

db-A*: Discontinuity-bounded Search for Kinodynamic Mobile Robot Motion Planning

Wolfgang Hönig, Joaquim Ortiz-Haro, and Marc Toussaint

Abstract—We consider time-optimal motion planning for dynamical systems that are translation-invariant, a property that holds for many mobile robots, such as differential-drives, cars, airplanes, and multirotors. Our key insight is that we can extend graph-search algorithms to the continuous case when used symbiotically with optimization. For the graph search, we introduce discontinuity-bounded A* (db-A*), a generalization of the A* algorithm that uses concepts and data structures from sampling-based planners. Db-A* reuses short trajectories, so-called motion primitives, as edges and allows a maximum user-specified discontinuity at the vertices. These trajectories are locally repaired with trajectory optimization, which also provides new improved motion primitives. Our novel kinodynamic motion planner, kMP-db-A*, has almost surely asymptotic optimal behavior and computes near-optimal solutions quickly. For our empirical validation, we provide the first benchmark that compares search-, sampling-, and optimization-based time-optimal motion planning on multiple dynamical systems in different settings. Compared to the baselines, kMP-db-A* consistently solves more problem instances, finds lower-cost initial solutions, and converges more quickly.

I. INTRODUCTION

Motion planning for robots with known kinodynamics remains challenging, especially when a time-optimal motion is desired. Consider the example in Fig. 1 of a simple dynamical model in 2D (unicycle, 3-dimensional state space and 2-dimensional control space). Finding the time-optimal solution is surprisingly challenging for state-of-the-art methods when constraining the control space to model a plane with a malfunctioning rudder, i.e., with a positive minimum speed and asymmetric angular velocity limits.

Current planning approaches are sampling-based, search-based, optimization-based, or hybrid. Each of these methods has their strengths and weaknesses. Sampling-based planners can find initial solutions quickly and have strong guarantees for convergence to an optimal solution. However, in practice the initial solutions are far from the optimum, the convergence rate is low, and the solutions typically require some post-processing. Search-based approaches can remedy those shortcomings by connecting precomputed trajectories, so-called *motion primitives*, using A* or related graph search algorithms. Yet, the seemingly strong theoretical guarantees only hold up to the selected discretization of the state space and the precomputed motions. Moreover, scaling this

The authors are with TU Berlin, Germany. {hoenig, joaquim.ortizdeharo, toussaint}@tu-berlin.de.

The research was funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) - 448549715 and by the German-Israeli Foundation for Scientific Research (GIF) grant I-1491-407.6/2019. Joaquim Ortiz-Haro thanks the International Max-Planck Research School for Intelligent Systems (IMPRS-IS) for the support.




Fig. 1. A kinodynamic planning problem, where a plane-like 2D robot with a malfunctioning rudder (no sharp right turns) has to move from the left configuration (filled blue) to the right configuration (blue outline) in minimum time. A) An initial solution by “stitching” motion primitives with bounded discontinuities (magnitude indicated by hue of red color). B) A refined solution using more primitives and a lower discontinuity bound δ . C) Final trajectory computed with optimization using B) as initial guess.

approach to higher dimensions has proved difficult and requires careful, frequently hand-crafted design of the motion primitives. This curse of dimensionality can be overcome by optimization-based planners, which scale polynomially rather than exponentially with the number of state dimensions. However, these planners are, in the general case, only locally optimal and thus require a good initial guess both for the trajectory and time horizon.

In this paper, we present a new approach for kinodynamic motion planning of mobile robots that combines key ideas and strengths of the aforementioned previous methods. We rely on a graph-search method, because it provides a theoretically grounded exploration/exploitation tradeoff, but we want to remedy its primary shortcoming of a predefined discretization, similar to sampling-based planning. The naive approach of simply increasing the number of primitives is intractable, due to the resulting infinite number of states and infinite branching factor. We solve this challenge with a combination of bounded-discontinuity search with nonlinear optimization. Introducing the discontinuity makes the search tractable: we can reuse the primitives and have a finite number of states to expand. While the resulting trajectory is not feasible, it can be used as initial guess of trajectory optimization that locally repairs the discontinuous trajectory into a valid trajectory. We execute search and optimization in an iterative fashion, where the value of the discontinuity bound decreases in every iteration. For large bounds, the search is very fast, but the optimizer might fail to find a valid solution. For very small bounds, the search requires a longer runtime, but the optimizer has an excellent initial guess. This combination results in an efficient anytime planner with probabilistic optimality guarantees.

More specifically, our first contribution is the introduction of *kMP-db-A**, a new kinodynamic motion planner that combines a novel search algorithm, discontinuity-bounded A* (db-A*), and trajectory optimization in an iterative fashion. Db-A* generalizes A* with ideas from sampling-based planning to obtain solution trajectories that may have discontinuities up to a user-specified bound. Our second contribution is the, to our knowledge, first benchmark that compares the three major kinodynamic motion planning techniques on the same problem instances with the identical objective of computing time-optimal trajectories. While we focus in our evaluation on the challenging case of time-optimality, our approach supports arbitrary cost functions.

II. PROBLEM DESCRIPTION

We consider a robot with state $\mathbf{x} = [\mathbf{x}^t, \mathbf{x}^r] \in \mathcal{X} \subset \mathbb{R}^{d_w} \times \mathbb{R}^{d_x - d_w}$, where the first d_w dimensions indicate the translation in the workspace ($d_w \in \{2, 3\}$) of the robot and the remaining $d_x - d_w$ dimensions may contain orientation or derivatives. The robot can be actuated by controlling actions $\mathbf{u} \in \mathcal{U} \subset \mathbb{R}^{d_u}$. We consider dynamics that are *translation invariant*, with

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}^r, \mathbf{u}), \quad (1)$$

where \mathbf{f} only depends on \mathbf{x}^r and not on \mathbf{x}^t . In order to employ gradient-based optimization, we assume that we can compute the Jacobian of \mathbf{f} with respect to \mathbf{x}^r and \mathbf{u} .

The robot is operating in a workspace $\mathcal{W} \subseteq \mathbb{R}^{d_w}$ that indicates the free space for safe navigation. The free state space then becomes $\mathcal{X}_{\text{free}} = \{\mathbf{x} = [\mathbf{x}^t, \mathbf{x}^r] \in \mathcal{X} | \mathbf{x}^t \in \mathcal{W}\}$.

Almost all generic kinodynamic motion planners assume a discrete-time formulation with zero-order hold, i.e., the applied action remains constant during a timestep. We can then frame the dynamics Eq. (1) as

$$\mathbf{x}_{k+1} \approx \text{step}(\mathbf{x}_k, \mathbf{u}_k) \equiv \mathbf{x}_k + \mathbf{f}(\mathbf{x}_k^r, \mathbf{u}_k)\Delta t, \quad (2)$$

using a small timestep Δt so that the Euler approximation holds sufficiently well.

