

Single-query Path Planning Using Sample-efficient Probability Informed Trees

Daniel Rakita¹, Bilge Mutlu, Michael Gleicher

Abstract—In this work, we present a novel sampling-based path planning method, called *SPRINT*. The method finds solutions for high dimensional path planning problems quickly and robustly. Its efficiency comes from minimizing the number of collision check samples. This reduction in sampling relies on heuristics that predict the likelihood that samples will be useful in the search process. Specifically, heuristics (1) prioritize more promising search regions; (2) cull samples from local minima regions; and (3) steer the search away from previously observed collision states. Empirical evaluations show that our method finds shorter or comparable-length solution paths in significantly less time than commonly used methods. We demonstrate that these performance gains can be largely attributed to our approach to achieve sample efficiency.

I. INTRODUCTION

Solving high-dimensional path planning problems, such as for robot manipulator motion planning, remains a challenging and important problem. Many aspects of this problem are often addressed with *sampling-based* path planners, which sample a collision-check function in order to probe the configuration space and inform a search strategy. While, in theory, many sampling-based planners are guaranteed to eventually find a solution if one exists, in practice, even state-of-the-art approaches are unable to consistently solve challenging problems in reasonable amounts of time as they require many costly collision-check samples [1].

In this work, we present a novel sampling-based path planning method, called *SPRINT* (Sample-efficient **PR**obability **IN**formed **T**rees) that finds solutions for high dimensional planning problems significantly faster than many state-of-the-art approaches. Our method uses heuristics designed to minimize the total number of collision-check samples required to find solutions by modeling how likely regions of the search space will yield useful samples. These heuristics, as part of newly proposed global and local tree searches, are used to (1) prioritize more promising search regions to foster a greedy first global search strategy, while still exploring broadly in the limit; (2) cull samples from local minima regions in the local search to avoid wasteful, unfruitful samples; and (3) steer the local search away from previously observed collision states toward search regions predicted to contain more useful free-space samples. Our heuristics can be implemented

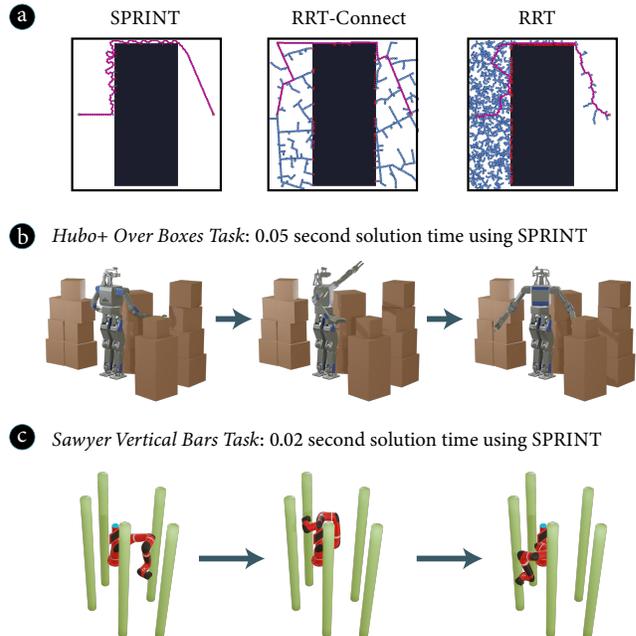


Fig. 1. (a) This 2D planning example shows the sample efficiency of the SPRINT planner, as many fewer collision-check samples were taken (left) compared to RRT-Connect (middle) and RRT (right). (b–c) This sample efficiency also makes path planning more efficient in high dimensions, such as for the illustrated robot manipulator motion planning problems.

with efficient information storage and retrieval methods that further accelerate the search. Specifically, our method does not perform expensive nearest neighbor checking each time a node is added to the search graph structure.

We assessed the efficacy of our method by running a large testbed of simulation experiments and compared our method to numerous path planners on single-arm and bimanual robot manipulation problems (§VII). Our method found shorter or comparable-length solution paths in significantly less time than the alternatives. We demonstrate that these performance advantages were, in large part, due to an increase in sample efficiency from our heuristics. We conclude our work by discussing the implications of our results in robotics applications and beyond. We provide open-source code for an implementation of our method.¹

II. RELATED WORKS

In general, there are three main paradigms for solving path planning problems: *sampling-based*, *local optimization*, and *search-based*. Sampling-based planners, many of which

¹Authors are with the Department of Computer Sciences, University of Wisconsin–Madison, Madison 53706, USA

[rakita|bilge|gleicher]@cs.wisc.edu

This work was supported by a Microsoft Research PhD Fellowship, National Science Foundation award 1830242, and a NASA University Leadership Initiative (ULI) grant awarded to the UW-Madison and The Boeing Company (Cooperative Agreement # 80NSSC19M0124).

¹<https://github.com/uwgraphics/lynx>

are extensions of Rapidly-Exploring Random Tree (RRT) [2], Probabilistic Roadmap (PRM) [3], or Expansive Space Tree (EST) [4] methods, often use random configuration-space samples to bootstrap a broad search strategy. These algorithms often have guarantees, such as probabilistic completeness or asymptotic optimality [5], [6], though they often *over-explore* the space with many unnecessary samples in order to achieve such guarantees.

Our work shares particular similarities with *informed-tree* sampling-based approaches that iteratively build search strategies that trade-off between exploration and exploitation based on information gathered at run-time [7]–[11]. Our work draws on these approaches, but differs in two main ways: (1) these approaches generally use linear local search segments between samples, whereas our work proposes a new local search designed to maneuver through narrow passages and avoid over-sampling; and (2) the search strategies for these approaches are typically more focused on achieving asymptotic optimality, whereas the search strategies for our method are more singularly focused on reaching a first feasible solution as quickly as possible.

Local optimization-based planners, such as CHOMP [12] or STOMP [13], use non-linear optimization to iteratively transform a trajectory into a higher quality path. These methods are designed to be greedier than sampling-based methods, though they often have fewer guarantees, strongly depend on the quality of the initial condition, and commonly do not converge on a feasible solution due to local minima.

Search-based planners prioritize exploration based on a set *heuristic*, such as best-first search, depth-first search, breadth-first search, or the A* heuristic [14]. However, because these searches are often structured in a discrete, grid-like fashion, higher dimensional planning often scales intractably in terms of memory and run-time complexity.

Our method draws on all of the above paradigms, and attempts to unite many of the discussed themes. To illustrate, SPRINT also uses random samples to bootstrap its search and explore broadly, but local optimization and search-based concepts are used to optimize and prioritize certain parts of the search to increase the probability of sample usefulness. We compare our method to other common approaches in our evaluation (§VII) to assess whether this bridging of concepts is effective for path planning.

