

Informed Sampling for Asymptotically Optimal Path Planning (Consolidated Version)

Jonathan D. Gammell, Member, IEEE,

Timothy D. Barfoot, Senior Member, IEEE, and Siddhartha S. Srinivasa, Senior Member, IEEE

Abstract—Anytime almost-surely asymptotically optimal planners, such as RRT*, incrementally find paths to every state in the search domain. This is inefficient once an initial solution is found as then only states that can provide a better solution need to be considered. Exact knowledge of these states requires solving the problem but can be approximated with heuristics.

This paper formally defines these sets of states and demonstrates how they can be used to analyze arbitrary planning problems. It uses the well-known L^2 norm (i.e., Euclidean distance) to analyze minimum-path-length problems and shows that existing approaches decrease in effectiveness factorially (i.e., faster than exponentially) with state dimension. It presents a method to address this curse of dimensionality by directly sampling the prolate hyperspheroids (i.e., symmetric n -dimensional ellipses) that define the L^2 informed set.

The importance of this direct informed sampling technique is demonstrated with Informed RRT*. This extension of RRT* has less theoretical dependence on state dimension and problem size than existing techniques and allows for linear convergence on some problems. It is shown experimentally to find better solutions faster than existing techniques on both abstract planning problems and HERB, a two-arm manipulation robot.

Index Terms—path planning, sampling-based planning, optimal path planning, informed sampling.

I. INTRODUCTION

THERE are many powerful path planning techniques in robotics. Popular approaches include graph-based searches, such as Dijkstra’s algorithm [1] and A* [2], and sampling-based methods, such as Probabilistic Roadmaps (PRM) [3], Expansive Space Trees (EST) [4], and Rapidly-exploring Random Trees (RRT) [5]. While sampling-based methods avoid the challenges of a priori discretizations, their stochastic nature limits their formal performance. They are said to be probabilistically complete if the probability of finding a solution, if one exists, approaches unity with an infinite number of samples. They are also said to be almost-surely asymptotically optimal if the probability of converging asymptotically to the optimum, if one exists, approaches unity with an infinite number of samples (e.g., RRT* [6]).

J. D. Gammell performed this work as a member of the Autonomous Space Robotics Lab at the University of Toronto Institute for Aerospace Studies. He is now with the Oxford Robotics Institute at the University of Oxford, Oxford, United Kingdom. Email: gammell@robotics.ox.ac.uk

T. D. Barfoot is with the Autonomous Space Robotics Lab at the University of Toronto Institute for Aerospace Studies, Toronto, Ontario, Canada. Email: timbarfoot@utoronto.ca

S. S. Srinivasa is with the School of Computer Science and Engineering, University of Washington, Seattle, Washington, USA. Email: siddh@cs.washington.edu

Manuscript submitted June 20, 2017.




Fig. 1. An illustration of how the set of states that can improve solution length shrinks as better solutions are found. Common estimates of this omniscient set are illustrated as informed sets. The L^2 informed set always contains the entire omniscient set (i.e., 100% recall) and shrinks along with it as a function of the current solution (i.e., high precision). It is exactly equal to the omniscient set in the absence of obstacles and constraints (i.e., 100% recall and precision). This paper shows that direct sampling this L^2 informed set is a necessary condition for almost-surely asymptotically optimal planners to scale effectively to high state dimensions. This technique is demonstrated with Informed RRT*.

RRT searches a planning problem by incrementally building a tree through free space. RRT* extends this procedure to incrementally rewire the tree during its construction. This rewiring locally optimizes every vertex in the tree and allows the algorithm to almost-surely converge asymptotically to the optimal path to every state in the problem domain. This is an inefficient way to find the optimal solution to a single planning query.

The only states that need to be considered in single-query scenarios are those that can provide a better solution [7]. While exact knowledge of these states requires solving the planning problem, they can often be approximated with heuristics (Fig. 1). These heuristics have previously been used to focus almost-surely asymptotically optimal search [8, 9] but can also provide insight into the optimal planning problem.

This paper uses the set of states that can provide a better solution to analyze incremental almost-surely asymptotically optimal planning. It formally defines this shrinking set as the omniscient set and shows that sampling it is a necessary condition for RRT*-style planners to improve a solution. It defines estimates of this set as informed sets and provides metrics to quantify them in terms of their compactness (i.e., precision) and completeness (i.e., recall). It uses these results to bound the probability of improving a solution to a holonomic planning problem by the probability of sampling an informed set with 100% recall.

The L^2 norm (i.e., Euclidean distance) is a well-known heuristic for problems seeking to minimize path length. It describes the omniscient set exactly in the absence of obstacles

and constraints (i.e., it is sharp¹) and always contains the omniscient set of a problem (i.e., it is universally admissible). This paper uses it to analyze the minimum-path-length problem and shows that existing focusing techniques (e.g., [8, 9]) are ineffective in high state dimensions. It is proven that these rejection-sampling approaches have a probability of improving a solution that goes to zero factorially (i.e., faster than exponentially) as state dimension increases.

This paper demonstrates how this minimum-path-length curse of dimensionality can be reduced by directly sampling the symmetric n -dimensional ellipse (i.e., prolate hyperspheroid), the L^2 informed set. The presented direct sampling approach always finds states that are believed to belong to a better solution regardless of the relative size of the L^2 informed set. It outperforms existing focusing techniques by orders of magnitude as state dimension increases.

The informed search approach is demonstrated with Informed RRT*. This extension of RRT* uses direct informed sampling and admissible graph pruning to focus the search for improvements. It is shown analytically to outperform existing techniques in terms of convergence rate, especially in high state dimensions, and to result in linear convergence on some problems. It is probabilistically complete and almost-surely asymptotically optimal. When the L^2 heuristic does not provide additional information (e.g., small planning problems and/or large informed sets) it is identical to RRT*. A version of Informed RRT* is publicly available in the Open Motion Planning Library (OMPL) [10].

Informed RRT* is evaluated experimentally on abstract problems and on the CMU Personal Robotic Lab's Home Exploring Robot Butler (HERB) [11], a 14-degree-of-freedom (DOF) mobile manipulation platform. These experiments show that it outperforms existing focusing techniques as state dimension increases, especially in problems with large planning domains.

This paper is organized as follows. Section II defines omniscient and informed sets and their associated precision and recall in preparation for the literature review presented in Section III. Section IV presents a direct informed sampling technique for problems seeking to minimize path length which is demonstrated with Informed RRT* in Section V. Section VI analyzes the expected convergence rate of RRT* algorithms and Section VII demonstrates the practical advantages of this improvement on abstract and simulated problems. Section VIII finally presents a closing discussion and thoughts on future work.