Let $\mathbf{X} = \langle \mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T \rangle$ be a sequence of states sampled at times $0, \Delta t, \dots, T\Delta t$ and $\mathbf{U} = \langle \mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{T-1} \rangle$ be a sequence of actions applied to the system for times $[0, \Delta t), [\Delta t, 2\Delta t), \dots, [(T-1)\Delta t, T\Delta t)$. Then our goal of moving the robot from its start state to a goal state can be framed as the following optimization problem:

$$\begin{aligned} & \min_{\mathbf{U}, \mathbf{X}, T} J(\mathbf{U}, \mathbf{X}, T) & (3) \\ \text{s.t.} & \begin{cases} \mathbf{x}_{k+1} = \text{step}(\mathbf{x}_k, \mathbf{u}_k) & \forall k \in \{0, \dots, T-1\} \\ \mathbf{u}_k \in \mathcal{U} & \forall k \in \{0, \dots, T-1\} \\ \mathbf{x}_k \in \mathcal{X}_{\text{free}} & \forall k \in \{0, \dots, T\} \\ \mathbf{x}_0 = \mathbf{x}_s; \mathbf{x}_T = \mathbf{x}_f, \end{cases} \end{aligned}$$

where $\mathbf{x}_s \in \mathcal{X}$ is the start state and $\mathbf{x}_f \in \mathcal{X}$ is the goal state. The objective function J is application specific; we will focus on time-optimal trajectories, i.e., $J(\mathbf{U}, \mathbf{X}, T) = T\Delta t$.

Example 1: Consider a unicycle robot with state $\mathbf{x} = [x, y, \theta] \in \mathcal{X} = SE(2) \subset \mathbb{R}^2 \times \mathbb{R}^1$, i.e., x, y are the position and θ is the orientation. The actions are $\mathbf{u} =$

$[v, \omega] \in \mathcal{U} \subset \mathbb{R}^2$, i.e., the speed and angular velocity can be controlled directly. The dynamics are translation invariant: $\dot{\mathbf{x}} = [v \cos \theta, v \sin \theta, \omega]$. The choice of \mathcal{U} can make this low-dimensional problem challenging to solve. For example, Fig. 1 shows a plane-like case (positive minimum speed, i.e., $0.25 \leq v \leq 0.5$ m/s) with a malfunctioning rudder (asymmetric angular speed, i.e., $-0.25 \leq \omega \leq 0.5$ rad/s).

III. RELATED WORK

There are several conceptually different algorithmic approaches to solving kinodynamic motion planning problems.

Search-based approaches rely on existing methods for discrete path planning, such as A* and variants. The common approach is to generate short trajectories (*motion primitives*) using a state lattice (i.e., pre-specified discrete state components) [1]. Each primitive starts and ends at a grid cell and swept cells can be precomputed for efficient collision checking. Once motion primitives are computed, existing algorithms such as A* or the anytime variant ARA* can be employed without modification.

These methods can solve Eq. (3) if \mathbf{x}_s and \mathbf{x}_f fall within the chosen lattice and retain very strong theoretical guarantees on both optimality and completeness with respect to the chosen primitives. The major challenge is to select and compute good motion primitives, especially for high-dimensional systems [2], [3].

Sampling-based approaches build a tree \mathcal{T} rooted at the start state \mathbf{x}_s . During tree expansion, i) a random state \mathbf{x}_{rand} in the state space is sampled, ii) an existing state $\mathbf{x}_{\text{expand}} \in \mathcal{T}$ is selected, and iii) a new state \mathbf{x}_{new} is added with a motion that starts at $\mathbf{x}_{\text{expand}}$ and moves towards \mathbf{x}_{rand} . The classic version of this approach, *kinodynamic RRT* [4], is probabilistically complete when using the correct variant [5]. Asymptotic optimality can be achieved when planning in state-cost space (*AO-RRT*) [6], [7], [8] or by computing a sparse tree (*SST**) [9]. These methods rely on a distance function $d: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ and often use fast nearest neighbor data structures such as k-d trees for efficiency. The mentioned algorithms work without solving a two-point boundary value problem, which is computationally expensive.

Sampling-based approaches are designed to explore the state space as fast as possible and typically do not use a heuristic function, unlike search-based methods. The exploration/exploitation tradeoff is typically controlled by using goal-biasing instead. These approaches cannot solve Eq. (3) directly, because the probability to reach \mathbf{x}_f by sampling is zero. Instead, the goal constraint is typically reformulated to $\mathbf{x}_T \in \mathcal{X}_f$ using a goal region \mathcal{X}_f rather than a goal state \mathbf{x}_f .

Optimization-based approaches locally optimize an initial trajectory using the gradients of J , unlike the previous gradient-free methods. Dynamics, collision avoidance, goal constraints, and control limits are modeled with piece-wise differentiable functions. In *CHOMP*, Hamiltonian Monte Carlo is used to perturb local solutions [10], while *TrajectoryOpt* [11] and *GuSTO* [12], [13] rely on sequential convex programming (SCP). Trajectories can also be computed with optimal control solvers that rely on Differential Dynamic

Programming [14] or extend the linear quadratic regulator to nonlinear systems [15].

Both *STOMP* [16] and *KOMO* [17] use only the (geometric) state sequence \mathbf{X}^t as decision variables and support kinodynamic systems via constraints. All optimization-based approaches require an initial guess as a starting trajectory, but this guess does not need to be kinodynamically feasible.

These approaches can solve Eq. (3) directly for a differentiable J and given number of timesteps T . The observed solution quality is significantly higher (e.g., in terms of smoothness) compared to sampling-based or search-based approaches. Moreover, optimization-based approaches do not suffer from the curse of dimensionality directly, although higher dimensions might result in more local optima.

Hybrid approaches combine two or more ideas. One can combine search and optimization [18], search and sampling [19], [20], [21], or combine sampling and optimization [22]. For some dynamical systems, using insights from control theory for the motion planning can also be beneficial [23], [24], [25], but requires domain knowledge. Motion planning can also benefit from using machine learning for computational efficiency [26], [27].

Our approach relies on the fact that the dynamics are translation-invariant for many mobile robots. The most similar related works are a method that reuses edges within a sampling-based planning framework [21] and a search-based approach that has an enhanced duplicate detection [28]. Unlike those works we also include trajectory optimization and reuse precomputed and online computed motion primitives for fast convergence in practice. Moreover, our discontinuity bound δ is not a fixed user-specified value but converges to 0, which allows us to provide stronger theoretical guarantees.