III. PRELIMINARIES

A. Problem Statement

Consider χ as a d-dimensional configuration space with states within χ denoted as \mathbf{q} . Suppose χ_{obs} is a subset of χ considered to be infeasible space (i.e., obstacles). Feasible space can then be defined as the subset of χ *not* in an obstacle region: $\chi_{\text{free}} \equiv \chi \setminus \chi_{\text{obs}}$. A *path* in χ is a continuous function $\Gamma[\mathbf{q}_a, \mathbf{q}_b] : [0, 1] \rightarrow \chi$, where $\Gamma[\mathbf{q}_a, \mathbf{q}_b](0) = \mathbf{q}_a$, $\Gamma[\mathbf{q}_a, \mathbf{q}_b](1) = \mathbf{q}_b$, and all points along the path are in χ . The goal in path planning is to find a feasible, C^0 continuous path $\Gamma[\mathbf{q}_{\text{init}}, \mathbf{q}_{\text{goal}}]$ such that $\Gamma[\mathbf{q}_{\text{init}}, \mathbf{q}_{\text{goal}}](0)$ is a given start state, \mathbf{q}_{init} ; $\Gamma[\mathbf{q}_{\text{init}}, \mathbf{q}_{\text{goal}}](1)$ is a given goal state, \mathbf{q}_{goal} ; and $\Gamma[\mathbf{q}_{\text{init}}, \mathbf{q}_{\text{goal}}](u) \in \chi_{\text{free}}, \forall u \in [0, 1]$.

B. Graph-based Planning Structure

A common structure to solve the problem described above is a graph-based search [2]–[6], [10]. This framework organizes the search into a set of collision-free edges, E , between pairs of nodes in a set N to form a graph $G = (N, E)$. The goal in these searches is to construct G such that a feasible path $\Gamma[\mathbf{q}_{\text{init}}, \mathbf{q}_{\text{goal}}]$ is contained as some connected sequence of edge traversals between nodes within the graph. The feasibility of each edge is checked prior to its addition to the graph using a collision-check function. SPRINT interleaves global and local graph-based searches, overviewed in §IV-A.

C. Search-space Regions and Useful Samples

The overall goal of SPRINT is to minimize the number of collision-check samples by incorporating probability heuristics that model how likely regions of the search space will yield useful samples.

A *search-space region* is a segment of the configuration space χ loosely thought of as states in the vicinity of the line segment between two specified end-points. These regions are used to characterize and model where the planner has previously searched and where the planner could search in the future. Their loose definition is sufficient for these use cases. We denote a search-space region as $R[\mathbf{q}_a, \mathbf{q}_b]$, where \mathbf{q}_a and \mathbf{q}_b serve as the end-point markers of the region.

Intuitively, a *useful sample* is a sample that plays an integral role in constructing a final solution path. More formally, we consider a sample to be δ -useful if it achieves two criteria: (1) it is in free-space, χ_{free} ; and (2) it ultimately lies within a distance of δ of the final solution path, $\Gamma[\mathbf{q}_{\text{init}}, \mathbf{q}_{\text{goal}}]$. Note that this definition implies that it is unknowable if a sample is δ -useful until a solution path is found. However, our method uses heuristics to predict the probability that a particular search-space region will contain at least one δ -useful sample, which we will denote as $\Pr(U_\delta(R[\mathbf{q}_a, \mathbf{q}_b]))$. The following sections will detail the structure and functionalities of these probability heuristics.

IV. TECHNICAL OVERVIEW

This section outlines the search structure and overall strategy of the SPRINT method. Pseudocode of the SPRINT planning sub-processes can be found in Alg. 1–7.

A. SPRINT Global planning Level

The global planning level in SPRINT (Alg. 1) uses a tree-graph as a search structure, $T_g = (N_g, E_g)$, rooted at \mathbf{q}_{init} . A set of sampled collision-free *milestone points*, M , serve as intermediary goals for the global search to reach en route to \mathbf{q}_{goal} , reminiscent of FMT* [6]. On each global planning level loop, the planner selects a best search-space region for the next local search, $R[\mathbf{q}_n^*, \mathbf{q}_m^*]$ ($\mathbf{q}_n^* \in N_g$, $\mathbf{q}_m^* \in M$). Milestones that are reached via local searches to form global edges in E_g from already established nodes in N_g are added as global nodes in N_g . If the current set of milestone points does not foster a path to \mathbf{q}_{goal} , more milestone points are added to M , and the search

Algorithm 1: `sprint (qinit, qgoal)`

```

1 Ng ← {qinit}; Eg ← ∅; M ← {qgoal}; Rlocalmin ← ∅; k ← 50
2 Loop either indefinitely or until some run-time limit is met
3   while |M| < k do
4     M ← + randomly sampled state from χfree
5   Loop
6     R[qn*, qm*] ← probability_heuristic_1(Ng, M, Rlocalmin)
7     if R[qn*, qm*] ≡ ∅ then
8       k ← 1.5 * k
9       break // go to line 2 to add more milestones to M
10    (success, Tg) ← local_tree_search(qn*, qm*)
11    if success then
12      Ng ← + qm*; Eg ← + Γ[qn*, qm*]; M ← - qm*
13      if qm* ≡ qgoal then
14        return Γ[qinit, qgoal] // solution found
15    else
16      Rlocalmin ← + R[qn*, qm*] // remember that this search-space
        region appears to be a local minimum trap

```

Algorithm 2: `local_tree_search (qn*, qm*)`

```

1 Ng ← {qn*}; Eg ← ∅; Nstack ← ∅; Ncheckpoints ← {qn*}; qx ← qn*
2 Hexploit ← ∅ // hashmap with (q, (float, float)) key-value pairs
3 Hexplore ← ∅ // hashmap with (q, (float, float)) key-value pairs
4 Hobs ← ∅ // hashmap with (q, {q}) sets key-value pairs
5 Hnum ← ∅ // hashmap with (q, float) key-value pairs
6 add_checkpoint ← true // the root node is always labeled as a checkpoint
7 Loop
8   if probability_heuristic_2(qx, Tg, Ωg) then
9     if add_checkpoint then
10      Ncheckpoints ← + qx // label qx as checkpoint
11      exploit ← ||qx - qm*||; explore ← ||qn* - qx||
12      Hexploit ← + (qx, (0, exploit)); Hexplore ← + (qx, (0, explore))
13      Hobs ← + (qx, ∅); Hnum ← + (qx, 1)
14      R[qx, qc*] ← probability_heuristic_3(qx, Tg, Ωg)
15      if in_collision(qc*) then
16        for all q ∈ Ncheckpoints ∩ Γ[qn*, qx] // all checkpoints on the
        path from qx back to qn*
17        do
18          Hexploit[q][0] += 1; Hexplore[q][0] += 1; Hobs[q] ← + qc*
19          if Nstack.is_empty() then
20            return (false, Tg) // if the stack is empty, the local
            search has failed
21          qx ← Nstack.pop_last_added()
22          add_checkpoint ← true
23        else
24          // if qc* is not in collision
25          Γ[qx, qc*] ← qx + u * (qc* - qx), u ∈ [0, 1]
26          Ng ← + qc*; Eg ← + Γ[qx, qc*]; Nstack ← + qx; qx ← qc*
27          if qc* ≡ qm* then
28            return (true, Tg) // local search solution was found
29            exploit ← ||qc* - qm*||; explore ← ||qn* - qc*||
30            for all q ∈ Ncheckpoints ∩ Γ[qn*, qx] // all checkpoints on the
            path from qx back to qn*
31            do
32              // update checkpoints with information about qc*
33              Hnum[q] += 1
34              if exploit < Hexploit[q][1] then
35                Hexploit[q][0] ← 0; Hexploit[q][1] ← exploit
36              else
37                Hexploit[q][0] += 1
38                if explore > Hexplore[q][1] then
39                  Hexplore[q][0] ← 0; Hexploit[q][1] ← explore
40              else
41                Hexplore[q][0] += 1
42            add_checkpoint ← false
43          else
44            // if Probability Heuristic 2 considered the most recent qx as
            not worthwhile to extend, try the next node on the stack
45            if Nstack.is_empty() then
46              return (false, Tg) // if the stack is empty, the local
              search has failed
47            qx ← Nstack.pop_last_added()

```

proceeds. This process iterates until \mathbf{q}_{goal} is reached via T_g , provided a solution exists.

