (Color figure online)

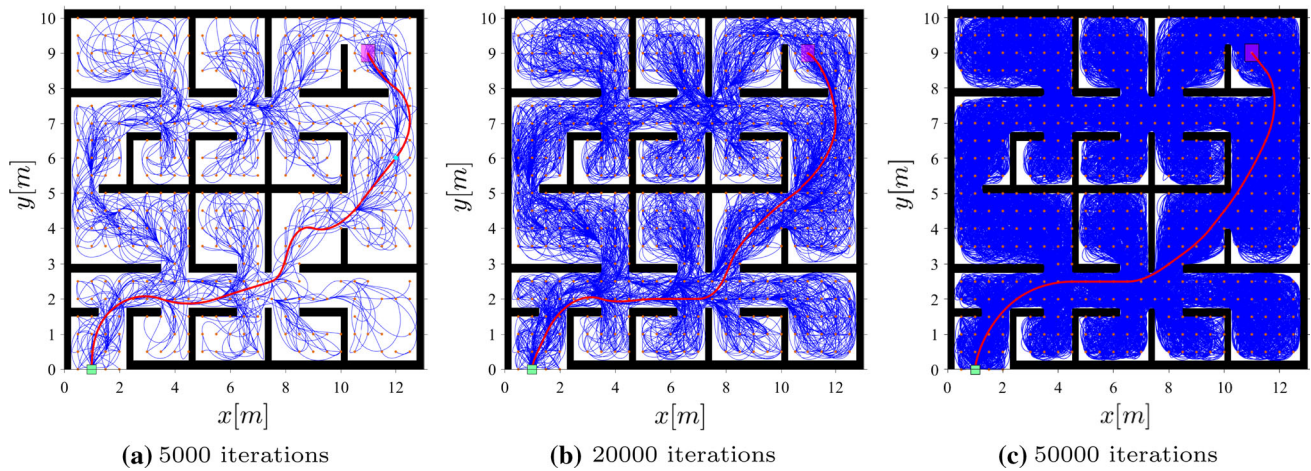


Fig. 11 Trajectories generated with an indoor map and a fine resolution uniform square gridding for various number of iterations. Magenta square is the goal region, and the optimal trajectories are represented in red (Color figure online)

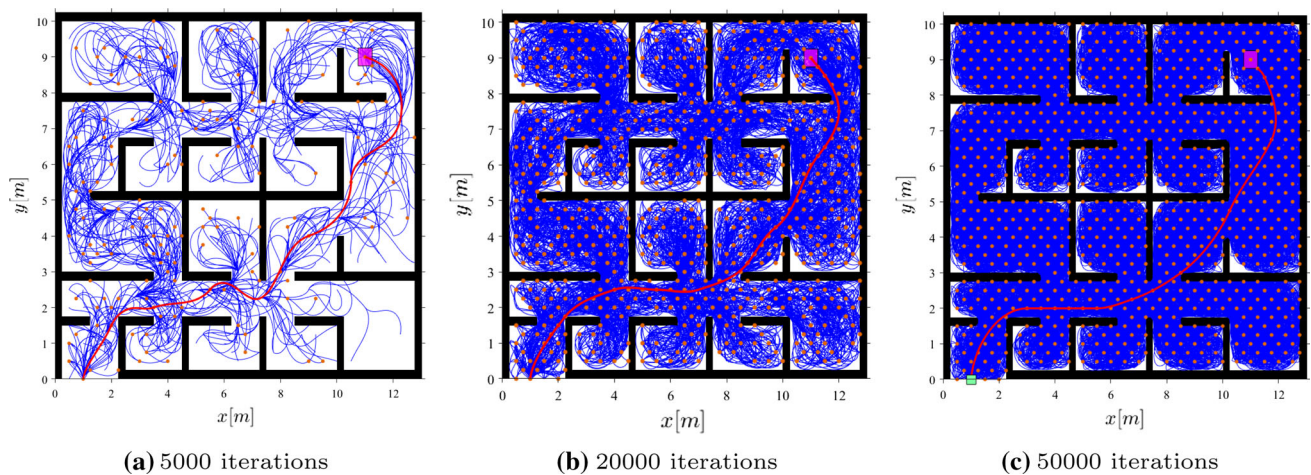


Fig. 12 Trajectories generated with an indoor map and a fine resolution uniform diamond gridding for various number of iterations. Magenta square is the goal region, and the optimal trajectories are represented in red (Color figure online)