IV. APPROACH

Our general approach is shown in Algorithm 1. We assume that we have access to a set of *motion primitives*, which are valid trajectories according to our dynamics.

Definition 1: A *motion primitive* is a tuple $\langle \mathbf{X}, \mathbf{U}, T, c \rangle$ with

$$\begin{aligned} \mathbf{X} &= \langle \mathbf{x}_0, \dots, \mathbf{x}_T \rangle & \mathbf{x}_0^t &= \mathbf{0}, \quad \mathbf{x}_k \in \mathcal{X} \\ \mathbf{U} &= \langle \mathbf{u}_0, \dots, \mathbf{u}_{T-1} \rangle & \mathbf{u}_k &\in \mathcal{U} \\ \mathbf{x}_{k+1} &= \text{step}(\mathbf{x}_k, \mathbf{u}_k) & \forall k \in \{0, \dots, T-1\} \\ c &= J(\mathbf{U}, \mathbf{X}, T). \end{aligned} \quad (4)$$

Typically, at least some of these are computed offline (e.g., using an optimization-based kinodynamic planner), while others can also be generated online. There is no requirement on the optimality of each motion primitive, although *optimal motion primitives*, where J is minimized for the respective start and goal states, are beneficial in our setting.

Our kinodynamic planning approach iteratively improves the solution. In every iteration, the following steps are executed: i) the set of used motion primitives grows and the bound that limits the maximum magnitude of discontinuous jumps that a solution may have is computed (Lines 4 to 5); ii) our discrete planner, db-A*, computes an initial solution that may include a bounded violation of some constraints

Algorithm 1: kMP-db-A*: Kinodynamic Motion Planning with db-A*

```

1  $\mathcal{M} \leftarrow \emptyset$  ▷ Set of motion primitives
2  $c_{\max} \leftarrow \infty$  ▷ Solution cost bound
3 for  $n = 1, 2, \dots$  do
4    $\mathcal{M} \leftarrow \mathcal{M} \cup \text{AddPrimitives}()$ 
5    $\delta \leftarrow \text{ComputeDelta}(\mathcal{M})$ 
6    $\mathbf{X}_d, \mathbf{U}_d, T_d \leftarrow \text{db-A}^*(\mathbf{x}_s, \mathbf{x}_f, \mathcal{X}_{\text{free}}, \mathcal{M}, \delta, c_{\max})$ 
7   if  $\mathbf{X}_d, \mathbf{U}_d$  successfully computed then
8      $\mathbf{X}, \mathbf{U}, T \leftarrow \text{Optimization}(\mathbf{X}_d, \mathbf{U}_d, T_d, c_{\max})$ 
9     if  $\mathbf{X}, \mathbf{U}$  successfully computed then
10      Report  $(\mathbf{X}, \mathbf{U}, T)$  ▷ New solution found
11       $c_{\max} \leftarrow \min(c_{\max}, J(\mathbf{X}, \mathbf{U}, T))$  ▷ cost bound
12       $\mathcal{M} \leftarrow \mathcal{M} \cup \text{ExtractPrimitives}(\mathbf{X}, \mathbf{U})$ 

```

(Line 6); iii) the result of db-A* warm-starts an optimization-based formulation (Line 8); and iv) additional motion primitives are extracted from the optimization (Line 12).

A. db-A*: Discontinuity-bounded A*

In the following, we rely on a user-specified *metric* $d : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, which measures the distance between two states. This is analogous to sampling-based planners, and we assume that (\mathcal{X}, d) is a metric space in order to use efficient nearest neighbor data structures, such as k-d trees.

Definition 2: Sequences $\mathbf{X} = \langle \mathbf{x}_0, \dots, \mathbf{x}_T \rangle$, $\mathbf{U} = \langle \mathbf{u}_0, \dots, \mathbf{u}_{T-1} \rangle$ are δ -*discontinuity-bounded* solutions to Eq. (3) iff the following conditions hold:

$$d(\mathbf{x}_{k+1}, \text{step}(\mathbf{x}_k, \mathbf{u}_k)) \leq \delta \quad k \in \{0, \dots, T-1\} \quad (5a)$$

$$\mathbf{u}_k \in \mathcal{U} \quad \forall k \in \{0, \dots, T-1\} \quad (5b)$$

$$\mathbf{x}_k \in \mathcal{X}_{\text{free}} \quad \forall k \in \{0, \dots, T\} \quad (5c)$$

$$d(\mathbf{x}_0, \mathbf{x}_s) \leq \delta \quad (5d)$$

$$d(\mathbf{x}_T, \mathbf{x}_f) \leq \delta. \quad (5e)$$

Intuitively, Definition 2 enforces that the sequences connect the start and goal states with a bounded error δ in the dynamics, which corresponds to “stitching” primitives together. By the definition of a metric space, \mathbf{X} and \mathbf{U} fulfill all constraints of Eq. (3) if $\delta = 0$.

Our approach to compute such sequences is *discontinuity-bounded A** (db-A*), see Algorithm 2. Db-A* is, like A*, an informed search that relies on a *heuristic* $h : \mathcal{X} \rightarrow \mathbb{R}$ to explore an implicitly defined directed graph efficiently. Nodes in the graph represent states and an edge between nodes indicates a δ -bounded motion that connects the states.

The algorithm keeps track of nodes to explore using a *priority queue*, which is sorted by the lowest $f(\mathbf{x}) = g(\mathbf{x}) + h(\mathbf{x})$ value, where $g(\mathbf{x})$ is the cost-to-come. The overall structure is the same as in A*: The OPEN priority queue \mathcal{O} is initialized with the start state (Line 3), the current node n is the removed first element of \mathcal{O} (Line 5), and that node is expanded in order to compute valid (collision-free) neighbors (Lines 10 to 12). Newly found nodes are added directly to \mathcal{O} (Line 17), while previously found nodes are updated if the solution cost is reduced (Lines 21 to 26).

Unlike A*, we consider two states to be identical for nonzero δ values. The major changes of db-A* from A*

are highlighted in Algorithm 2. We use the notation $\mathbf{x} \oplus m$ to indicate that a motion m is applied to state \mathbf{x} ; that is, we shift m by the translational part of \mathbf{x} . For efficient search, we adopt two k-d trees (rather than a hashmap in A*). Namely, we use \mathcal{T}_m (Line 1) to index the start states of all provided motion primitives, which can be done once at the beginning. In order to reuse the same distance metric d , we use the translation-invariance property and set \mathbf{x}^t of a given state \mathbf{x} to $\mathbf{0}$. This data structure allows us to efficiently find suitable motions extending from a given state (Line 9). The second k-d tree \mathcal{T}_n (Line 2) contains the states of all explored nodes and grows dynamically (Line 18). It is used to find nearby previously explored states (Line 15) in order to limit the graph size and enable rewiring. The discontinuity with a magnitude of up to δ may occur in two cases. First, when we select suitable motion primitives for expansion (Line 9) and second, when we prune a potential new node in favor of already existing states (Line 15). The tradeoff between the two can be selected by a user-specified parameter $\alpha \in (0, 1)$. For most search-based algorithms, collision checking is achieved using a binary occupancy grid, which makes the choice of the grid size a critical decision. Instead, we rely on broadphase collision checking. The required data structures can be efficiently precomputed for the environment and each motion primitive. For the collision check in Line 11, we only need to shift the data structure for the selected motion primitive, before executing the broadphase collision check.