In general, the global planning level is intended to be a simple wrapper around the more sophisticated local search, outlined below. While the global planning level achieves probabilistic completeness in the limit (§VI), in practice, the

default starting set of 50 milestone points is sufficient to quickly solve all problems that we have ever tried.

B. SPRINT Local Planning Level

The local planning level in SPRINT (Alg. 2) also uses a tree-graph as a search structure, $T_\ell = (N_\ell, E_\ell)$. The root of the local tree is the first boundary point of the search-space region selected by the global planning level, \mathbf{q}_n^* , and the goal state that the local tree is trying to reach is the second boundary point of this region, \mathbf{q}_m^* .

The local planning level in SPRINT is a greedy, depth-first-search-like algorithm that uses heuristics to intelligently select branching directions, backtrack to fruitful parts of the search tree, and stop as early as possible when a solution is unlikely to be found. The local search was designed to be particularly adept at steering around approximately convex-shaped obstacle sections and navigating through narrow passages (as seen in Figure 1a). While the local search has no guarantees and needs the global search to route around local minima regions, in practice, it is effective and solves many problems on its own, even in high dimensions.

The local planning level progresses in three steps: (1) Select a node from the tree to extend from, $\mathbf{q}_x \in N_\ell$; (2) Decide if \mathbf{q}_x is worthwhile to extend (if not, return to step 1); and (3) If \mathbf{q}_x is worthwhile to extend, calculate a candidate node to extend toward, \mathbf{q}_c^* . After \mathbf{q}_c^* is calculated, it is checked for collision (Alg. 2, line 14). If \mathbf{q}_c^* is in free-space, it is added as a node to the tree, the algorithm assesses and stores information about how the search has improved or worsened given the addition of the new node, and it becomes the next \mathbf{q}_x to extend (Alg. 2, lines 23–39). If \mathbf{q}_c^* is in collision, the collision-point is stored by the algorithm to help prevent the tree from colliding with the same region again and a next \mathbf{q}_x is popped off of a stack of nodes to extend from, N_{stack} (Alg. 2, lines 15–21). These steps repeat until either the local goal is reached or N_{stack} becomes empty.

The heuristics used in the local search often rely on its ability to efficiently assess and characterize sub-trees within the tree. Therefore, the local search tree stores information allowing fast responses to queries such as “has the sub-tree rooted at node \mathbf{q} recently progressed toward the goal \mathbf{q}_m^* ?”; or “what previously observed collision points are in close proximity to the sub-tree rooted at node \mathbf{q} ?”. To achieve these quick assessments, our method labels certain nodes in the local search tree as data *checkpoints*, then stores and accesses information at these checkpoints throughout the search. It would be highly inefficient for all nodes to be labeled as checkpoints. Thus, only the root node, \mathbf{q}_n^* , and nodes that have more than one child node are labeled as checkpoints, which reduces the overall number of checkpoints and structurally places them as roots of salient sub-trees where distinct branching decisions were made.

When pertinent information is processed during the search, e.g., a node is added or a collision point is detected, this information is back-propagated through the tree and stored at all checkpoint nodes that reside on the path from the current extend node to the local tree root, \mathbf{q}_n^* . This process can be

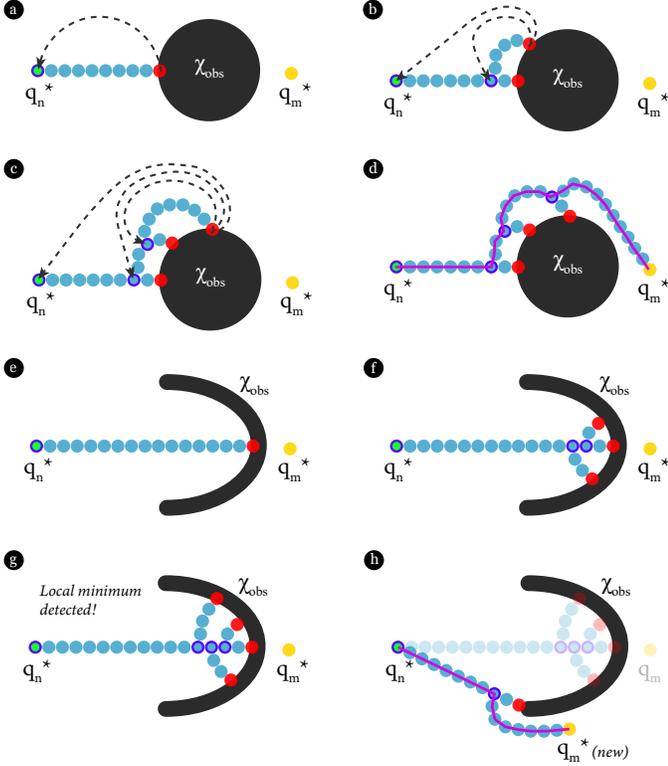


Fig. 2. (a) - (d) The detected collision states (red dots) get stored (dotted lines) in all checkpoint nodes (dots with purple outlines) along the path from the current extend node, \mathbf{q}_x to the root, \mathbf{q}_n^* . (e) - (f) Each added node passes its exploitation and exploration progress information back to its predecessor checkpoints until (g) the search decides that it is stuck in a local minimum trap based on lack of progress. (h) A new search-space region is selected by the global search to route around the local minimum region.

seen in Alg. 2 lines 17 and 30–39 where information is updated in four hashmap data structures (H_{exploit} , H_{explore} , H_{obs} , H_{num} , explained more below) on checkpoint nodes $\mathbf{q} \in N_{\text{checkpoints}} \cap \Gamma[\mathbf{q}_n^*, \mathbf{q}_x]$. This procedure ensures that checkpoints have access to information that has been processed “downstream” at newer, more distal parts of its sub-tree (illustrated in Figure 2).

V. PROBABILITY HEURISTICS

In this work, our method uses three probability heuristics, one in the global planning level and two in the local planning level (highlighted in green in Alg. 1 and 2). Each heuristic tries to guide the search into search-space regions with a high probability of containing δ -useful samples. In this section, we describe each of these three heuristics.