A. Statement of Contributions

This paper is a continuation of ideas first published in [12] and associated technical reports [13, 14] and makes the following specific contributions:

- Formally defines omniscient and informed sets (Definitions 3 and 7) and demonstrates how precision and recall can be used to quantify the performance of informed sampling (Definitions 8 and 9).

- Provides upper bounds on the probability that an incremental sampling-based planner improves a solution to a holonomic planning problem (Theorems 6 and 13).

- Bounds the expected next-iteration cost for RRT* algorithms on any minimum-path-length planning problem (Lemma 17) and shows that existing formulations of these algorithms for holonomic planning have a probability of improving a solution that decreases factorially with state dimension (Theorem 14).

- Develops a method to reduce this minimum-path-length curse of dimensionality by directly sampling the ellipsoidal L^2 informed set defined by a goal or countable set of goals and the current solution (Algs. 1–5).

- Proves that a planning algorithm using this approach, Informed RRT*, has better theoretical convergence (Theorems 18–20) and experimental performance than existing focused planning algorithms on holonomic problems.

II. OMNISCIENT AND INFORMED SETS

A formal discussion of the optimal planning problem is presented in support of the literature review. It includes definitions of the states that can provide a better solution, the omniscient set (Definition 3), and estimates of this set, informed sets, quantified by precision and recall (Definitions 7–10). These sets provide theoretical upper bounds on the probability of improving a solution to a holonomic problem that are used throughout the remainder of the paper (Theorems 6 and 13).

Finding the optimal path from a start to a goal is formally defined as the optimal planning problem (Definition 1). The given definition is similar to [6].

Definition 1 (Optimal planning). Let $X \subseteq \mathbb{R}^n$ be the state space of the planning problem, $X_{\text{obs}} \subseteq X$ be the states in collision with obstacles, and $X_{\text{free}} = \text{cl}(X \setminus X_{\text{obs}})$ be the resulting set of permissible states, where $\text{cl}(\cdot)$ represents the closure of a set. Let $x_{\text{start}} \in X_{\text{free}}$ be the initial state and $X_{\text{goal}} \subseteq X_{\text{free}}$ be the set of desired goal states. Let $\gamma : [0, 1] \rightarrow X_{\text{free}}$ be a sequence of states through collision-free space that can be executed by the robot (i.e., a collision-free feasible path) and \mathcal{P} be the set of all such nontrivial paths.

The optimal planning problem is then formally defined as the search for a path, $\gamma \in \mathcal{P}$, that minimizes a given cost function, $c : \mathcal{P} \rightarrow \mathbb{R}_{\geq 0}$, while connecting x_{start} to $x_{\text{goal}} \in X_{\text{goal}}$,

$$\gamma = \arg \min_{\gamma \in \mathcal{P}} c(\gamma) \quad \gamma(0) = x_{\text{start}}; \quad \gamma(1) \in X_{\text{goal}};$$

where $\mathbb{R}_{\geq 0}$ is the set of non-negative real numbers.

Many sampling-based planners, such as RRT*, probabilistically converge towards the optimum of these problems. Such planners are described as probabilistically complete and almost-surely asymptotically optimal (Definition 2).

Definition 2 (Almost-sure asymptotic optimality). A planner is said to be almost-surely asymptotically optimal if, with an infinite number of samples, the probability of converging asymptotically to the optimum (Definition 1), if one exists, is one,

$$\mathbb{P} \limsup_{q \rightarrow 1} c(\gamma_q) = c(\gamma^*) = 1;$$

¹A bound or estimate is said to be sharp if it is exactly equal to the true value (i.e., has perfect precision and recall) in at least one case.




Fig. 2. An example of how RRT* almost-surely converges asymptotically to the optimum by incrementally building and rewiring a tree through the entire problem domain. RRT* incrementally expands the tree into the problem domain and improve its connections. By continuing this process indefinitely, it almost-surely converges asymptotically to the optimal solution by asymptotically improving every path in the tree. This is inefficient in single-query planning scenarios.

where q is the number of samples, q is the path found by the planner from those samples, x^* is the optimal solution to the planning problem, and $c(\cdot)$ is the cost of a path.

Once any solution is found, the set of states that can provide a better solution can be defined as the omniscient set (Definition 3).

Definition 3 (Omniscient set). Let $g(x)$ be the cost of the optimal path from the start to a state, $x \in X_{\text{free}}$, the optimal cost-to-come,

$$g(x) := \min_{\gamma} f(c(\gamma)) \quad (0) = x_{\text{start}}; \quad (1) = xg;$$

and $h(x)$ be the cost of the optimal path from x to the goal region, the optimal cost-to-go,

$$h(x) := \min_{\gamma} f(c(\gamma)) \quad (0) = x; \quad (1) \in X_{\text{goal}}g;$$

The cost of the optimal path from x_{start} to x_{goal} constrained to pass through x is then given by $f(x) := g(x) + h(x)$. This defines the subset of states that can belong to a solution better than the current solution, c_i , as

$$X_f := \{x \in X_{\text{free}} \mid f(x) < c_i\} \quad (1)$$

Exact knowledge of X_f requires exact knowledge of the entire planning problem so we refer to it as the omniscient set.

RRT* builds a tree by incrementally adding states from the problem domain (Fig. 2). A necessary condition for it to improve a solution is that the newly added state belongs to the omniscient set (Lemma 4).

Lemma 4 (The necessity of adding states in the omniscient set). Adding a state from the omniscient set, $x_{\text{new}} \in X_f$, is a necessary condition for RRT* to improve the current solution, c_i ,

$$c_{i+1} < c_i \Rightarrow x_{\text{new}} \in X_f;$$

This condition is necessary but not sufficient to improve the solution as the ability of states in X_f to provide better solutions at any iteration depends on the structure of the tree (i.e., its optimality).

Proof. The proof of Lemma 4 from the supplementary online material appears in Appendix A-A. \square

The state added by RRT* at each iteration, x_{new} , is generated from a randomly sampled state, x_{rand} , and the nearest vertex in the existing tree,

$$v_{\text{nearest}} := \arg \min_{v \in V} f(kx_{\text{rand}} - vk_2g); \quad (2)$$

through expansion and differential constraints (i.e., the Steer function). Absent any constraints (i.e., in holonomic planning) this takes the form

$$x_{\text{new}} := \arg \min_{y \in X} f(kx_{\text{rand}} - yk_2 \mid ky - v_{\text{nearest}}k_2g); \quad (3)$$

where ℓ is a user-selected maximum edge length.

The number of tree vertices in the problem domain increases indefinitely with RRT* iterations. With an infinite number of iterations, eventually all reachable states will be no more than ℓ away from the nearest vertex in the tree. After these iterations, sampling the omniscient set is a necessary condition to add a state from the omniscient set and improve the solution (Lemma 5).