We now discuss the theoretical properties of db-A*.

Theorem 1: Sequences \mathbf{X} and \mathbf{U} returned by db-A* (Algorithm 2) are a δ -discontinuity-bounded solution to the given motion planning problem.

Proof: Algorithm 2 only returns a sequence in Line 7. Due to the condition in Line 6, Eq. (5e) holds.

By Definition 1, we have $d(\mathbf{x}_{k+1}, \text{step}(\mathbf{x}_k, \mathbf{u}_k)) = 0 \leq \delta$, $\mathbf{u}_k \in \mathcal{U}$, and $\mathbf{x}_k \in \mathcal{X}$ for each motion primitive $m \in \mathcal{M}$. Thus, Eq. (5b) holds. During the search, we expand motions whose start states are at most $\alpha\delta$ away from the current state $n.\mathbf{x}$ (Line 9). There are two cases. First, the motion corresponds to an edge leaving from the current state $n.\mathbf{x}$ (Line 17), in which case we have $d(m^0, n.\mathbf{x}) \leq \alpha\delta$, where m^0 is the (translated) first state of motion m . Second, the motion becomes an edge leaving from some neighbor state $n'.\mathbf{x}$ that is at most $(1 - \alpha)\delta$ away from $n.\mathbf{x}$ (Line 15), in which case we have $d(m^0, n'.\mathbf{x}) \leq \alpha\delta + (1 - \alpha)\delta = \delta$, using the triangle inequality of our metric space. Thus, Eq. (5a) holds for all edges.

We already know that $\mathbf{x}_k \in \mathcal{X}$. Motions are only used as edges, if the entire motion is in $\mathcal{X}_{\text{free}}$ (Line 11), thus Eq. (5c) holds. Finally, Eq. (5d) holds because \mathcal{O} is initialized with \mathbf{x}_s (Line 3) and Eq. (5a) holds. ■

Remark 1: Db-A* is incomplete and suboptimal if $\delta > 0$.

Proof: Consider an example where a robot has to move through a narrow door to navigate to an adjacent room. Even if a δ -discontinuity-bounded solution for the problem exists, db-A* may not find it, because motions are added in a random order and only if no other node is within $(1 - \alpha)\delta$ (Line 15). Since db-A* is incomplete, it cannot guarantee

Algorithm 2: db-A*

Input: $\mathbf{x}_s, \mathbf{x}_f, \mathcal{X}_{\text{free}}, \mathcal{M}, \delta, c_{\text{max}}$
Result: $\mathbf{X}_d, \mathbf{U}_d$ or Infeasible

```

1  $\mathcal{T}_m \leftarrow \text{NearestNeighborInit}(\mathcal{M})$   $\triangleright$  Use start states of motions (excl. position)
2  $\mathcal{T}_n \leftarrow \text{NearestNeighborInit}(\{\mathbf{x}_s\})$   $\triangleright$  capture explored vertices (incl. position)
3  $\mathcal{O} \leftarrow \{\text{Node}(\mathbf{x} : \mathbf{x}_s, g : 0, h : h(\mathbf{x}_s), p : \text{None}, a : \text{None})\}$ 
 $\triangleright$  Initialize open priority queue
4 while  $|\mathcal{O}| > 0$  do
5    $n \leftarrow \text{PriorityQueuePop}(\mathcal{O})$   $\triangleright$  Lowest  $f$ -value
6   if  $d(n.\mathbf{x}, \mathbf{x}_f) \leq \delta$  then
7     return  $\mathbf{X}_d, \mathbf{U}_d, T_d$   $\triangleright$  Lowest  $f$ -value
8    $\triangleright$  Find applicable motion primitives with discontinuity up to  $\alpha\delta$ 
9    $\mathcal{M}' \leftarrow \text{NearestNeighborQuery}(\mathcal{T}_m, n.\mathbf{x}^t, \alpha\delta)$ 
10  foreach  $m \in \mathcal{M}'$  do
11    if  $n.\mathbf{x} \oplus m \notin \mathcal{X}_{\text{free}}$  then
12      continue  $\triangleright$  entire motion is not collision-free
13     $g_t \leftarrow n.g + \text{cost}(m)$   $\triangleright$  tentative  $g$  score for this action
14     $\triangleright$  find already explored nodes within  $(1 - \alpha)\delta$ 
15     $\mathcal{N}' \leftarrow \text{NearestNeighborQuery}(\mathcal{T}_n, n.\mathbf{x} \oplus m, (1 - \alpha)\delta)$ 
16    if  $\mathcal{N}' = \emptyset$  then
17       $\text{PriorityQueueInsert}(\mathcal{O}, \text{Node}(\mathbf{x} : n.\mathbf{x} \oplus m, g : g_t, h : h(n.\mathbf{x} \oplus m), p : n, a : m))$ 
18       $\text{NearestNeighborAdd}(\mathcal{T}_n, n.\mathbf{x} \oplus m)$ 
19    else
20      foreach  $n' \in \mathcal{N}'$  do
21        if  $g_t < n'.g$   $\triangleright$  This motion is better than a known motion
22          then
23             $n'.g = g_t$   $\triangleright$  Update cost
24             $n'.p = n$   $\triangleright$  Update parent
25             $n'.a = m$   $\triangleright$  Update action
26             $\text{PriorityQueueUpdate}(\mathcal{O}, n')$ 
27 return Infeasible

```

that no better δ -discontinuity-bounded solution exists, once it finds one. ■

We note that Remark 1 uses a very strong definition of completeness in continuous state space. Other possible definitions include δ -robust completeness [9]; we leave the analysis regarding that property to future work. For the purpose of this paper, it is important to recognize that the strong properties of A* hold in the limit, i.e., as $\delta \rightarrow 0$.