Fig. 13 Cost with respect to the number of iterations. Coarse resolution square gridding (red line), fine resolution square gridding (green line) and diamond gridding (blue line) as a function of number of iterations (Color figure online)




Fig. 14 Tree cardinality with respect to the number of iterations, for coarse (red line) and fine (blue line) resolution square gridding (Color figure online)

$\{(0, 0)\}$ and each square cell has a size of half a meter. The initial and final orientations can take 24 equally spaced values in the range $[0, 2\pi)$ rad. The initial and final velocities are selected among 5 equally spaced values in the range $[0, 4]$ m/s.

Finally, the third database is based on a uniform diamond grid (Fig. 8c), characterised by the same initial and final states as the previous one, plus some additional final states at $(x_f, y_f) \in [-1.75, 1.75] \times [-1.75, 1.75]$ with a discretization step of half a meter in each direction. These additional states are characterized by the same orientation and velocity of the rest of the database.

Simulations are performed on an IntelCore i7@2.40 GHz personal computer with 8Gb RAM and the algorithm has been implemented in MATLAB.

An indoor map is considered (Figs. 10, 11, and 12), setting the robot initial pose at $(1, 0, \pi/2)$ with zero velocity. The goal area is defined as a square of half a meter side and centred at $(9, 11)$. The robot should stop at the end of the trajectory.

Figures 10, 11, and 12 show the trees and the optimal trajectories obtained for different number of iterations and for the three different gridding strategies. Correspondingly, Fig. 9 shows the velocity and the actuation profiles for the optimal trajectories computed with 3000 and 50,000 iterations, and reported in Figs. 10c, 11c and 12c, clearly demonstrating that the velocity constraint and the actuation bounds are satisfied. Note that the velocity profile that corresponds to the coarse resolution square gridding exhibits a




Fig. 15 Computation time with respect to the number of iterations, for coarse (red line) and fine (blue line) resolution square gridding (Color figure online)

jerky acceleration behaviour, due to the fact that the velocity at each node is constrained to be exactly one of the values in the database. This demonstrates that the velocity discretization step has to be accurately selected if a smoother velocity profile is required.

The same planning problem has been solved for 10 independent simulation runs. Figure 13 shows the average cost evolutions related to the coarse resolution, fine resolution and diamond gridding, as the number of iterations increases. As expected, the cost reduces increasing the number of iterations, and converges to the *resolution optimal Δ -cost*: once this minimum is achieved the solution will not further improve.

As can be easily seen, motion primitives computed using a denser grid provide lower cost plans. Moreover, the resolution optimal Δ -cost achieved using the fine resolution square and diamond grids are similar, demonstrating that the choice of the discretization step is strictly related to the specific problem. Finally, it is worth mentioning that as the resolution of the grid increases, the cardinality of Q_{free}^Δ increases as well, slowing down the convergence to the resolution optimal Δ -cost.

In order to assess the impact of the grid resolution on the size of the search space, in Table 1 we report the number of nodes corresponding to the grids of the three adopted databases (see Fig. 8), together with the corresponding minimum and maximum branching factors, i.e., the number of neighbors that each node is connected to. The computed number of nodes is an upper bound on the cardinality of G_{free}^Δ . Yet, from the figures in Table 1, it should be clear the combinatorial nature of the problem, which makes it hard building the whole graph of motion primitives and applying a graph search. As a matter of fact, the most commonly used lattice-based approaches use graph search algorithms that resort to some heuristic (see for example dynamic A* (D*) and anytime repairing A* (ARA*) by Stentz (1994) and Likhachev et al. (2008) respectively). Note that RRT*_MotionPrimitives does not need to adopt any heuristic, but if a smart heuristic were available, it could be

Table 1 Number of nodes of the state space grids in Fig. 8 corresponding to different resolutions together with the minimum and maximum branching factors

Type of grid	Number of nodes	Branching factor	
		Min	Max
Figure 8a	2460	324	370
Figure 8b	68,000	7300	9100
Figure 8c	131,000	12,000	15,000




Fig. 16 Histogram of the computation time for solving a single TPBVP, median of the values is represented with the red line. The histogram is determined based on 1000 boundary pairs randomly selected from the same subspace used for building the motion primitives (Color figure online)

integrated within `RRT*_MotionPrimitives` to speed up the search.