Lemma 5 (The necessity of sampling states in the omniscient set in holonomic planning). Sampling the omniscient set, $x_{\text{rand}} \in X_f$, is a necessary condition for RRT* to improve the current solution to a holonomic problem, c_i , after an initial iterations,

$$8i \quad c_{i+1} < c_i \Rightarrow x_{\text{rand}} \in X_f;$$

for any sample distribution that maintains a nonzero probability over the entire omniscient set.

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints with appropriate assumptions.

Proof. The proof of Lemma 5 from the supplementary online material appears in Appendix A-B. \square

This result provides an upper limit on the probability of RRT* improving a solution at any iteration (Theorem 6).

Theorem 6 (An upper bound on the probability of improving a solution to a holonomic planning problem given knowledge of the omniscient set). The probability that an iteration of RRT* improves the current solution to a holonomic problem, c_i , is bounded by the probability of sampling the omniscient set, X_f ,

$$8i \quad P(c_{i+1} < c_i) \leq P(x_{\text{rand}} \in X_f);$$

for any iteration, i , after a sufficient vertex density is achieved in the initial iterations.

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. Proof of Theorem 6 follows directly from Lemma 5. Sampling a state in X_f is a necessary but not sufficient condition to improve the solution after iterations; therefore, the probability of sampling such a state bounds the probability of improving the solution. \square

Knowledge of an omniscient set requires solving the planning problem; however, these results can be extended to estimates of the omniscient set defined by solution cost heuristics (Definition 7).

Definition 7 (Informed set). Let $\hat{h}(x)$ represent a heuristic estimate of the solution cost constrained to go through a state, $x \in X$. A heuristic estimate of the omniscient set can then be defined as

$$X_{\hat{h}} := \{x \in X \mid \hat{h}(x) < c_i\} :$$

Such a set will be referred to as an informed set.

There are an infinite number of potential informed sets for any planning problem and choosing the ‘best’ set requires methods to quantify their performance. In binary classification, estimates are evaluated in terms of their precision and recall (Fig. 3). Analogue terms can be defined in sampling-based planning to quantify the ability of informed sets to estimate the omniscient set (Definitions 8 and 9).

Definition 8 (Precision). The precision of an informed sampling technique is the probability that random samples drawn from the informed set could also be drawn from the omniscient set (e.g., the percentage of states drawn from the informed set, $X_{\hat{h}}$, that belong to the omniscient set, X_f). For uniform sampling of an informed set, this is a ratio of measures,

$$\text{Precision } X_{\hat{h}} := \frac{|X_{\hat{h}} \cap X_f|}{|X_{\hat{h}}|} :$$

Any informed set with nonzero sampling probability that is a subset of the omniscient set will have 100% precision.

Definition 9 (Recall). The recall of an informed sampling technique is the probability that random states drawn from the omniscient set could also be sampled from the informed set (e.g., the percentage of states that belong to the omniscient set, X_f , with a nonzero probability of being sampled from the informed set, $X_{\hat{h}}$). For uniform sampling of an informed set, this is a ratio of measures,

$$\text{Recall } X_{\hat{h}} := \frac{|X_{\hat{h}} \cap X_f|}{|X_f|} :$$

Any informed set with nonzero sampling probability that is a superset of the omniscient set will have 100% recall.

Informed sets with 100% recall (Definition 10) are important in almost-surely asymptotically optimal planning as less-than-perfect recall may exclude the optima to some problems.

Definition 10 (Admissible informed set). A heuristic is said to be admissible if it never overestimates the true value of the function,

$$\forall x \in X; \hat{h}(x) \leq f(x) :$$

Any informed set defined by such an admissible heuristic will contain all possibly better solutions and have 100% recall, i.e., $X_{\hat{h}} \supseteq X_f$:

This set will be referred to as an admissible estimate of the omniscient set, or an admissible informed set. If the heuristic




Fig. 3. An illustration of the precision and recall of estimating an oblong omniscient set, X_f , with a rectangular informed set, $X_{\hat{h}}$. The informed set is coloured to highlight where it is correct (light grey) incorrect (dark grey) or missing the omniscient set (white). Precision is the likelihood of correctly sampling the omniscient set by sampling the informed set. Recall is the coverage of the omniscient set by the informed set. For uniform distributions, both these terms are ratios of Lebesgue measures.

is an admissible estimate of the cost function for all possible problems then the set will be referred to as a universally admissible informed set.

These definitions allow the probability of improving a solution to a holonomic problem to be bounded by the probability of sampling any admissible informed set (Lemmas 11 and 12 and Theorem 13). The tightness of this bound will depend on the precision of the chosen estimate.

Lemma 11 (The necessity of adding states in an admissible informed set). Adding a state from an admissible informed set, $x_{\text{new}} \in X_{\hat{h}} \setminus X_f$, is a necessary condition for RRT* to improve the current solution, c_i ,

$$c_{i+1} < c_i \Rightarrow x_{\text{new}} \in X_{\hat{h}} \setminus X_f :$$

Proof. Lemma 11 follows directly from Lemma 4 given that $X_{\hat{h}} \supseteq X_f$. \square

Lemma 12 (The necessity of sampling states in an admissible informed set in holonomic planning). Sampling an admissible informed set, $x_{\text{rand}} \in X_{\hat{h}} \setminus X_f$, is a necessary condition for RRT* to improve the current solution to a holonomic problem, c_i , after an initial iterations,

$$8i \quad c_{i+1} < c_i \Rightarrow x_{\text{rand}} \in X_{\hat{h}} \setminus X_f ;$$

for any sample distribution that maintains a nonzero probability over the entire informed set.

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. Lemma 12 follows directly from Lemma 5 given that $X_{\hat{h}} \supseteq X_f$. \square

Theorem 13 (An upper bound on the probability of improving a solution to a holonomic planning problem given knowledge of an admissible informed set). The probability that an iteration of RRT* improves the current solution to a holonomic problem,

c_i , is bounded by the probability of sampling an admissible informed set, $X_{fb} \subseteq X_f$,

$$8i \quad P(c_{i+1} < c_i) = P(x_{rand} \in X_f) = P(x_{rand} \in X_{fb});$$

for any iteration, i , after an initial iterations.

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. Theorem 13 follows directly from Theorem 6 given that $X_{fb} \subseteq X_f$. \square

III. PRIOR WORK ACCELERATING RRT* CONVERGENCE

A review of previous work to improve the convergence rate of RRT* is presented using the results and terminology of Section II. All these techniques attempt to increase the real-time rate of searching the omniscient set by exploiting additional information. Most can be viewed as versions of sample biasing, sample rejection, and/or graph pruning (Sections III-A–III-D).