B. Kinodynamic Optimization

For the Optimization subroutine, we rely on *k-Order Motion Optimization* (KOMO) [17], which solves the following optimization problem:

$$\begin{aligned}
& \min_{\mathbf{X}} \sum_{l=1}^T \hat{J}(\mathbf{x}_{l-k:l}) & (6) \\
& \text{s.t.} \begin{cases} \mathbf{g}_l(\mathbf{x}_{l-k:l}) \leq \mathbf{0} & \forall l \in \{1, \dots, T\} \\ \mathbf{h}_l(\mathbf{x}_{l-k:l}) = \mathbf{0} & \forall l \in \{1, \dots, T\} \end{cases}
\end{aligned}$$

Here, $\mathbf{x}_{l-k:l}$ denotes the sequence $\mathbf{x}_{l-k}, \mathbf{x}_{l-k+1}, \dots, \mathbf{x}_l$ and the inequality constraints \mathbf{g}_l and equality constraints \mathbf{h}_l only depend on the current and up to k prior states. This k -order Markov assumption allows us to solve the nonlinear optimization problem efficiently e.g., using the augmented Lagrangian method, because k is typically small (1 to 3).

When using the Euler approximation in Eq. (2), we can transform Eq. (3) for a given T into Eq. (6) by encoding the dynamics, start, and goal constraints using \mathbf{h}_l and the action and state constraints into \mathbf{g}_l . Since \mathbf{U} is not a decision variable in this formulation, the dynamics constraint has to be encoded by using state constraints or by augmenting the state space. We note that if T and Δt are fixed and $J = T\Delta t$, we can use any \hat{J} to optimize in the nullspace of J . This allows us to include arbitrary regularization terms (in our case, smoothness) to guide the optimization and improve the convergence and success rate of the optimizer.

Some optimization methods may refine the given T_d in Line 8 of Algorithm 1 either by adding Δt as an optimization variable [13], [29] (which introduces additional nonlinearities), or by applying a linear search over multiple potential values of T that are around T_d , e.g., $T \in \langle 0.8T_d, T_d, 1.2T_d \rangle$. We use the latter approach for Algorithm 1.

When no estimate of T is available, we can use a linear search over T . For some dynamics, e.g., differentially-flat systems, it is also possible to use a modified binary search, where the first exponential search identifies an upper bound and the following binary search finds the optimal T . We use the latter approach for our baseline.

C. Motion Primitive Generation

Instead of sampling control sequences at random, we solve two-point boundary value problems with random start and goal configurations in free space with nonlinear optimization, which results in a superior primitive distribution. Specifically, we generate motion primitives offline using the following steps. First, random sampling of a start and goal configuration in free space; second, solving Eq. (6) using linear search over T ; and third, splitting the resulting motion into multiple pieces of a desired length. We sort the primitives using an iterative greedy method that approximately minimizes the dispersion. Let \mathcal{M} be the set of all motions, \mathcal{M}_s be the set of sorted motions, and $\mathcal{M}_r = \mathcal{M} \setminus \mathcal{M}_s$ be the set of remaining motions. We initialize $\mathcal{M}_s = \{\operatorname{argmax}_{m \in \mathcal{M}} d(m^0, m^f)\}$, where m^0 refers to the initial state of the motion and m^f to the final state of the motion. Then, we add an element to \mathcal{M}_s in each iteration selected by

$$\operatorname{argmax}_{m_r \in \mathcal{M}_r} \left(\min_{m_s \in \mathcal{M}_s} d(m_r^0, m_s^0) + \min_{m_s \in \mathcal{M}_s} d(m_r^f, m_s^f) \right). \quad (7)$$

Thus, we pick the motion in each iteration that maximizes the minimum distance to other, already picked motions.

For `AddPrimitives` we add motions from the precomputed sequence \mathcal{M}_s . Additional motions can be generated online using the same procedure.

Instead of letting users manually specify δ , we use the automatic procedure `ComputeDelta`, which estimates δ given a desired branching factor b_d . First, we initialize a k-d tree \mathcal{T}_m of all motions, as in Line 1 of Algorithm 2. Second, we randomly sample a state \mathbf{x}_{rand} . Third, we use \mathcal{T}_m to find the b_d -closest motions that could be applied from \mathbf{x}_{rand} . Fourth, we record the distance $\delta_r = \max d(m, \mathbf{x}_{\text{rand}})$, where m is one of the b_d -closest motions. The estimated value

of δ is the average over multiple δ_r values. This procedure reduces δ as the number of motion primitives increases in expectation and is easy to tune at the same time.

Motion primitives can also be extracted online in Algorithm 1. The `ExtractPrimitives` procedure uses the output of the optimization regardless of the constraint satisfaction and works as follows. First, intervals of valid sub-trajectories are computed by checking if all the constraints are fulfilled. Longer intervals can be split up as in the offline computation. The resulting primitives can be particularly useful for the planning problem at hand, because they are computed using the full knowledge of the environment.

D. Properties

We conjecture that the approach in Algorithm 1 will eventually compute the optimal solution, because as the number of iterations n increases, we add more primitives \mathcal{M} , which, by definition of `ComputeDelta`, reduces δ . Thus, as $n \rightarrow \infty$, we have $\delta \rightarrow 0$. For $\delta = 0$, db-A* as described in Algorithm 2 becomes regular A*, which is known to be complete and optimal. The major flaw of this argument is that, in the limit, we also have an infinite number of motion primitives and thus an infinite branching factor.

Formally, we can follow [6, Th. 3] to establish almost surely asymptotic optimality (which also implies probabilistic completeness) under the assumption that we have a non-zero probability of our `Optimization` method to find a solution if one exists. This assumption is justified by the fact that the nonlinear trajectory optimization has a region of attraction $\Delta > 0$ and for small $\delta > 0$ our initial guess will fall in this region of attraction, allowing the optimization method to eventually compute a solution if one exists.

Theorem 2: The kMP-db-A* motion planner in Algorithm 1 is asymptotically optimal, i.e.

$$\lim_{n \rightarrow \infty} P(\{c_n - c^* > \epsilon\}) = 0, \quad \forall \epsilon > 0, \quad (8)$$

where c_n is the best cost in iteration n and c^* is the optimal cost.

Proof: We closely follow [6, Th. 3]. Let S_1, \dots, S_n be random variables denoting the suboptimality $c_n - c^*$. In every iteration of Algorithm 1 we either reduce the cost if we find a new solution, or we remain with the same cost, i.e., $c_{n+1} \leq c_n$. Each iteration, we add more motion primitives and thus reduce δ . Using the assumption of the non-zero probability for our `Optimization` method to find a solution, we have $E[S_n | s_{n-1}] \leq (1 - \omega)s_{n-1}$, i.e., in expectation the solution improves by a constant amount $\omega > 0$ every iteration. Then, we have

$$\begin{aligned} E[S_n] &= \int E[S_n | s_{n-1}] P(s_{n-1}) ds_{n-1} \\ &\leq (1 - \omega) \int s_{n-1} P(s_{n-1}) ds_{n-1} \\ &= (1 - \omega) E[S_{n-1}] = (1 - \omega)^{n-1} E[S_1]. \end{aligned} \quad (9)$$

With the Markov inequality we have $P(S_n > \epsilon) \leq E[S_n] / \epsilon = (1 - \omega)^{n-1} E[S_1] / \epsilon$, which approaches 0 as $n \rightarrow \infty$. ■

V. EXPERIMENTAL RESULTS

We compare different motion planners, including ours, on the same problem scenarios. For fair comparison, we share code and data structures as much as possible, use the respective state-of-the-art open-source implementations, and focus on settings where the dynamics and not the collision-checking create challenges.