A. Probability Heuristic 1

Probability Heuristic 1 is used by the global search to select which search-space region would be best to next attempt a local search. This assessment is based on two criteria: (1) A good search-space region should have an end-point $\mathbf{q}_m^* \in M$ that gets closer to \mathbf{q}_{goal} . Intuitively, global edges that make progress toward the goal are more likely to have useful samples; and (2) A good search-space region should be far from any regions that already fostered local searches that did not reach their respective goals. The goal of this second

Algorithm 3: probability_heuristic_1 ($N_g, M, R_{\text{localmin}}$)

```

1  $R[\mathbf{q}_n^*, \mathbf{q}_m^*] \leftarrow \emptyset$ ;  $best\_value \leftarrow -\infty$ 
2  $c_1 \leftarrow 1$ ;  $c_2 \leftarrow 0.25$ ;  $w_1 \leftarrow 1$ ;  $w_2 \leftarrow 1.5$ 
3 for all  $\mathbf{q}_n \in N_g$  do
4   for all  $\mathbf{q}_m \in M$  do
5     if  $R[\mathbf{q}_n, \mathbf{q}_m] \notin R_{\text{localmin}}$  then
6        $x_1 \leftarrow \frac{\|\mathbf{q}_m - \mathbf{q}_{\text{goal}}\|}{\|\mathbf{q}_n - \mathbf{q}_{\text{goal}}\|}$  // characterizes how much closer region
           would be getting to  $\mathbf{q}_{\text{goal}}$ 
7        $g_1 \leftarrow \exp(-x_1^2 / 2c_1^2)$  // first Gaussian term
8        $x_2 \leftarrow 1$  // initialization, iteratively constructed in for
           loop below
9       for all  $R[\mathbf{q}_a, \mathbf{q}_b] \in R_{\text{localmin}}$  do
10         $\mathbf{p}_1 \leftarrow \text{proj}_+(\mathbf{q}_n, R[\mathbf{q}_b, \mathbf{q}_a])$  // projection point 1
11         $\mathbf{p}_2 \leftarrow \text{proj}_+(\mathbf{q}_m, R[\mathbf{q}_a, \mathbf{q}_b])$  // projection point 2
12         $d_1 \leftarrow \|\mathbf{p}_1 - \mathbf{q}_n\|$ ;  $d_2 \leftarrow \|\mathbf{p}_2 - \mathbf{q}_m\|$  // distances to
           projection points
13         $x_2 \leftarrow x_2 * (\frac{d_1 + d_2}{0.5 * \|\mathbf{q}_{\text{init}} - \mathbf{q}_{\text{goal}}\|})$ 
14         $g_2 \leftarrow (-\exp(-x_2^2 / 2c_2^2) + 1)$  // second Gaussian term
15        if  $(w_1 * g_1) * (w_2 * g_2) > best\_value$  then
16           $best\_value = (w_1 * g_1) * (w_2 * g_2)$ 
17           $R[\mathbf{q}_n^*, \mathbf{q}_m^*] \leftarrow R[\mathbf{q}_n, \mathbf{q}_m]$  // update best search-space
           region found thus far
18 return  $R[\mathbf{q}_n^*, \mathbf{q}_m^*]$  // return best overall search-space region found

```

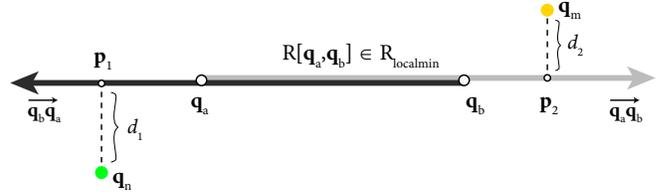


Fig. 3. Projection procedure of a search-space region $R[\mathbf{q}_n, \mathbf{q}_m]$ onto a region already deemed to be a local minimum trap, $R[\mathbf{q}_a, \mathbf{q}_b]$.

criterion is to de-prioritize the search from repeating a local search in a region already estimated to be a local minimum trap (stored in the R_{localmin} set in Alg. 1).

We model the two criteria outlined above as a mixture of two respective Gaussian functions:

$$\phi^{-1} * (U_{\delta}(R[\mathbf{q}_g, \mathbf{q}_m]) | T_g) = \phi^{-1} * [(w_1 * g_1(\mathbf{q}_n, \mathbf{q}_m, T_g)) * (w_2 * g_2(\mathbf{q}_n, \mathbf{q}_m, T_g))] \quad (1)$$

These Gaussian functions can be seen in Alg. 3, lines 7 and 8–14. The ϕ scalar is a normalization constant, and the w scalars signify weights that adjust the peaks of their Gaussian terms. The default values for w_1 and w_2 were hand-tuned, though our experiments below show that performance is not overly sensitive to these parameters. The c scalar is the standard deviation over these Gaussian functions, which we set as the Euclidean distance from start to goal.

The g_1 function uses Euclidean distances to achieve the first criterion discussed above. Specifically, x_1 , the input variable of the g_1 Gaussian function, is designed such that the closer a search-space region gets to \mathbf{q}_{goal} compared to where it started, the higher the probability value will be from g_1 . The g_2 function achieves the second criterion by projecting all candidate search-space regions onto rays cast from the end-points of prior unfruitful search-space regions (stored in the set R_{localmin}) and deprioritizing candidate regions with end-points close to any of these projections (Alg. 3, lines 9–14, seen illustrated in Figure 3).

B. Probability Heuristic 2

Probability Heuristic 2 is used by the local search to decide whether a particular node in the local tree search, \mathbf{q}_x , is

Algorithm 4: `probability_heuristic_2` ($\mathbf{q}_x, T_\ell, \Omega_\ell$)

```
1 if  $\mathbf{q}_x \in N_{\text{checkpoints}}$  then
2   | return false // prevents a node from being extended more than twice
3  $\kappa \leftarrow 0.3$  // lower bound for allowable probability
4 for all  $\mathbf{q} \in N_{\text{checkpoints}} \cap \Gamma[\mathbf{q}_n^*, \mathbf{q}_x]$  // all checkpoints on the path from  $\mathbf{q}_x$ 
   | back to  $\mathbf{q}_n^*$ 
5 do
6    $x_1 \leftarrow H_{\text{exploit}}[\mathbf{q}][0] / H_{\text{num}}[\mathbf{q}]; x_2 \leftarrow H_{\text{explore}}[\mathbf{q}][0] / H_{\text{num}}[\mathbf{q}]$ 
7    $x \leftarrow \min(x_1, x_2)$ 
8    $c \leftarrow \log_2(H_{\text{num}}[\mathbf{q}])^{-1}$ 
9    $g \leftarrow \exp(-x^2 / 2c^2)$  // Gaussian value
10  if  $g < \kappa$  then
11    | return false // if even one checkpoint is below the probability
      | threshold, the node is considered invalid to extend
12 return true
```

worthwhile to extend. This decision is made by checking information stored at all checkpoints along the path from \mathbf{q}_x back to \mathbf{q}_n^* . At a high level, \mathbf{q}_x is deemed worthwhile to extend if either (1) all of the sub-trees rooted at checkpoints along the path from \mathbf{q}_x back to \mathbf{q}_n^* have shown recent progress in terms of getting closer to the goal, \mathbf{q}_m^* . Throughout this work, we refer to this kind of progress as *exploitation* (stored and tracked in the H_{exploit} hashmap in Alg. 2); or (2) all of the sub-trees rooted at checkpoints along the path from \mathbf{q}_x back to \mathbf{q}_n^* have shown recent progress in terms of getting further from the tree root, \mathbf{q}_n^* . Throughout this work, we refer to this kind of progress as *exploration* (stored and tracked in the H_{explore} hashmap in Alg. 2). If the sub-trees that \mathbf{q}_x is a part of show no sign of recent exploitation or exploration progress, the heuristic considers this region of the local tree trapped in a local minimum, and \mathbf{q}_x is culled prior to extension.