Figure 14 shows the cardinality of the tree for the coarse and fine resolution square griddings, as the number of iterations increases. Since the nodes that can be added to the tree are limited to the elements of the grid, the tree cardinality converges to the cardinality of \mathcal{G}_{free}^A : once this value is achieved no more nodes can be added.

Figure 15 shows the computation time evolution for the coarse and fine resolution griddings, as the number of iterations increases. Though `RRT*_MotionPrimitives` does not necessarily add a new node at each iteration, but can only change the existing connections, showing the computation time evolution with respect to the number of iterations allows to easily relate this quantity to the achieved cost, and thus the degree of sub-optimality of the planned trajectory.

It is also worth mentioning that the resulting computation time is promising, even for online replanning in the case of dynamic and partially known environments. Code optimisation and a C/C++ implementation can be considered for a further speed up.

To better emphasize the advantage of using a pre-computed database, we report in Fig. 16 the histogram of the computation time for solving a single TPBVP of the considered example using the GPOPS commercial numerical solver (Patterson and Rao 2014). As can be seen from this figure, it typically takes around 400 ms to get a solution for a single TPBVP while a trajectory can be extracted from the database in a time of the order of 0.01 ms (values ranged between 0.008

ms and 0.015 ms over 100 trials). Note that at each iteration of the standard RRT* algorithm, a set of TPBVPs that corresponds to the set of tentative trajectories connecting \mathbf{q}_{rand} to a set of nearby nodes has to be solved. When the TPBVP is not easy to be solved (like in the considered example), the applicability of RRT* to dynamic and partially known environments is hampered.

6 Conclusions

In this paper, a variant of RRT*, named `RRT*_MotionPrimitives`, that allows to introduce motion primitives in the RRT* planning framework is presented. In particular, a set of pre-computed trajectories, named motion primitives, is used to substitute the computationally challenging step of solving for a steering action. Then, in order to ensure that for any queried steering action a pre-computed trajectory exists, a grid representation of the state space has been introduced.

This newly conceived algorithm is supported by an accurate theoretical analysis, demonstrating the optimality and probabilistic completeness.

The performance of `RRT*_MotionPrimitives` has been verified in simulation, showing promising results in terms of quality of the planned trajectory and computation time, that is particularly important for an online usage in the case of dynamic environments that require repeated replanning. The results show also that as the grid size gets smaller, asymptotic optimality is achieved. Having a fine resolution, however, increases the size of the database and the number of iterations required to converge to the resolution optimal trajectory. Nevertheless, one advantage of adopting a sampling based approach is the possibility of computing a feasible though sub-optimal solution first, and then, in case more time is available, improve it. One should indeed choose the best compromise between computing time and performance, according to the application at hand.

Acknowledgements This work is supported by the European Commission under the project UnCoVerCPS with Grant Number 643921.

References

- Bertsekas, D. P., & Tsitsiklis, J. N. (2002). *Introduction to probability* (Vol. 1). Belmont, MA: Athena Scientific Belmont.
- Boggs, P. T., & Tolle, J. W. (1995). Sequential quadratic programming. *Acta Numerica*, 4, 1–51.
- Choset, H. M. (2005). *Principles of robot motion: Theory, algorithms, and implementation*. Cambridge: MIT press.
- Dolgov, D., Thrun, S., Montemerlo, M., & Diebel, J. (2010). Path planning for autonomous vehicles in unknown semi-structured environments. *The International Journal of Robotics Research*, 29(5), 485–501.