A. Dynamical Systems

Unicycle (1st order) has a 3-dimensional state space $[x, y, \theta] \in SE(2)$ and a 2-dimensional $[v, \omega] \in \mathcal{U} \subset \mathbb{R}^2$ control space with dynamics defined in [30, Eq. (13.18)]. The simplest version (v0) uses bounds $v \in [-0.5, 0.5]$ m/s and $\omega \in [-0.5, 0.5]$ rad/s. More interesting variants are a plane-like version (v1) using a positive minimum speed of 0.25 m/s, and a plane-like version with a rudder damage (v2) ($\omega \in [-0.25, 0.5]$ rad/s).

Unicycle (2nd order) has a 5-dimensional state space $[x, y, \theta, v, \omega] \in \mathcal{X} \subset \mathbb{R}^5$, a 2-dimensional $[\dot{v}, \dot{\omega}] \in \mathcal{U} \subset \mathbb{R}^2$ control space, and dynamics defined in [30, Eq. (13.46)]. Our version (v0) uses $v \in [-0.5, 0.5]$ m/s, $\omega \in [-0.5, 0.5]$ rad/s, $\dot{v} \in [-0.25, 0.25]$ m/s², and $\dot{\omega} \in [-0.25, 0.25]$ rad/s².

Car with trailer has a 4-dimensional state space $[x, y, \theta_0, \theta_1] \in \mathcal{X} \subset \mathbb{R}^4$, a 2-dimensional $[v, \phi] \in \mathcal{U} \subset \mathbb{R}^2$ control space, and dynamics and visualization given in [30, Eq. (13.19), Fig. 13.6]. We add an additional constraint $|\angle(\theta_0, \theta_1)| < \pi/4$ that avoids that the angle between the car and the trailer exceeds a threshold. Our version (v0) uses $v \in [-0.1, 0.5]$ m/s, $\phi \in [-\pi/3, \pi/3]$, $L = 0.25$ m, and $d_1 = 0.5$ m, where L and d_1 are defined in [30].

Quadrotor has a 13-dimensional state space (pose and first order derivatives using a Quaternion representation), a 4-dimensional control space (force for each of the four motors), and dynamics defined in [31, Eq. (1)]. We use the parameters of the Crazyflie quadrotor with limits on the motor forces, velocity, and angular velocity. Note that the low thrust-to-weight ratio of 1.4 is very challenging for kinodynamic motion planning and that problem settings with a harsh initial condition prevent the use of specialized methods [32], [33].

We use $\Delta t = 0.1$ s for all dynamical systems except the quadrotor, which uses $\Delta t = 0.01$ s due to the fast rotational dynamics.

B. Environments

For most of the dynamical systems, we consider three environments (see Fig. 2), which are inspired by the common use-cases in the related literature. For the v2 unicycle, we use the *wall* environment as shown in Fig. 1. For the quadrotor, we use an *empty* environment without obstacles. The scenario requires the quadrotor to recover from a harsh initial condition with an upside-down initial rotation and nonzero initial first derivatives. All environments only use simple geometric box shapes for efficient collision checking. The environments are bounded, where the bounds only limit the translational part of the state, i.e., parts of the robots are allowed to be outside. One such example is visible in the park solution of Fig. 2.




Fig. 2. Example environments park, kink, and bugtrap (left to right) for the car with trailer dynamics. The start state is indicated in filled blue, the goal with a blue outline, and a near-optimal solution of the (x, y) -part of the state as computed by kMP-db-A* as a time-colored line.

C. Algorithms

For a **search-based** approach, we rely on SBPL¹ (Search-based Planning Library), a commonly used C++ library with integration in the Robot Operating System (ROS). SBPL contains an example for unicycles, although the used dynamics do not match the ones from [30, eq. 13.18]. Thus, we generate our own primitives using the formulation in Section IV-B. Moreover, we make minor adjustments to the heuristic to enable time-optimal anytime planning using the provided implementation of ARA* in SBPL. Due to limits in SBPL², we limit our evaluation to the v0 first order unicycle.

For a **sampling-based** approach, we rely on OMPL [34] (Open Motion Planning Library), a widely used C++ library with integration in ROS through MoveIt. OMPL implements several kinodynamic planners, including SST* [9], which we use. As part of this work, we contribute minor changes to allow time as an optimization objective. Since sampling-based kinodynamic approaches cannot reach a goal state, we use a goal region instead that we verify to be small enough such that an optimizer can find an exact solution.

For an **optimization-based** approach, we rely on RAI³ (Robotic AI), a C++ library that implements KOMO and nonlinear optimization algorithms. For each of the dynamical systems, we implement the appropriate constraints and their derivative computation. In case of the trailer and the quadrotor, we add parts of the actions as decision variables (angle ϕ and motor forces, respectively); otherwise the decision variables are the state sequences only. As an initial guess, we use a geometric solution as found by RRT* of OMPL. We then use the modified binary search method as outlined in Section IV-B. This combination of *geometric RRT*+KOMO* is anytime like the other approaches we compare to.

For **db-A***, we implement Algorithm 2 in C++ using the data structures provided in OMPL to represent states and for nearest neighbor computation. Algorithm 1 is implemented in Python that executes C++ binaries for subroutines when necessary. As heuristic h , we use Euclidean distance divided by the upper bound of the speed. For the `AddPrimitives` function, we precompute 10 000 motion primitives for most

¹<https://github.com/sbpl/sbpl>

²The official documentation states: “[For custom scenarios], you will have to implement your own environment (a very involved topic that might be covered in the future).” <http://sbpl.net/node/47>

³<https://github.com/MarcToussaint/rai>

TABLE I

BENCHMARK RESULTS COMPARING SUCCESS RATE (p), TIME FOR THE FIRST FOUND SOLUTION (t^{st}), THE COST OF THE FIRST FOUND SOLUTION (J^{st}), AND THE COST OF THE SOLUTION AFTER 5 min (J^f). TIME AND COST ARE THE MEDIAN OVER 10 TRIALS. BEST RESULTS ARE **BOLD**.