A. Sample Biasing

Increasing the likelihood of sampling an informed set improves RRT* performance. This sample biasing creates a nonuniform sample distribution that will increase exploration of the informed set but invalidates the assumptions used to prove almost-sure asymptotic optimality. One method to maintain these formal performance guarantees is to calculate the random geometric graph (RGG) connection limit from a subset of samples that are uniformly distributed [16]. This maintains almost-sure asymptotic optimality but increases the required number of rewirings.

It is common to bias sampling around the current solution. This path biasing increases the likelihood of sampling a state that can improve the current solution but reduces the likelihood of finding solutions in other homotopy classes (i.e., it increases precision by decreasing recall; Fig 4a). The ratio of path biasing to global search is frequently a user-chosen parameter that must be tuned for each problem.

Akgun and Stilman [8] use path biasing in their dual-tree version of RRT*. Once an initial solution is found the algorithm spends a user-specified percentage of iterations refining the current solution. It does this by explicitly sampling near a randomly selected state on the current path. This increases the probability of improvement at the expense of decreasing the exploration of other homotopy classes. Their algorithm also employs sample rejection in exploring the state space (see Section III-B).

Nasir et al. [17] combine path biasing with smoothing in their RRT*-Smart algorithm. Solution paths are simplified and then used as biases for further sampling around the solution. Their path smoothing rapidly improves the current solution but the path biasing decreases the likelihood of finding a solution in a different homotopy class.

Kiesel et al. [18] use a two-stage sampling process in their f-biasing technique. Samples are generated by randomly selecting a region of the planning problem and then uniformly sampling it. The probability of selecting a region is calculated by solving




Fig. 4. A illustration of the precision and recall of informed sampling techniques on the omniscient set depicted in Fig. 1(b). The informed sets are coloured to highlight where they are correct (light grey), incorrect (dark grey), or missing the omniscient set (white). Path biasing, (a), generally has high precision but low recall, especially in the presence of multiple homotopy classes. Global or bounded sampling, (b), generally has full recall but low precision, especially in large relative planning problems or high state dimensions. Direct sampling of the L^2 informed set, (c), has full recall and high precision, regardless of the size of the omniscient set and is exactly equal to the omniscient set in the absence of obstacles and constraints.

a simple discretization of the planning problem with Dijkstra's algorithm [1]. The regions along the discrete solution are given a higher selection probability but all regions maintain a nonzero probability to compensate for the incompleteness of the discretization. This technique provides a sampling bias for the entire RRT* search but once a solution is found it continues to sample states that cannot provide a better solution. It is stated that almost-sure asymptotic optimality is maintained but it is not discussed how to modify the rewiring neighbourhood to do so.

Kim et al. [19] also use a two-stage sampling process in their Cloud RRT* algorithm. They generate uniform samples from a series of collision-free, possibly overlapping, spheres defined by a generalized Voronoi graph [20]. New spheres are added on solution paths and the probability of selecting them is updated so that samples from the homotopy class of the solution are biased around the path while maintaining the probability of sampling other homotopy classes. Cloud RRT* successfully finds better solutions faster than other algorithms but continues to sample states that cannot improve the solution and its effect on almost-sure asymptotic optimality is not discussed.

Unlike sample biasing methods, the direct informed sampling used by Informed RRT* does not consider states that are known to be unable to improve a solution. It does result in a nonuniform sample distribution over the problem domain but it is still almost-surely asymptotically optimal as it has a uniform distribution in the informed set being searched.

B. Sample Rejection

Ignoring samples outside an informed set improves RRT* performance. This sample rejection decreases the computational cost of states that cannot improve a solution but does not increase the probability of finding ones that can. If this probability is low (i.e., if the informed set is small relative to the sampling domain) then convergence will not be improved (Fig. 4b). It is shown that this probability decreases factorially with state dimension (i.e., faster than exponentially) in existing formulations of the holonomic minimum-path-length problem (Theorem 14).

Akgun and Stilman [8] use global rejection sampling in addition to sample biasing in their dual-tree algorithm. As samples

are drawn from the entire problem domain, performance will decrease rapidly as the solution improves and/or in large or high-dimensional planning problems.

Otte and Correll [9] use adaptive rejection sampling in their parallelized Coupled Forest of Random Engrafting Search Trees (C-FOREST) algorithm. Samples are generated from a rectangular subset of the planning domain that bounds the ellipsoidal L^2 informed set and rejected using the L^2 heuristic. This increases sampling precision and improves performance in large planning problems but its effectiveness still decreases factorially with state dimension (Theorem 14).

Unlike sample rejection methods, the direct informed sampling used by Informed RRT* maintains high precision and 100% recall regardless of the relative sizes of the informed set and problem domain. It focuses its search in response to solution improvements and does not decrease in effectiveness in large planning domains. It scales more effectively than existing approaches to high-dimensional planning problems

C. Graph Pruning

Limiting the tree to an informed set improves RRT* performance. This graph pruning removes states that can no longer improve the existing solution and reduces the computational cost of basic operations (e.g., nearest neighbour searches). It can also be used to reject potential new states given their connection and any constraints, e.g., (3). After a sufficient number of iterations, this incremental pruning is equivalent in holonomic planning to rejection sampling with the same heuristic (Lemma 12) but with the additional computational costs of expanding towards the sample.

Karaman et al. [21] use graph pruning to implement an online version of RRT* that improves solutions during path execution. They remove vertices whose current cost-to-come plus a heuristic estimate of cost-to-go is higher than the current solution. As current cost-to-come overestimates a vertex's optimal cost-to-come (i.e., it is an inadmissible heuristic), this approach may erroneously remove vertices that could provide a better solution.

Arsalan and Tsiotras [22, 23] combine incremental graph-pruning and incremental graph search techniques with Rapidly exploring Random Graphs (RRG) [6] to reject samples in their RRT[#] algorithm. This incremental pruning focuses the search but its performance will also decrease rapidly as the solution improves or when used on large or high-dimensional planning problems. Some of the rejection criteria also use the current cost-to-come of vertices and may reject samples that could later improve the solution.

Unlike rejecting states with incremental graph pruning, the direct informed sampling used by Informed RRT* wastes no computational effort on states that are known to be unable to improve the solution. Its admissible graph pruning algorithm to remove unnecessary states also only removes vertices from the tree if doing so does not negatively affect the search.

D. Other Techniques

Some techniques to improve RRT/RRT* performance do not fit neatly into the previous categories. Many of these methods could be further accelerated through direct informed sampling.

Urmson and Simmons [24] uses rejection sampling to create a “probabilistic implementation of heuristic search concepts” in their Heuristically Guided RRT (hRRT). At each iteration, a uniformly distributed sample is probabilistically kept or rejected as a function of its heuristic value relative to the existing tree. This iteratively biases RRT expansion towards regions of the problem domain believed to contain high-quality solutions and often finds better solutions than RRT, especially on problems with continuous cost functions (e.g., path length [24]); however, it results in nonuniform sample distributions.