We model Probability Heuristic 2 as the minimum of a set of probability functions:

$$\Pr_2(\mathcal{U}_\delta(\mathbf{R}[\mathbf{q}_x, \mathbf{q}_m^*]) \mid T_\ell) = \min_{\mathbf{q} \in \mathcal{B}} (g(\mathbf{q}_x, \mathbf{q}, T_\ell)), \mathcal{B} \equiv N_{\text{checkpoints}} \cap \Gamma[\mathbf{q}_n^*, \mathbf{q}_x] \quad (2)$$

The probability functions are modeled as Gaussian functions, constructed in Alg. 4, lines 4 – 9. Here, x_1 and x_2 are the number of collision-check samples since a node has made exploitation and exploration progress, respectively. These values are stored and updated in the H_{exploit} and H_{explore} hashmaps in Alg. 2, lines 17 and 31–38. We model this sub-tree assessment process by taking the minimum of x_1 and x_2 in Alg. 4, line 7. Taking the minimum ensures that if either exploitation or exploration progress are occurring in the given sub-tree (*i.e.*, either x_1 or x_2 are near 0), the sub-tree as a whole will still be considered promising. This value is inputted into the Gaussian function in line 9 to output the approximate probability value $g \in [0, 1]$. The standard-deviation c value is attenuated based on the number of nodes in a given sub-tree (stored and updated in the hashmap H_{num} in Alg. 2) such that new sub-trees with few nodes do not fail too abruptly even if they do not immediately show exploitation or exploration progress. In our prototype system, we use a probability cut-off of $\kappa = 0.3$, though our experiments below suggest that performance is not overly sensitive to this parameter selection.

C. Probability Heuristic 3

Probability Heuristic 3 is used by the local search to calculate which search-space region the local tree should extend

Algorithm 5: Common Functions

```
1 proj_scalar( $\mathbf{q}, \mathbf{R}[\mathbf{q}_a, \mathbf{q}_b]$ ) =  $\frac{(\mathbf{q}-\mathbf{q}_a) \cdot (\mathbf{q}_b-\mathbf{q}_a)}{(\mathbf{q}_b-\mathbf{q}_a) \cdot (\mathbf{q}_b-\mathbf{q}_a)}$ 
2 proj( $\mathbf{q}, \mathbf{R}[\mathbf{q}_a, \mathbf{q}_b]$ ) =  $\mathbf{q}_a + \text{proj\_scalar}(\mathbf{q}, \mathbf{R}[\mathbf{q}_a, \mathbf{q}_b]) * (\mathbf{q}_b - \mathbf{q}_a)$ 
3 proj_+( $\mathbf{q}, \mathbf{R}[\mathbf{q}_a, \mathbf{q}_b]$ ) =
    $\mathbf{q}_a + \max(\text{proj\_scalar}(\mathbf{q}, \mathbf{R}[\mathbf{q}_a, \mathbf{q}_b]), 0) * (\mathbf{q}_b - \mathbf{q}_a)$ 
4 hvs( $x$ ) =  $\begin{cases} 0 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0 \end{cases}$  // Heaviside step function
```

Algorithm 6: `probability_heuristic_3` ($\mathbf{q}_x, T_\ell, \Omega_\ell$)

```
1  $w_1 \leftarrow 0.5; w_2 \leftarrow 1; w_3 \leftarrow 1.2$ 
2  $\mathbf{q}_c \leftarrow \mathbf{q}_x + (\mathbf{q}_x - \mathbf{q}_p)$  // initial condition prior to gradient ascent
3  $N_{\text{obs}} \leftarrow \text{get\_nearby\_collision\_points}(\mathbf{q}_x, T_\ell, \Omega_\ell)$  if
   |  $|N_{\text{obs}}| > 0$  then
4   |  $\mathbf{q}_c += \xi$  // small amount of noise
5 for 0..num_iters do
6   |  $\mathbf{q}_c += (w_1 \frac{\partial g_1}{\partial \mathbf{q}_c}) + (w_2 \frac{\partial g_2}{\partial \mathbf{q}_c}) + (\text{hvs}(|N_{\text{obs}}|) * w_3 \frac{\partial g_3}{\partial \mathbf{q}_c})$  // Equations 4,
   | 5, and 6, respectively
7   |  $\mathbf{q}_c \leftarrow \mathbf{q}_x + \lambda * \frac{\mathbf{q}_c - \mathbf{q}_x}{\|\mathbf{q}_c - \mathbf{q}_x\|}$  // re-set to length  $\lambda$ 
8 return  $\mathbf{R}[\mathbf{q}_x, \mathbf{q}_c]$ 
```

into from an extend node \mathbf{q}_x . This region is computed based on three criteria: (1) A good search-space region should follow a straight line path with respect to the predecessor edge of \mathbf{q}_x . This criterion prevents the local search from making drastic turns and unnecessarily turning back on its own path; (2) A good search-space region should get closer to the local goal, \mathbf{q}_m^* . Intuitively, making progress toward the goal is estimated to be useful; and (3) A good search-space region should move away from previously observed collision points nearby (stored in the H_{obs} hashmap in Alg. 2). A search-space region that is pointing in the direction of a previously observed collision point is also likely to collide, and thus, is less likely to lead to a useful sample.

Instead of explicitly modeling a \Pr_3 model, we implicitly define the model through its gradient with respect to a search-space region end-point, \mathbf{q}_c : $\frac{\partial \Pr_3(\mathcal{U}_\delta(\mathbf{R}[\mathbf{q}_x, \mathbf{q}_c]) \mid T_\ell)}{\partial \mathbf{q}_c}$. Then, given this gradient, we perform gradient ascent to steer local edges in a way that would approximately optimize a \Pr_3 . We model this gradient as proportional to the sum of three weighted sub-term gradients, each trying to achieve one of the three criteria from Probability Heuristic 3 above, respectively:

$$\frac{\partial \Pr_3(\mathcal{U}_\delta(\mathbf{R}[\mathbf{q}_x, \mathbf{q}_c]) \mid T_\ell)}{\partial \mathbf{q}_c} \propto w_1 \frac{\partial g_1}{\partial \mathbf{q}_c} + w_2 \frac{\partial g_2}{\partial \mathbf{q}_c} + w_3 \frac{\partial g_3}{\partial \mathbf{q}_c} \quad (3)$$

The $\frac{\partial g_1}{\partial \mathbf{q}_c}$ sub-term gradient pulls the search-space region in a straight line with respect to its predecessor edge. Here, \mathbf{q}_p denotes the predecessor node of the extend node, \mathbf{q}_x :

$$\frac{\partial g_1}{\partial \mathbf{q}_c} = \frac{\mathbf{q}_x - \mathbf{q}_p}{\|\mathbf{q}_x - \mathbf{q}_p\|} \quad (4)$$

The $\frac{\partial g_2}{\partial \mathbf{q}_c}$ sub-term gradient pulls the search-space region toward the local goal, \mathbf{q}_m^* :

$$\frac{\partial g_2}{\partial \mathbf{q}_c} = \psi_2 * \frac{\mathbf{q}_m^* - \mathbf{q}_c}{\|\mathbf{q}_m^* - \mathbf{q}_c\|} \quad (5)$$
$$\psi_2 = \exp(-\|\mathbf{q}_m^* - \mathbf{q}_c\|^2 / 4\lambda^2) + 1$$

Here, ψ_2 increases the pull from $\frac{\partial g_2}{\partial \mathbf{q}_c}$ as the local tree approaches its goal which helps avoid overshoot when doing gradient ascent. The λ scalar is the fixed length of all local search edges.