#	System	Instance	SST*				SBPL				geom. RRT*+KOMO				kMP-db-A*			
			p	t^{st} [s]	J^{st} [s]	J^f [s]	p	t^{st} [s]	J^{st} [s]	J^f [s]	p	t^{st} [s]	J^{st} [s]	J^f [s]	p	t^{st} [s]	J^{st} [s]	J^f [s]
1	unicycle 1 st order, v0	park	1.0	1.2	5.9	3.5	1.0	0.0	6.2	6.2	1.0	3.1	5.3	3.2	1.0	0.9	3.2	3.1
2		kink	1.0	1.2	47.5	17.9	1.0	0.2	22.6	22.6	1.0	9.9	24.8	21.7	1.0	5.5	15.4	13.1
3		bugtrap	1.0	1.1	63.3	30.0	1.0	1.4	36.8	36.6	1.0	10.8	40.3	22.2	1.0	21.6	23.8	22.1
4	unicycle 1 st order, v1	kink	0.8	38.5	43.2	34.3					0.0	—	—	—	1.0	11.9	23.9	23.7
5	unicycle 1 st order, v2	wall	0.8	29.9	45.2	37.4					0.0	—	—	—	1.0	7.3	20.0	18.0
6	unicycle 2 nd order	park	1.0	4.7	14.4	7.5					1.0	5.5	10.6	6.1	1.0	7.9	6.8	6.1
7		kink	1.0	2.6	71.0	59.7					0.5	18.1	24.6	21.1	1.0	10.7	20.4	19.8
8		bugtrap	1.0	5.2	66.8	51.1					1.0	23.2	39.9	27.2	1.0	40.6	33.9	25.9
9	car with trailer	park	0.6	24.6	13.6	13.6					1.0	15.6	10.8	6.2	1.0	10.0	5.7	5.4
10		kink	1.0	9.1	70.8	66.5					0.1	217.5	174.4	130.8	1.0	94.8	34.1	24.2
11		bugtrap	1.0	0.7	47.5	43.8					0.0	—	—	—	1.0	8.3	21.5	19.4
12	quadrotor	empty	0.0	—	—	—					1.0	207.2	2.6	2.6	1.0	131.0	1.6	1.6

dynamical systems (30 000 for the quadrotor) and only add a subset per iteration. Generating the primitives took about 8 h per dynamical system utilizing all CPU cores.

The benchmark infrastructure is written in Python and all tuning parameters can be found in the open-source repository⁴. Collision checking is done using FCL (Flexible Collision Library) [35] in all cases. All approaches use the Euler integration Eq. (2), although KOMO uses an implicit formulation by design.

D. Benchmark

We execute our benchmark on a desktop computer with AMD Ryzen 9 3900X (3.8 GHz) and 32 GB RAM. Our results are summarized in Table I.

We summarize the main results as follows. **SBPL** can compute results very quickly and consistently. The initial solution quality is very high, but due to the limited number of primitives, the solution does not improve much over time (rows 1 – 3, Table I). The approach is not as general as the other ones, and we were unable to use it for all of our dynamical systems. **SST*** can find an initial solution very quickly; however the solution quality is initially poor, especially with higher-dimensional systems (rows 6–11). The convergence is slow – our 5 min timeout was not sufficient for **SST*** to fully converge in any of the cases. **Geometric RRT*+KOMO** can find near-optimal initial solutions, but does not work well in instances that require long trajectories and fails if the geometric initial guess is not close to a dynamically feasible motion. For example, finding an initial solution in the kink and bugtrap examples (rows 7, 8) took significantly longer than parallelpark. Another drawback is that this approach is incomplete, as visible for the v1 and v2 unicycle systems (row 4 and 5) and not globally optimal, e.g., row 10 and 11 show a very poor solution quality after 5 min. **kMP-db-A*** converged to the lowest-cost solution during the time limit in all cases. At the same time, it found the highest-quality first solution in all cases, although




Fig. 3. Success rate and solution cost over runtime (log-scale) for the unicycle (1st order) dynamical system in the kink environment (row 2 in Table I). The line is the median and the shaded region shows the first/third quartile over all trials that found a solution so far (up to 10 in total).

it often took more time to compute an initial solution than the other algorithms. We found that this is mostly caused by the challenging scenario of time-optimal planning: most motion primitives are time-optimal, i.e., result in bang-bang control. When allowing discontinuities, the estimated time horizon is often too short for the optimizer to find a solution, requiring multiple iterations in Algorithm 1 to report the first solution.

For brevity, Table I does not include any standard deviation. In general, we found that SBPL has almost no variance, SST* has a very high variance, and KOMO and kMP-db-A* are somewhere in between the two extremes. One example that includes the convergence behavior as well as the variance is shown in Fig. 3.

The runtime of the individual components of kMP-db-A* can vary widely, depending on δ . For example, in the bugtrap example for the trailer (row 11) it takes around 2 s for db-A* to find a solution with $\delta = 0.33$ and 14 s for the optimization, while during later iterations db-A* requires 46 s ($\delta = 0.12$) and the optimization only 6 s.

⁴<https://github.com/IMRCLab/kinodynamic-motion-planning-benchmark>

VI. CONCLUSION

We present a new kinodynamic motion planning technique, kMP-db-A*, that uses a novel graph-search method with trajectory optimization in an iterative fashion. For the graph search, we introduce db-A*, a generalization of A* that reuses motion primitives to compute trajectories with a bounded discontinuity. Then, we warm-start trajectory optimization using the output of the graph search and compute new motion primitives online. KMP-db-A* combines ideas and advantages of sampling-based, search-based, and optimization-based kinodynamic motion planners: it converges asymptotically to the optimal solution, directly solves for the time horizon, finds a near-optimal solution quickly, and does not require any additional post-processing.

The major limitation of kMP-db-A* is that it sometimes requires a long time to compute an initial solution. We believe that this is not a fundamental issue and that it can be improved using the following techniques in the future. First, we are interested in using stronger heuristics and bounded suboptimal and incremental graph search techniques to reuse information between iterations. Second, we plan to investigate the use of optimizers that do not operate over the full trajectory time horizon. Finally, we believe that our work also lays the foundation for novel kinodynamic multi-robot motion planners.