- Donald, B., Xavier, P., Canny, J., & Reif, J. (1993). Kinodynamic motion planning. *J ACM*, 40(5), 1048–1066. <https://doi.org/10.1145/174147.174150>.
- Gammell, J.D., Srinivasa, S.S., & Barfoot, T.D. (2014). BIT*: Batch informed trees for optimal sampling-based planning via dynamic programming on implicit random geometric graphs. Tech. rep., Tech. Report TR-2014-JDG006, ASRL, University of Toronto.
- Goretkin, G., Perez, A., Platt, R., & Konidaris, G. (2013). Optimal sampling-based planning for linear-quadratic kinodynamic systems. In: IEEE International Conference on Robotics and Automation, IEEE, (pp. 2429–2436).
- Houska, B., Ferreau, H., & Diehl, M. (2011). ACADO Toolkit—An open source framework for automatic control and dynamic optimization. *Optimal Control Applications and Methods*, 32(3), 298–312.
- hwanJeon, J., Karaman, S., & Frazzoli, E. (2011). Anytime computation of time-optimal off-road vehicle maneuvers using the RRT. In: IEEE Conference on Decision and Control and European Control Conference, (pp. 3276–3282).
- Karaman, S., & Frazzoli, E. (2010). Optimal kinodynamic motion planning using incremental sampling-based methods. In: IEEE Conference on Decision and Control, IEEE, (pp. 7681–7687).
- Karaman, S., & Frazzoli, E. (2011). Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30(7), 846–894.
- Karaman, S., & Frazzoli, E. (2013). Sampling-based optimal motion planning for non-holonomic dynamical systems. In: IEEE International Conference on Robotics and Automation, (pp. 5041–5047).
- Kavraki, L. E., Svestka, P., Latombe, J. C., & Overmars, M. H. (1996). Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4), 566–580.
- Khalil, H. K. (1996). *Nonlinear systems* (Vol. 2). New Jersey: Prentice-Hall.
- Kuderer, M., Sprunk, C., Kretschmar, H., & Burgard, W. (2014). Online generation of homotopically distinct navigation paths. In: IEEE International Conference on Robotics and Automation, (pp. 6462–6467).
- LaValle, S. M. (2006). *Planning algorithms*. Cambridge: Cambridge University Press.
- LaValle, S. M., & Kuffner, J. J., Jr. (2001). Randomized kinodynamic planning. *The International Journal of Robotics Research*, 20(5), 378–400.
- Li, Y., Littlefield, Z., & Bekris, K. E. (2016). Asymptotically optimal sampling-based kinodynamic planning. *The International Journal of Robotics Research*, 35(5), 528–564.
- Likhachev, M., & Ferguson, D. (2009). Planning long dynamically feasible maneuvers for autonomous vehicles. *The International Journal of Robotics Research*, 28(8), 933–945.
- Likhachev, M., Ferguson, D., Gordon, G., Stentz, A., & Thrun, S. (2008). Anytime search in dynamic graphs. *Artificial Intelligence*, 172(14), 1613–1643.
- Park, J., Karumanchi, S., & Iagnemma, K. (2015). Homotopy-based divide-and-conquer strategy for optimal trajectory planning via mixed-integer programming. *IEEE Transactions on Robotics*, 31(5), 1101–1115.
- Patterson, M. A., & Rao, A. V. (2014). GPOPS-II: A matlab software for solving multiple-phase optimal control problems using hp-adaptive Gaussian quadrature collocation methods and sparse nonlinear programming. *ACM Transactions on Mathematical Software*, 41(1), 1.
- Pearl, J. (1984). *Heuristics: intelligent search strategies for computer problem solving*. Boston: Addison.
- Perez, A., Platt, R., Konidaris, G., Kaelbling, L., & Lozano-Perez, T. (2012). LQR-RRT*: Optimal sampling-based motion planning with automatically derived extension heuristics. In: IEEE International Conference on Robotics and Automation, (pp. 2537–2542).
- Pivtoraiko, M., Knepper, R. A., & Kelly, A. (2009). Differentially constrained mobile robot motion planning in state lattices. *Journal of Field Robotics*, 26(3), 308–333.
- Reif, J.H. (1979). Complexity of the mover’s problem and generalizations. In: Annual Symposium on Foundations of Computer Science, (pp. 421–427).
- Stentz, A. (1994). Optimal and efficient path planning for partially-known environments. In: 1994 IEEE International Conference on Robotics and Automation, 1994. Proceedings., IEEE, (pp 3310–3317).
- Stoneman, S., & Lampariello, R. (2014). Embedding nonlinear optimization in RRT* for optimal kinodynamic planning. In: 2014 IEEE 53rd Annual Conference on Decision and Control (CDC), IEEE, (pp 3737–3744).
- Webb, D.J., & van den Berg, J. (2013). Kinodynamic RRT*: Asymptotically optimal motion planning for robots with linear dynamics. In: 2013 IEEE International Conference on Robotics and Automation (ICRA), IEEE, (pp. 5054–5061).