Ferguson and Stentz [7] recognize that an existing solution defines the set of states that could provide better solutions. Their Anytime RRTs algorithm attempts to incrementally find better solutions by searching a decreasing series of these ellipses. This shrinking search ignores some expensive solutions but does not guarantee better ones will be found.

Alterovitz et al. [25] add path refinement to RRT* in their Rapidly exploring Roadmap (RRM) algorithm. Once an initial solution is found, each iteration of RRM either samples a new state or selects an existing state from the current solution and refines it. Path refinement connects the selected state to its neighbours and results in a graph instead of a tree. The ratio of refinement to exploration is a user-tuned parameter.

Shan et al. [26] find an initial solution with RRT, simplify and rewire it using their RRT*_S algorithm, and then continue the search with RRT*. This can find better solutions faster than RRT* alone but the resulting search is not focused and continues to consider states that cannot provide better solutions.

Salzman and Halperin [27] relax performance to asymptotic near optimality in their Lower Bound Tree RRT (LBT-RRT). Rewirings are only considered if they are required to maintain the desired tolerance to the optimum. This can reduce computational complexity but does not focus the search.

Devaurs et al. [28] use ideas from stochastic optimization to explore complex cost functions in their Transition-based RRT* (T-RRT*) and Anytime Transition-based RRT (AT-RRT) algorithms. Transition tests accept or reject a potential new state depending on its cost relative to its parent. These tests help reduce the integral or mechanical work of the path in a cost space; however, for problems seeking to minimize path length are equivalent to graph pruning.

These algorithms, and those designed for more advanced purposes (e.g., RRT^X [29]), can be improved with the direct informed sampling and admissible graph pruning techniques illustrated in Informed RRT*.

E. Direct Informed Sampling for Path Length

This paper presents Informed RRT* as a demonstration of how direct sampling of L^2 informed sets increases the rate at which RRT* improves solutions for problems seeking to minimize path length. Unlike sample biasing, this approach considers all homotopy classes that could provide better solutions (i.e., 100% recall) while maintaining uniform sample distribution over a subplanning problem. Unlike sample rejection or graph pruning, it is effective regardless of the relative size of the informed set or the state dimension (i.e., high precision). In situations where the heuristic does not




Fig. 5. The L^2 informed set, X_ϕ , for a R^2 problem seeking to minimize path length is an ellipse with the initial state, x_{start} , and the goal state, x_{goal} , as focal points. The shape of the ellipse depends on both the initial and goal states, the theoretical minimum cost between the two, c_{min} , and the cost of the best solution found to date, c_i . The eccentricity of the ellipse is given by $c_{\text{min}} = c_i$.

provide substantial information (i.e., small planning problems and/or large informed sets), it performs identically to RRT*.

IV. THE L^2 INFORMED SET

A universally admissible heuristic is well defined for problems seeking to minimize path length in R^n and is commonly used in sampling-based planners (e.g., [7–9]). The cost of a solution constrained to pass through any state, $x \in X$, is bounded from below by the L^2 norm (i.e., Euclidean distance) between it, the start, x_{start} , and the goal, x_{goal} ,

$$\phi(x) = \|x - x_{\text{start}}\|_2 + \|x - x_{\text{goal}}\|_2 : \quad (4)$$

The set of states that could provide a better solution than the current solution cost, c_i , can then be referred to as the L^2 informed set,

$$X_\phi = \{x \in X_{\text{free}} \mid \|x - x_{\text{start}}\|_2 + \|x - x_{\text{goal}}\|_2 < c_i\} :$$

This informed set is a universally admissible estimate of the omniscient set and is exact in the absence of obstacles and constraints (i.e., it is sharp over all minimum-path-length problems). The size of this informed set will decrease as solutions improve.

The L^2 informed set is the intersection of the free space, X_{free} , and a n -dimensional hyperellipsoid symmetric about its transverse axis (i.e., a prolate hyperspheroid),

$$X_\phi = X_{\text{free}} \cap X_{\text{PHS}} ;$$

where

$$X_{\text{PHS}} := \{x \in R^n \mid \|x - x_{\text{start}}\|_2 + \|x - x_{\text{goal}}\|_2 < c_i\} :$$

The prolate hyperspheroid has focal points at x_{start} and x_{goal} , a transverse diameter of c_i , and conjugate diameters of $c_i^2 - c_{\text{min}}^2$, where

$$c_{\text{min}} := \|x_{\text{goal}} - x_{\text{start}}\|_2 ;$$

is the theoretical minimum cost (Fig. 5). The Lebesgue measure of the informed set is

$$|X_\phi| = |X_{\text{PHS}}| = \frac{c_i (c_i^2 - c_{\text{min}}^2)^{\frac{n-1}{2}}}{2^n} ; \quad (5)$$

where $| \cdot |$ is the Lebesgue measure of a n -dimensional unit ball,

$$| \cdot | := \frac{\pi^{\frac{n}{2}}}{\Gamma(\frac{n}{2} + 1)} ; \quad (6)$$

and $\Gamma(\cdot)$ is the gamma function, an extension of factorials to real numbers [30].

The probability of uniformly sampling this informed set by sampling any superset (e.g., a bounding box), $X_{\text{samp}} \supset X_\phi$, can be written as a ratio of measures,

$$P(x_{\text{rand}} \in X_\phi \mid x_{\text{rand}} \in X_{\text{samp}}) = \frac{|X_{\text{PHS}}|}{|X_{\text{samp}}|} = \frac{\frac{\pi^{\frac{n}{2}}}{2^n} c_i (c_i^2 - c_{\text{min}}^2)^{\frac{n-1}{2}}}{\frac{\pi^{\frac{n}{2}}}{2^n} \Gamma(\frac{n}{2} + 1) |X_{\text{samp}}|} ; \quad (7)$$

which can be combined with Theorem 13 to bound the probability of improving a solution to a holonomic problem,

$$8i \quad P(c_{i+1} < c_i \mid x_{\text{rand}} \in X_{\text{samp}}) \leq \frac{\frac{\pi^{\frac{n}{2}}}{2^n} c_i (c_i^2 - c_{\text{min}}^2)^{\frac{n-1}{2}}}{\frac{\pi^{\frac{n}{2}}}{2^n} \Gamma(\frac{n}{2} + 1) |X_{\text{samp}}|} ; \quad (8)$$

This probability becomes arbitrarily small for (i) costs, c_i , near the theoretical limit, c_{min} , (ii) large sampling domains, $|X_{\text{samp}}|$, or (iii) high state dimensions, n . While the solution cost and sampling domain size may vary during the search of a problem, the state dimension is constant throughout. This motivates investigating the effect of state dimension on existing formulations of the holonomic minimum-path-length planning problem (Theorem 14).