REFERENCES

- [1] M. Pivtoraiko and A. Kelly, "Kinodynamic motion planning with state lattice motion primitives," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2011, pp. 2172–2179.
- [2] M. Pivtoraiko, "Differentially constrained motion planning with state lattice motion primitives," Ph.D. dissertation, Carnegie Mellon University, 2012.
- [3] L. Jarin-Lipschitz, J. Paulos, R. Bjorkman, and V. Kumar, "Dispersion-minimizing motion primitives for search-based motion planning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2021, pp. 12 625–12 631.
- [4] S. M. LaValle and J. J. Kuffner, "Randomized kinodynamic planning," *I. J. Robotics Res.*, vol. 20, no. 5, pp. 378–400, 2001.
- [5] T. Kunz and M. Stilman, "Kinodynamic RRTs with fixed time step and best-input extension are not probabilistically complete," in *Int. Worksh. on Alg. Found.*, 2015, pp. 233–244.
- [6] K. Hauser and Y. Zhou, "Asymptotically optimal planning by feasible kinodynamic planning in a state-cost space," *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1431–1443, 2016.
- [7] M. Kleinbort, E. Granados, K. Solovey, R. Bonalli, K. E. Bekris, and D. Halperin, "Refined analysis of asymptotically-optimal kinodynamic planning in the state-cost space," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2020, pp. 6344–6350.
- [8] F. Grothe, V. N. Hartmann, A. Orthey, and M. Toussaint, "ST-RRT*: Asymptotically-Optimal Bidirectional Motion Planning through Space-Time," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2022.
- [9] Y. Li, Z. Littlefield, and K. E. Bekris, "Asymptotically optimal sampling-based kinodynamic planning," *I. J. Robotics Res.*, vol. 35, no. 5, pp. 528–564, 2016.
- [10] M. Zucker, N. D. Ratliff, A. D. Dragan, M. Pivtoraiko, M. Klingensmith, C. M. Dellin, J. A. Bagnell, and S. S. Srinivasa, "CHOMP: Covariant hamiltonian optimization for motion planning," *I. J. Robotics Res.*, vol. 32, no. 9–10, pp. 1164–1193, 2013.
- [11] J. Schulman, Y. Duan, J. Ho, A. X. Lee, I. Awwal, H. Bradlow, J. Pan, S. Patil, K. Goldberg, and P. Abbeel, "Motion planning with sequential convex optimization and convex collision checking," *I. J. Robotics Res.*, vol. 33, no. 9, pp. 1251–1270, 2014.
- [12] D. Malyuta, T. P. Reynolds, M. Szmuk, T. Lew, R. Bonalli, M. Pavone, and B. Acikmese, "Convex optimization for trajectory generation."
- [12] R. Bonalli, A. Cauligi, A. Bylard, and M. Pavone, "GuSTO: Guaranteed sequential trajectory optimization via sequential convex programming," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2019, pp. 6741–6747.
- [14] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocodyl: An efficient and versatile framework for multi-contact optimal control," in *Proc. IEEE Int. Conf. Robot. Autom.*, pp. 2536–2542.
- [15] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems," in *ICINCO*, 2004, pp. 222–229.
- [16] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, "STOMP: Stochastic trajectory optimization for motion planning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2011, pp. 4569–4574.
- [17] M. Toussaint, "A tutorial on newton methods for constrained trajectory optimization and relations to slam, gaussian process smoothing, optimal control, and probabilistic inference," in *Geometric and Numerical Foundations of Movements*, 2017, vol. 117, pp. 361–392.
- [18] R. Natarajan, H. Choset, and M. Likhachev, "Interleaving graph search and trajectory optimization for aggressive quadrotor flight," *IEEE Trans. Robot. Autom. Lett.*, vol. 6, no. 3, pp. 5357–5364, 2021.
- [19] B. Sakcak, L. Bascetta, G. Ferretti, and M. Prandini, "Sampling-based optimal kinodynamic planning with motion primitives," *Auton. Robots*, vol. 43, no. 7, pp. 1715–1732, 2019.
- [20] Z. Littlefield and K. E. Bekris, "Efficient and asymptotically optimal kinodynamic motion planning via dominance-informed regions," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2018, pp. 1–9.
- [21] R. Shome and L. E. Kavraki, "Asymptotically optimal kinodynamic planning using bundles of edges," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2021, pp. 9988–9994.
- [22] S. Choudhury, J. D. Gammell, T. D. Barfoot, S. S. Srinivasa, and S. Scherer, "Regionally accelerated batch informed trees (RABIT*): A framework to integrate local information into optimal path planning," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2016, pp. 4207–4214.
- [23] A. Perez, R. P. Jr, G. D. Konidaris, L. P. Kaelbling, and T. Lozano-Pérez, "LQR-RRT*: Optimal sampling-based motion planning with automatically derived extension heuristics," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2012, pp. 2537–2542.
- [24] D. Zheng and P. Tsotras, "Accelerating kinodynamic RRT* through dimensionality reduction," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2021, pp. 3674–3680.
- [25] A. Wu, S. Sadraddini, and R. Tedrake, "R3T: Rapidly-exploring random reachable set tree for optimal kinodynamic planning of nonlinear hybrid systems," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2020, pp. 4245–4251.
- [26] H.-T. L. Chiang, J. Hsu, M. Fiser, L. Tapia, and A. Faust, "RL-RRT: Kinodynamic motion planning via learning reachability estimators from rl policies," *IEEE Trans. Robot. Autom. Lett.*, vol. 4, no. 4, pp. 4298–4305, 2019.
- [27] B. Ichter and M. Pavone, "Robot motion planning in learned latent spaces," *IEEE Trans. Robot. Autom. Lett.*, vol. 4, no. 3, pp. 2407–2414, 2019.
- [28] W. Du, S.-K. Kim, O. Salzman, and M. Likhachev, "Escaping local minima in search-based planning using soft duplicate detection," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2019, pp. 2365–2371.
- [29] B. Ponton, M. Khadiv, A. Meduri, and L. Righetti, "Efficient multicontact pattern generation with sequential convex approximations of the centroidal dynamics," *IEEE Trans. Robot.*, vol. 37, no. 5, pp. 1661–1679, 2021.
- [30] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
- [31] G. Shi, W. Hönig, Y. Yue, and S.-J. Chung, "Neural-swarm: Decentralized close-proximity multirotor control using learned interactions," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2020, pp. 3241–3247.
- [32] S. Liu, K. Mohta, N. Atanasov, and V. Kumar, "Search-based motion planning for aggressive flight in se(3)," *IEEE Trans. Robot. Autom. Lett.*, vol. 3, no. 3, pp. 2439–2446, 2018.
- [33] B. Zhou, F. Gao, L. Wang, C. Liu, and S. Shen, "Robust and efficient quadrotor trajectory generation for fast autonomous flight," *IEEE Trans. Robot. Autom. Lett.*, vol. 4, no. 4, pp. 3529–3536, 2019.
- [34] I. A. Şucan, M. Moll, and L. E. Kavraki, "The open motion planning library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, 2012.
- [35] J. Pan, S. Chitta, and D. Manocha, "FCL: A general purpose library for collision and proximity queries," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2012, pp. 3859–3866.