Algorithm 7: `get_nearby_collision_points`($\mathbf{q}_x, T_\ell, \Omega_\ell$)

```

1  $N_{\text{obs}} \leftarrow \emptyset$ 
2 for  $\mathbf{q}_{\text{checkpoint}} \in N_{\text{checkpoints}} \cap \Gamma[\mathbf{q}_n^*, \mathbf{q}_x]$  // Ordered set of checkpoints from  $\mathbf{q}_x$ 
   back to  $\mathbf{q}_n^*$ 
3 do
4   for  $\mathbf{q}_{\text{obs}} \in H_{\text{obs}}[\mathbf{q}_{\text{checkpoint}}]$  do
5     if  $\mathbf{q}_{\text{obs}} \notin N_{\text{obs}}$  then
6        $N_{\text{obs}} \leftarrow \mathbf{q}_{\text{obs}}$  // add collision point to  $N_{\text{obs}}$ 
7       if  $|N_{\text{obs}}| > m$  then
8         //  $m$  is some maximum capacity of  $N_{\text{obs}}$ , e.g. 10
9         return  $N_{\text{obs}}$ 

```

The $\frac{\partial g_3}{\partial \mathbf{q}_c}$ sub-term gradient pushes the candidate edge away from a set of nearby, previously observed collision points $\mathbf{q}_{\text{obs}} \in N_{\text{obs}}$ (utility functions used in this sub-term gradient can be found in Alg. 5):

$$\frac{\partial g_3}{\partial \mathbf{q}_c} = \frac{1}{|N_{\text{obs}}|} \sum_{\mathbf{q}_{\text{obs}} \in N_{\text{obs}}} \psi_{3,1} * \psi_{3,2} * \frac{\text{proj}(\mathbf{q}_{\text{obs}}, R[\mathbf{q}_x, \mathbf{q}_c]) - \mathbf{q}_{\text{obs}}}{\|\text{proj}(\mathbf{q}_{\text{obs}}, R[\mathbf{q}_x, \mathbf{q}_c]) - \mathbf{q}_{\text{obs}}\|} \quad (6)$$

$$\psi_{3,1} = \text{hvs}(\text{proj_scalar}(\mathbf{q}_{\text{obs}}, R[\mathbf{q}_x, \mathbf{q}_c]))$$

$$\psi_{3,2} = 5 * \exp(-\|\text{proj}(\mathbf{q}_{\text{obs}}, R[\mathbf{q}_x, \mathbf{q}_c]) - \mathbf{q}_{\text{obs}}\|^2 / 4\lambda^2)$$

The push from a particular collision point, \mathbf{q}_{obs} , is given in the direction of \mathbf{q}_{obs} towards its projection onto the candidate search-space region, $R[\mathbf{q}_x, \mathbf{q}_c]$. If this projection point lies “behind” \mathbf{q}_x , *i.e.*, $\text{proj_scalar}(\mathbf{q}_{\text{obs}}, R[\mathbf{q}_x, \mathbf{q}_c]) < 0$, the collision point is considered already passed, and the $\psi_{3,1}$ term removes this collision point’s effect on the gradient. If \mathbf{q}_{obs} lies “ahead” of \mathbf{q}_x , $\psi_{3,2}$ attenuates the push strength away from \mathbf{q}_{obs} such that a smaller distance between \mathbf{q}_{obs} and its projection point corresponds to a stronger push, and vice versa. The scalar 5 coefficient on the $\psi_{3,2}$ term raises the peak of its corresponding Gaussian to further strengthen the push away from \mathbf{q}_{obs} if it is close to its projection onto $R[\mathbf{q}_x, \mathbf{q}_c]$. The set of nearby collision states, N_{obs} , is efficiently accessed in constant time, *i.e.*, without using a heavy data structure like a kd-tree, by using the checkpoint data back-propagation technique outlined in §IV-B (Alg. 7). Here, collision points are stored and accessed at checkpoint nodes in the H_{obs} hashmap throughout Alg. 2. Note that if $|N_{\text{obs}}| \equiv 0$, the effect of $\frac{\partial g_3}{\partial \mathbf{q}_c}$ is cancelled by the Heaviside step coefficient in Alg. 6, line 6, and thus, $\frac{\partial g_3}{\partial \mathbf{q}_c}$ does not need to be computed in this case.

The $\frac{\partial \text{Pr}_3(U_\delta(R[\mathbf{q}_x, \mathbf{q}_c]))}{\partial \mathbf{q}_c}$ gradient is used in gradient ascent iterations to steer the local tree, as seen in Alg. 6. Because this gradient ascent occurs at such a performance critical inner-loop process, we only do one or two iterations per local search loop in practice. By default, we use the point $\mathbf{q}_x + (\mathbf{q}_x - \mathbf{q}_p)$ as the initial condition for \mathbf{q}_c in the gradient ascent. Also, if $|N_{\text{obs}}| > 0$, we add a small amount of noise to this initial condition. In our prototype system, this noise vector is drawn from a d-dimensional uniform distribution $\xi \sim U^d(\frac{-\lambda}{100}, \frac{\lambda}{100})$. Without this small random push, the local search would just stay on a single plane embedded in the configuration space. We also re-set the edge length to λ after each gradient ascent iteration (Alg. 6, line 7).

VI. ANALYSIS

In this section, we show the probabilistic completeness property of SPRINT in \mathbb{R}^d .

Sketch Proof: Consider $\Gamma[\mathbf{q}_{\text{init}}, \mathbf{q}_{\text{goal}}]$ as any feasible solution path found by a probabilistically complete planner parameterized by a set of k nodes, $N \equiv \{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_k\}$, connected by linear edges, E . For example, solution paths found by RRT or PRM would fit this definition. Now, suppose d -dimensional open balls of radii $r_1 \dots r_k$, are centered around nodes $\mathbf{q}_1 \dots \mathbf{q}_k$, which we will denote as $\mathcal{B}_{r_1}(\mathbf{q}_1) \dots \mathcal{B}_{r_k}(\mathbf{q}_k)$. Each radius r_j will be selected such that all of $\mathcal{B}_{r_j}(\mathbf{q}_j)$ lies in free-space and all points in $\mathcal{B}_{r_j}(\mathbf{q}_j)$ can be connected by collision-free lines with all points in $\mathcal{B}_{r_{j+1}}(\mathbf{q}_{j+1}) \forall j \in [1, 2, \dots, k-1]$. Note that because $\Gamma[\mathbf{q}_{\text{init}}, \mathbf{q}_{\text{goal}}]$ is feasible, a radius of $r_j > 0$ must exist for all $\mathcal{B}_{r_j}(\mathbf{q}_j)$ [15].