Theorem 14 (The minimum-path-length curse of dimensionality). The probability that RRT* improves a solution to holonomic problems seeking to minimize path length decreases factorially (i.e., faster than exponentially) as state dimension increases,

$$8i \quad P(c_{i+1} < c_i \mid x_{\text{rand}} \in X_{\text{rect}}) \leq \frac{\frac{\pi^{\frac{n}{2}}}{2^n}}{\Gamma(\frac{n}{2} + 1)} ; \quad (9)$$

when uniformly sampling a (hyper)rectangle bounding the L^2 informed set, $X_{\text{rect}} \supset X_{\text{PHS}} \supset X_\phi \supset X_f$.

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. Theorem 14 is proven for RRT* but holds for any algorithm for which an equivalent to Theorem 13 exists.

The smallest possible X_{rect} that completely contains X_{PHS} is a (hyper)rectangle with widths corresponding to the diameters of the prolate hyperspheroid (Fig. 6a). The measure of any $X_{\text{rect}} \supset X_{\text{PHS}}$ is therefore bounded from below as

$$|X_{\text{rect}}| \geq c_i (c_i^2 - c_{\text{min}}^2)^{\frac{n-1}{2}} ; \quad (10)$$

When substituted into (8) this gives

$$8i \quad P(c_{i+1} < c_i \mid x_{\text{rand}} \in X_{\text{rect}}) \leq \frac{\frac{\pi^{\frac{n}{2}}}{2^n}}{\Gamma(\frac{n}{2} + 1)} ;$$

proving Theorem 14 for all rectangular sets, X_{rect} , such that $X_{\text{rect}} \supset X_{\text{PHS}} \supset X_\phi \supset X_f$. \square

Theorem 14 is an upper bound on the utility of rectangular rejection sampling in holonomic planning and is illustrated




Fig. 6. An illustration of state dimension on problems seeking to minimize path length. The best case performance of an admissible rectangular sampling, e.g., [9], occurs when the rectangle tightly bounds the prolate hyperspheroid defined by the current solution cost, $X_{\text{rect}} \supset X_{\text{PHS}} \supset X_{\text{pb}}$, (a). The probability of sampling this L^2 informed set (i.e., its relative measure) decreases factorially (i.e., faster than exponentially) with state dimension, n , (b), meaning that existing formulations of RRT* do not scale effectively to high state dimensions. Direct informed sampling, Alg. 1, scales more efficiently as illustrated by the average per-sample time versus state dimension, (c). Samples in the unit n -ball for Alg. 2 were generated with Boost 1.58.

by plotting (9) versus state dimension (Fig. 6b). The results show that while rectangular rejection sampling may be 79% successful in R^2 , its success decreases factorially as state dimension increases and is only 2% in R^8 and $4 \cdot 10^{-4}\%$ in R^{16} . These numbers represent the best-case for rectangular rejection sampling and actual performance will depend on the size and orientation of the informed set relative to the sampling domain. This motivates a need for a direct method to sample the prolate hyperspheroid regardless of size, orientation, and state dimension.

A. Direct Sampling

A direct method to generate uniformly distributed samples in the L^2 informed set is adapted from techniques to sample hyperellipsoids [31].

Let $S \in R^{n \times n}$ be a symmetric, positive-definite matrix (the hyperellipsoid matrix) such that the interior of a hyperellipsoid, X_{ellipse} , is defined as

$$X_{\text{ellipse}} := \{x \in R^n \mid (x - x_{\text{centre}})^T S^{-1} (x - x_{\text{centre}}) < 1\}; \quad (11)$$

where x_{centre} is the centre point of the hyperellipsoid. Uniformly distributed samples in the hyperellipsoid, $x_{\text{ellipse}} \in U(X_{\text{ellipse}})$, can be generated from uniformly distributed samples in the interior of a unit n -dimensional ball, $x_{\text{ball}} \in U(X_{\text{ball}})$, by

$$x_{\text{ellipse}} = Lx_{\text{ball}} + x_{\text{centre}}; \quad (12)$$

where $L \in R^{n \times n}$ is the lower-triangular Cholesky decomposition of the hyperellipsoid matrix such that

$$LL^T = S$$

and

$$X_{\text{ball}} := \{x \in R^n \mid \|x\|_2 < 1\};$$

For hyperellipsoids with orthogonal axes, there exists a coordinate frame in which the hyperellipsoid matrix is diagonal,

$$S^0 := \text{diag}(r_1^2; r_2^2; \dots; r_n^2);$$

where r_j is the radius of j -th axis of the hyperellipsoid and $\text{diag}(\cdot)$ constructs a diagonal matrix. A rotation from

this hyperellipsoid-aligned frame to the world frame, $C_{\text{we}} \in SO(n)$, can be used to write (11) in terms of S^0 as

$$X_{\text{ellipse}} := \{x \in R^n \mid (x - x_{\text{centre}})^T C_{\text{we}} S^{0^{-1}} C_{\text{we}}^T (x - x_{\text{centre}}) < 1\};$$

and (12) as

$$x_{\text{ellipse}} = C_{\text{we}} L^0 x_{\text{ball}} + x_{\text{centre}}; \quad (13)$$

given the orthogonality of rotation matrices, $C_{\text{we}}^{-1} = C_{\text{we}}^T$, and that $L^0 L^0{}^T = S^0$.

The rotation between frames can be solved directly as a general Wahba problem [32] even when underspecified [33]. Generally, the rotation matrix from one set of axes, $f a_j g$, to another set of axes, $f b_j g$, is given by

$$C_{ba} = U V^T; \quad (14)$$

where $U, V \in R^{n \times n}$ is

$$U := \text{diag}(1; \dots; 1; \det(U) \det(V));$$

and $\det(\cdot)$ is the matrix determinant. The terms $U \in R^{n \times n}$ and $V \in R^{n \times n}$ are unitary matrices such that $U V^T = M$ via singular value decomposition and $M \in R^{n \times n}$ is given by the outer product of the $j = 1 \dots n$ corresponding axes,

$$M := [a_1; a_2; \dots; a_j] [b_1; b_2; \dots; b_j]^T; \quad (15)$$

In problems seeking to minimize path length, the hyperellipsoid is a prolate hyperspheroid described by

$$x_{\text{centre}} := \frac{x_{\text{start}} + x_{\text{goal}}}{2}; \quad (16)$$

$$S^0 := \text{diag} \left(\frac{c_1^2}{4}; \frac{c_1^2 - c_{\min}^2}{4}; \dots; \frac{c_1^2 - c_{\min}^2}{4} \right);$$

and therefore,

$$L^0 = \text{diag} \left(\frac{c_1}{2}; \sqrt{\frac{c_1^2 - c_{\min}^2}{4}}; \dots; \sqrt{\frac{c_1^2 - c_{\min}^2}{4}} \right); \quad (17)$$