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



jectory optimization.

Basak Saccak was born in Ankara, Turkey. She received the B.Sc. degree in Mechatronics Engineering at Sabanci University, Turkey, in June 2010. Following that she received the M.Sc. degree in Automation Engineering and Ph.D. degree in Information Technology in 2013 and 2018 respectively from Politecnico di Milano, Italy. Currently, she is a postdoctoral researcher at the same university. Her research interests include industrial and mobile robotics, in particular motion planning and tra-



jectory optimization.

Luca Bascetta was born in Milano (Italy) in 1974. He received the Laurea degree cum laude in Computer Science Engineering and the Doctorate degree in Information Technology in 1999 and 2004, respectively, both from Politecnico di Milano, Italy. From 2015 he has been Associate Professor in Automatic Control at the School of Industrial and Information Engineering of Politecnico His research interests include autonomous vehicles, robotics, visual servoing and all the problems related to motion control in general. Professor Bascetta is a senior member of IEEE Robotics and Automation Society since 2017.



Gianni Ferretti was born in Cremona (Italy) in 1962. He received the 'Laurea' degree in Electronics Engineering in 1987 from Politecnico di Milano. From 1990 to 1998 he was Assistant Professor at the Department of Electronics and Computer Science of Politecnico di Milano and in 1994 he was a visiting researcher at the Laboratoire d'Automatique de Grenoble. From 1998 to 2005 he was Associate Professor at Politecnico di Milano where he is currently Full Professor in Systems and Control.

He is Vice Rector for the Cremona campus of Politecnico di Milano and teaches Fundamentals of Automatic Control, Digital Controls and Simulation Techniques and Tools. His principal research interests are: control of servomechanisms, modelling and simulation of mechatronic systems, force/position control of robots, modelling and simulation of automotive systems, modelling and control of biogas power plants. He is coauthor of more than 100 papers, most of which published in international journals or presented in international conferences, and was involved in many research projects with industries and institutions.



Maria Prandini received her laurea degree in Electrical Engineering (summa cum laude) from Politecnico di Milano (1994) and her Ph.D. degree in Information Technology from Università degli Studi di Brescia, Italy (1998). From 1998 to 2000 she was a postdoctoral researcher at the Department of Electrical Engineering and Computer Sciences, University of California at Berkeley. She also held visiting positions at Delft University of Technology (1998), Cambridge University (2000), University of California at Berkeley (2005), and Swiss Federal Institute of Technology Zurich (2006).

In 2002, she started as an assistant professor in systems and control at Politecnico di Milano, where she is currently an full professor. Her research interests include stochastic hybrid systems, randomized algorithms, constrained control, system abstraction and verification, nonlinear identification, distributed optimization, and the application of control theory to air traffic management and energy systems. She serves on the editorial board of Cyber Physical Systems, and previously of European Journal of Control, IEEE Transactions on Automatic Control, IEEE Transactions on Control Systems Technology and Nonlinear Analysis: Hybrid Systems. From 2013 to 2015, she has been editor for Electronic Publications of the IEEE CSS. She was elected member of the IEEE CSS Board of Governors for a triennium (2015–17), and CSS Vice-President for Conference Activities in 2016 and 2017.