Observation 1: Because milestone points are uniformly sampled from χ_{free} and each radius $r_j > 0$, each $\mathcal{B}_{r_j}(\mathbf{q}_j)$ will become dense with milestones in the limit [15]. **Observation 2:** SPRINT will process a local search between all current global node and milestone point pairs each loop through Algorithm 1 lines 2–16. **Observation 3:** If a collision-free, straight-line path exists between a global node \mathbf{q}_n and a milestone \mathbf{q}_m , the local search in SPRINT resolves to a straight-line. From these observations, it can be seen that the following recursive sequence will always happen in SPRINT: a milestone point will eventually be sampled in $\mathcal{B}_{r_j}(\mathbf{q}_j)$, which will be reached by a straight line local search from a global node in $\mathcal{B}_{r_{j-1}}(\mathbf{q}_{j-1}) \forall j \in [2, 3, \dots, k]$. Note that \mathbf{q}_{init} , *i.e.*, \mathbf{q}_1 , and \mathbf{q}_{goal} , *i.e.*, \mathbf{q}_k , start as a global node and milestone point, respectively, so these points bootstrap the recursive process and never have to be exactly sampled. Therefore, SPRINT must inherit the probabilistic completeness of any algorithm that would eventually find a feasible path $\Gamma[\mathbf{q}_{\text{init}}, \mathbf{q}_{\text{goal}}]$ as a homotopically equivalent path will always eventually be found by SPRINT. \square

VII. EVALUATION

We have assessed the efficacy of our approach in three empirical experiments. Although we demonstrate that SPRINT is able to solve low-dimensional testbed problems used to validate planning algorithms, as illustrated in Figure 1, we focus our assessment on higher dimensional problems with robot manipulators. In this section, we overview our experiments and share results.

A. Implementation Details

Our prototype SPRINT implementation is implemented in the Rust programming language. Robot manipulator self-collisions and environment collisions are detected using the `ncollide` Rust library. Configurations that exceed joint position limits are also considered in a collision state. All evaluations throughout this work were run on a Lenovo Legion laptop with an i7-9750H processor and 32GB RAM. All parameters used in SPRINT were held constant at their default values for all problems in our evaluation.

B. Evaluation Benchmark

We developed a set of seven benchmark tasks to compare our method against alternative path planners. The benchmark

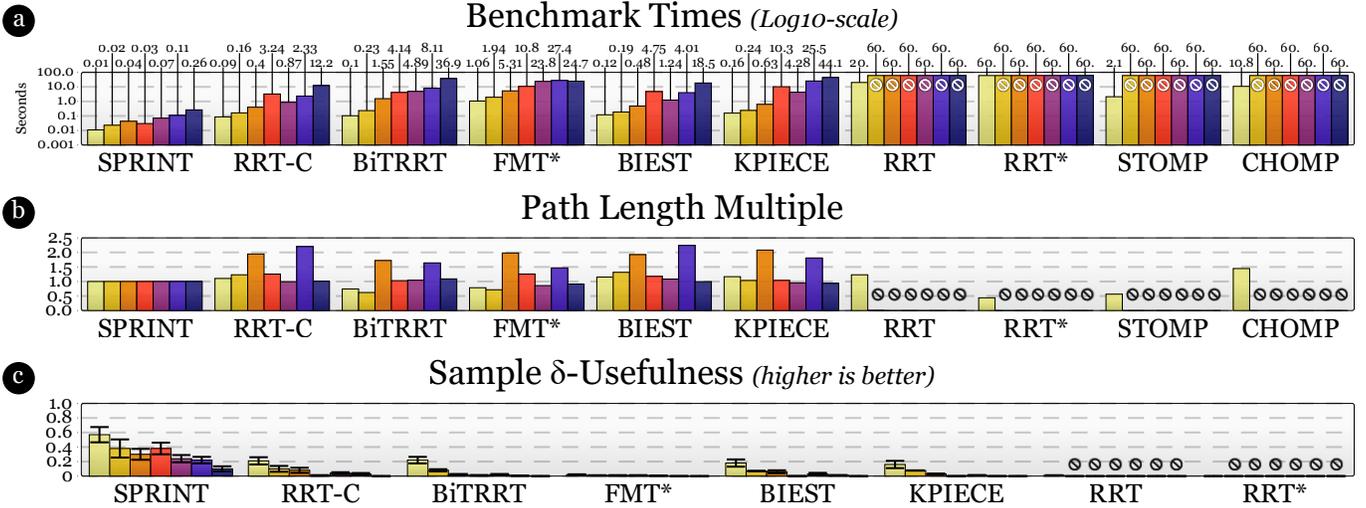


Fig. 4. Results from Experiments 1 and 2. Bar colors denote different tasks, with ordering from left to right being *Single Box* (UR5), *Table* (UR5), *Vertical Bars* (Sawyer), *Narrow Passage Box* (Sawyer), *Over Boxes* (Hubo+), *Bookshelf* (Hubo+), and *Arms around Table* (Hubo+). Error bars denote standard error. The values above the log10-scale bars in (a) denote the standard decimal values of the bars (in seconds). A circle with a line through it indicates that no solution was found on the given task in the allotted time.

was designed to test the planners on a wide variety of tasks that range in dimensionality and topological structure. Our benchmark consists of two tasks on a Universal Robots UR5 (6-DOF), two tasks on a Rethink Robotics Sawyer (7-DOF), and three bimanual tasks on the DRC-Hubo+ (15-DOF). Two of these tasks can be seen illustrated in Figure 1(b–c). The UR5 tasks included (1) *Single Box*, where the robot moves over a large box; and (2) *Table*, where the robot moves its arm up and around a small table. The Sawyer tasks included (3) *Vertical Bars* where the robot maneuvers around six vertical bars; and (4) *Narrow Passage Box* where the robot maneuvers out of a cramped box through a narrow passage. Finally, the Hubo tasks included (5) *Over Boxes* where the robot moves its arm up and over a set of boxes; (6) *Bookshelf* where the robot moves an item from the top shelf to the bottom shelf; and (7) *Arms around Table* where the robot maneuvers both of its arms up and around a table to reach objects on top. Our benchmark consisted of 100 trials through each task. All tasks had a maximum evaluation time of one minute.

C. Experiment 1: Comparisons with Alternative Planners

In our first experiment, we compared performance on our benchmark tasks to nine commonly used path planners: RRT [16], RRT-Connect [17], RRT* [5], FMT* [6], KPIECE [18], BiEST [4], BiTRRT [19], CHOMP [12], and STOMP [13]. All methods were tested using their MoveIt! implementations with default planner options. Also, we implemented RRT, RRT-Connect, FMT*, and EST within our framework to ensure that the planning primitives in SPRINT, *e.g.*, collision checking or graph operations, were not eliciting unintended performance gains. Our implementations used kd-trees for nearest neighbor checking. We used both our own implementations and MoveIt! in our evaluation and only report the better of the two results for fairness. Across all planners and tasks, the highly optimized implementations in MoveIt! were faster.

1) *Results—Computation Time*: Figure 4a provides an overview of the average times needed to complete our bench-

mark tasks. SPRINT computed solution paths for the benchmark problems often one to three orders of magnitude faster than alternative approaches. While many of the planners did well on the somewhat simpler, lower dimensional UR5 problems, SPRINT still out-performed the other methods on these tasks.