Its local coordinate system is underspecified in the conjugate directions due to symmetry, making (15) just

$$M = a_1 1_1^T; \quad (18)$$

where 1_1 the first column of the identity matrix and the transverse axis in the world frame is

$$a_1 = (x_{\text{goal}} - x_{\text{start}}) = \|x_{\text{goal}} - x_{\text{start}}\|_2 \cdot k_2;$$

Alg. 1: Sample ($x_{\text{start}} \in X$; $x_{\text{goal}} \in X$; $c_i \in \mathbb{R}_{>0}$)

```

1 repeat
2   if  $\|x_{\text{PHS}} - x_{\text{start}}\|_2 < c_i$  then
3      $x_{\text{rand}} \leftarrow \text{SamplePHS}(x_{\text{start}}; x_{\text{goal}}; c_i)$ ;
4   else
5      $x_{\text{rand}} \leftarrow \text{SampleProblem}(X)$ ;
6 until  $x_{\text{rand}} \in X_{\text{free}} \setminus X_{\text{PHS}}$ ;
7 return  $x_{\text{rand}}$ ;

```

Alg. 2: SamplePHS ($x_{\text{start}} \in X$; $x_{\text{goal}} \in X$; $c_i \in \mathbb{R}_{>0}$)

```

1  $c_{\min} \leftarrow \|x_{\text{goal}} - x_{\text{start}}\|_2$ ;
2  $x_{\text{centre}} \leftarrow (x_{\text{start}} + x_{\text{goal}})/2$ ;
3  $a_1 \leftarrow (x_{\text{goal}} - x_{\text{start}})/c_{\min}$ ;
4  $[U; V] \leftarrow \text{SVD}(a_1 a_1^T)$ ;
5  $\text{diag}(1; \dots; 1; \det(U) \det(V))$ ;
6  $C_{\text{we}} \leftarrow U V^T$ ;
7  $r_1 \leftarrow c_i/2$ ;
8 for  $j = 2; \dots; n$ 
9    $P \leftarrow \frac{c_i^2 - c_{\min}^2}{c_i^2 - c_{\min}^2} = 2$ ;
10   $L \leftarrow \text{diag}(r_1; r_2; \dots; r_n)$ ;
11  $x_{\text{ball}} \leftarrow \text{SampleUnitBall}(n)$ ;
12  $x_{\text{rand}} \leftarrow C_{\text{we}} L x_{\text{ball}} + x_{\text{centre}}$ ;
13 return  $x_{\text{rand}}$ ;

```

Samples distributed uniformly in the L^2 informed set, $X_{\text{fb}} = X_{\text{PHS}} \setminus X_{\text{free}}$, can therefore be generated by using (13) to transform samples drawn uniformly from a unit n -ball. These samples are mapped to the prolate hyperspheroid through scaling, (17), rotation, (14) and (18), and translation, (16).

Sun and Farooq [31] investigate various methods to generate samples in hyperellipsoids and provide the following lemma regarding the uniform sample density of this technique.

Lemma 15 (The uniform distribution of samples transformed into a hyperellipsoid from a unit n -ball. Originally Lemma 1 in [31]). If the random points distributed in a hyperellipsoid are generated from the random points uniformly distributed in a hypersphere through a linear invertible nonorthogonal transformation, then the random points distributed in the hyperellipsoid are also uniformly distributed.

Proof. For brevity, [31] only presents anecdotal proofs of Lemma 15. The full proof from the supplementary online material appears in Appendix B. \square

1) Algorithm: The L^2 informed set is an arbitrary intersection of the prolate hyperspheroid and the problem domain. It can be sampled efficiently by considering the relative measure of the two sets and sampling the smaller set until a sample belonging to both sets is found. These procedures are presented algorithmically in Algs. 1 and 2 and are publicly available in OMPL. Note that for most problems Alg. 2, Lines 1–6 are constant and only need to be calculated once.

The function $\text{SVD}(\cdot)$ denotes the singular value decomposition of a matrix and $\text{SampleUnitBall}(n)$ returns uniformly distributed samples from the interior of an n -dimensional unit ball. The measure of the prolate hyperspheroid, $\|x_{\text{PHS}}\|_2$, is

given by (5) and SampleProblem returns samples uniformly distributed over the entire planning domain. Implementations of SVD and SampleUnitBall can be found in common C++ libraries.

2) Practical Performance: Direct informed sampling (Alg. 1) is compared to the best-case performance of rectangular rejection sampling. The average computational time required to find a sample in the L^2 informed set is calculated by generating 10^6 samples at each dimension (Fig. 6c). The results show that while rejection sampling may outperform direct informed sampling in low state dimensions (e.g., R^2 : 7.3×10^{-8} vs. 3.5×10^{-7} seconds), it becomes orders of magnitude slower as state dimension increases (e.g., R^{16} : 4.0×10^{-2} vs. 7.2×10^{-7} seconds). These per-sample times are small but significant. Generating 10^5 samples in R^{16} requires less than a second with direct informed sampling (7.2×10^{-2} seconds) but over an hour with rectangular rejection sampling (3953 seconds).

This experiment represents optimistic results for both constant (e.g., the problem domain) and adaptive (e.g., [9]) rectangular rejection sampling. Constant sampling domains rarely provide tight bounds on the informed set and will generally have higher rejection rates than the experiment. Adaptive sampling domains may tightly bound the informed set but must account for its alignment relative to the state space. This requires either a larger rectangular sampling domain or a rotation between frames that increases the rejection rate or computational cost compared to the experiment, respectively.

B. Extension to Multiple Goals

Many planning problems seek the minimum-length path that connects a start to any state in a goal region, X_{goal} . In these situations the omniscient set is all states that could provide a better solution to any goal. The multigoal L^2 informed set is

$$X_{\text{fb}} := \{x \in X_{\text{free}} \mid \|x - x_{\text{start}}\|_2 + \|x - x_{\text{goal};j}\|_2 < c_i \text{ for any } x_{\text{goal};j} \in X_{\text{goal}}\};$$

For a countable goal region, $X_{\text{goal}} := \{x_{\text{goal};j} \mid j = 1, \dots, z\}$, this set is the union of the individual informed sets of each goal,

$$X_{\text{fb}} = \bigcup_{j=1}^z X_{\text{fb};j};$$

where z is the number of goals and

$$X_{\text{fb};j} := \{x \in X_{\text{free}} \mid \|x - x_{\text{start}}\|_2 + \|x - x_{\text{goal};j}\|_2 < c_i\};$$

is the L^2 informed set of an individual $(x_{\text{start}}; x_{\text{goal};j})$ pair. If the individual informed sets do not intersect, then a uniform sample distribution can be generated by randomly selecting an individual subset, j , in proportion to its relative measure,

$$p(j \mid z) := \frac{|X_{\text{fb};j}|}{\sum_{k=1}^z |X_{\text{fb};k}|};$$

and then generating a uniformly distributed sample inside the selected subset, $X_{\text{fb};j}$.