2) *Results—Path Length*: Figure 4b provides an overview of average path lengths per each planner and task. The paths computed by SPRINT are shorter or at most similar in length compared to other sampling-based planners. Thus, the performance advantages in computation time do not come at the expense of lower quality paths. While optimizing planners, such as RRT* and STOMP, did find higher quality paths on the UR5 *Single Box* task, these results came at a significantly higher computation cost. Further, these planners were not able to find any solutions in the allotted sixty seconds for any of the other benchmark tasks.

D. Experiment 2: Analysis of Performance Gains

In our second experiment, our goal was to investigate whether the performance gains exhibited by SPRINT in Experiment 1 could be attributed to enhanced sample efficiency. We assessed the ratio of points from the solution paths from Experiment 1 that were δ -Useful for each tested planner, which we defined in §III-C. We defined δ to be 2λ for this evaluation, where λ was the minimum edge length in the tree or graph structure. We only report this δ -Usefulness metric on the sampling-based planners as it is not applicable to the local-optimization-based planners (STOMP and CHOMP).

Figure 4c shows the δ -Useful metric results. We observe that SPRINT performs much better on this metric across all tasks, indicating that many more of its samples were useful. These results suggest that the performance gains from SPRINT were indeed due to our intended goal of enhancing sample efficiency.

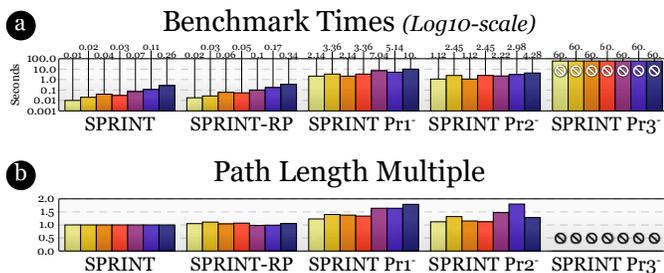


Fig. 5. Results from Experiment 3. Bar colors denote different tasks, with ordering from left to right indicating *Single Box* (UR5), *Table* (UR5), *Vertical Bars* (Sawyer), *Narrow Passage Box* (Sawyer), *Over Boxes* (Hubo+), *Bookshelf* (Hubo+), and *Arms around Table* (Hubo+). The values above the log10-scale bars in (a) denote the standard decimal values of the bars (in seconds). A circle with a line indicates that no solution was found.

E. Experiment 3: Analysis of Heuristics and Parameters

In our final experiment, our goal was to observe how the performance of SPRINT changes if we systematically perturb its parameters or remove its probability heuristics. We assessed four conditions: (1) SPRINT-RP, which randomly offsets all parameters such that each varies by up to 25% above or below its original hand selected value. These random offsets were drawn from a uniform distribution; SPRINT Pr₁, which replaced the Pr₁ heuristic with a random global edge selection; (2) SPRINT Pr₂, which replaced the Pr₂ heuristic with a random, 50% chance of extending a node; and (4) SPRINT Pr₃, which replaced the Pr₃ heuristic with a random extend direction process.

Figure 5 overviews the results from Experiment 3. First, we found that SPRINT-RP was comparable in terms of computation time and average path length, suggesting that SPRINT is not overly sensitive to parameter tuning. However, we see that performance significantly degrades when even one probability heuristic is removed, suggesting that all of our probability heuristics are integral and work in tandem in order to enhance sample-efficiency and, in turn, boost performance. Probability Heuristic 3 appears to be particularly important, as no solutions were found in the allotted sixty seconds without this heuristic model.

VIII. DISCUSSION

In this work, we presented a path planning approach that is able to quickly and reliably solve high-dimensional path planning problems. Through the notion of enhanced sample usefulness afforded by a set of probability heuristics, we show that our approach achieves significant performance gains over standard approaches.

Limitations—We note a number of limitations of our work that suggest future extensions. First, our probability heuristics were constructed by observation and intuition. While our three heuristics serve as proofs of concept for our overall premise, we believe that better models, either hand-engineered or data-driven, could improve results. Our approach also does not guarantee asymptotic optimality. Lastly, our work does not currently accommodate kinodynamic planning, though we believe that extensions to our work could explore kinodynamic steering functions compatible with Probability Heuristic 3.

ACKNOWLEDGMENTS

The authors thank Lydia Kavraki for valuable discussions regarding this work.

REFERENCES

- [1] K. Hauser, “Lazy collision checking in asymptotically-optimal motion planning,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 2951–2957.
- [2] S. M. LaValle, “Rapidly-exploring random trees: A new tool for path planning,” 1998.
- [3] L. E. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, “Probabilistic roadmaps for path planning in high-dimensional configuration spaces,” *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [4] D. Hsu, J.-C. Latombe, and R. Motwani, “Path planning in expansive configuration spaces,” in *Proceedings of International Conference on Robotics and Automation*, vol. 3. IEEE, 1997, pp. 2719–2726.
- [5] S. Karaman and E. Frazzoli, “Sampling-based algorithms for optimal motion planning,” *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [6] L. Janson, E. Schmerling, A. Clark, and M. Pavone, “Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions,” *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 883–921, 2015.
- [7] B. Burns and O. Brock, “Toward optimal configuration space sampling,” in *Robotics: Science and Systems*. Citeseer, 2005, pp. 105–112.
- [8] R. Alterovitz, S. Patil, and A. Derbakova, “Rapidly-exploring roadmaps: Weighing exploration vs. refinement in optimal motion planning,” in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 3706–3712.
- [9] B. Akgun and M. Stilman, “Sampling heuristics for optimal motion planning in high dimensions,” in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2011, pp. 2640–2645.
- [10] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, “Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic,” in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 2997–3004.
- [11] —, “Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3067–3074.
- [12] N. Ratliff, M. Zucker, J. A. Bagnell, and S. Srinivasa, “CHOMP: Gradient optimization techniques for efficient motion planning,” in *2009 IEEE International Conference on Robotics and Automation*. IEEE, 2009, pp. 489–494.
- [13] M. Kalakrishnan, S. Chitta, E. Theodorou, P. Pastor, and S. Schaal, “STOMP: Stochastic trajectory optimization for motion planning,” in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 4569–4574.
- [14] S. Aine, S. Swaminathan, V. Narayanan, V. Hwang, and M. Likhachev, “Multi-heuristic A*,” *The International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 224–243, 2016.
- [15] H. M. Choset, S. Hutchinson, K. M. Lynch, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun, *Principles of robot motion: theory, algorithms, and implementation*. MIT press, 2005.
- [16] S. M. LaValle and J. J. Kuffner Jr, “Randomized kinodynamic planning,” *The International Journal of Robotics Research*, vol. 20, no. 5, pp. 378–400, 2001.
- [17] J. J. Kuffner Jr and S. M. LaValle, “RRT-connect: An efficient approach to single-query path planning,” in *ICRA*, vol. 2, 2000.
- [18] I. A. Şucan and L. E. Kavraki, “Kinodynamic motion planning by interior-exterior cell exploration,” in *Algorithmic Foundation of Robotics VIII*. Springer, 2009, pp. 449–464.
- [19] L. Jaillet, J. Cortés, and T. Siméon, “Sampling-based path planning on configuration-space costmaps,” *IEEE Transactions on Robotics*, vol. 26, no. 4, pp. 635–646, 2010.