If individual sets do intersect, then this approach will oversample states that belong to multiple sets (Fig. 7a). In

Alg. 3: **Sample**($x_{\text{start}} \in X; X_{\text{goal}} \subseteq X; c_i \in \mathbb{R}_{>0}$)

```

1 repeat  $\frac{1}{z} \sum_{j=1}^z (X_{\text{PHS};j}) < (X)$  then
2   if  $\frac{1}{z} \sum_{j=1}^z (X_{\text{PHS};j}) < (X)$  then
3      $x_{\text{goal};j} \leftarrow \text{RandomGoal}(x_{\text{start}}; X_{\text{goal}}; c_i);$ 
4      $x_{\text{rand}} \leftarrow \text{SamplePHS}(x_{\text{start}}; x_{\text{goal};j}; c_i);$ 
5   else
6      $x_{\text{rand}} \leftarrow \text{SampleProblem}(X);$ 
7 until  $x_{\text{rand}} \in X_{\text{free}} \setminus \bigcup_{j=1}^z X_{\text{PHS};j}$ 
   and
    $\text{KeepSample}(x_{\text{rand}}; x_{\text{start}}; X_{\text{goal}}; c_i);$ 
8 return  $x_{\text{rand}};$ 

```

Alg. 4: **RandomGoal**($x_{\text{start}} \in X; X_{\text{goal}} \subseteq X; c_i \in \mathbb{R}_{>0}$)

```

1 a ← 0;
2 forall  $x_{\text{goal};k} \in X_{\text{goal}}$  do
3   a ← a +  $(X_{\text{PHS};k})$ ;
4 p ← U[0; 1];
5 j ← 0;
6 repeat
7   j ← j + 1;
8   p ← p -  $(X_{\text{PHS};j})$ ;
9 until p ≤ 0;
10 return  $x_{\text{goal};j};$ 

```

these situations, uniform sample density can be maintained by probabilistically rejecting samples in proportion to their membership in individual sets. This creates a uniform sample distribution for multigoal L^2 informed sets defined by arbitrarily overlapping individual informed sets (Fig. 7b).

1) Algorithm: The algorithm is described in Algs. 3–5 as modifications to the sampling technique for a single-goal L^2 informed set, with changes highlighted in red (cf. Alg. 1). The measure of individual informed sets, $(X_{\text{PHS};j})$, is calculated from (5) using the appropriate goal, $x_{\text{goal};j}$. This same technique can also be applied to problems with a countable start region.

V. INFORMED RRT*

Informed RRT* is an extension of RRT* that demonstrates how informed sets can be used to improve anytime almost-surely asymptotically optimal planning. It performs the same as RRT* until a solution is found after which the search is focused to the informed set through direct informed sampling and admissible graph pruning (Fig. 8). This increases the likelihood of sampling states that can improve the solution and increases the convergence rate towards the optimum regardless of the relative size of the informed set (e.g., near-minimum solutions or large problem domains) or the state dimension.

Informed RRT* uses direct informed sampling (Alg. 3), admissible graph pruning (Section V-B), and an updated calculation of the rewiring neighbourhood (Section V-C) to focus the search. The complete algorithm is presented in Algs. 6 and 7 as modifications to RRT*, with changes highlighted in red. It can also be integrated into other sampling-based planners, such as RRT^X [29] and Batch Informed Trees (BIT*) [34–36].

Alg. 5: **KeepSample**($x_{\text{rand}} \in X; x_{\text{start}} \in X; X_{\text{goal}} \subseteq X; c_i \in \mathbb{R}_{>0}$)

```

1 a ← 0;
2 forall  $x_{\text{goal};k} \in X_{\text{goal}}$  do
3   if  $kx_{\text{rand}} - x_{\text{start}}k_2 + kx_{\text{goal};k} - x_{\text{rand}}k_2 < c_i$  then
4     a ← a + 1;
5 p ← U[0; 1];
6 return p ≤ a;

```

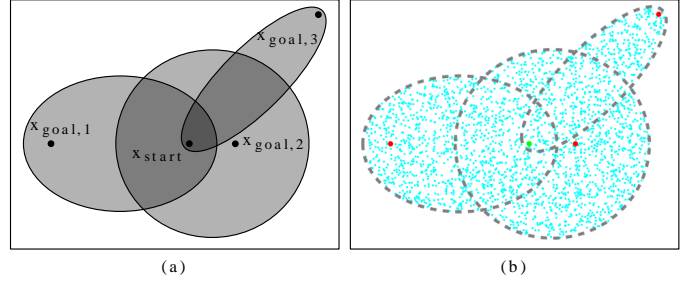


Fig. 7. An illustration of the multigoal L^2 informed set for a problem seeking to minimize path length from a start at $[0; 0]^T$, to any of three goals at $[-0.75; 0]^T$, $[0.25; 0]^T$, and $[0.7; 0.7]^T$, and a current solution cost of $c_i = 1.05$. Each ellipse illustrates the L^2 informed set for a start-goal pair. Combining the uniform distributions of these individuals (light grey) would result in a nonuniform distribution (dark grey). (a). By probabilistically rejecting samples in proportion to their individual membership, Alg. 3 uniformly samples complex sets of arbitrary intersections, as illustrated with 2500 random samples, (b).

At each iteration, Informed RRT* calculates the current best solution (Alg. 6, Line 4) from the vertices in the goal region (Alg. 6, Lines 2, 9–10). This defines a shrinking L^2 informed set that is used to both focus sampling (Alg. 6, Line 5; Alg. 3) and prune the graph (Alg. 6, Line 27; Alg. 7). This process continues for as long as time allows or until a suitable solution is found.

Informed RRT* retains the probabilistic completeness and almost-sure asymptotically optimality of RRT*. It is probabilistically complete since it does not modify the search for an initial solution. It is almost-surely asymptotically optimal as it maintains a uniform sample distribution over a subset of the planning problem in which it uses a local rewiring neighbourhood that satisfies the bounds presented in [6].

A. Notation

The tree, $T := (V; E)$, is defined by a set of vertices, V , X_{free} , and edges, $E = \{f(v; w)g\}$, for some $v, w \in V$. The function $g_T(v)$ represents the cost to reach a vertex, $v \in V$, from the start given the current tree (the cost-to-come). The function $c(v; w)$ represents the cost of a path connecting the states $v, w \in X_{\text{free}}$, and corresponds to the edge cost between those two states if they are connected as vertices in the tree. The notation $X \stackrel{+}{=} f x g$ and $X \stackrel{-}{=} f x g$ is used to compactly represent the compounding set operations $X \leftarrow X \cup f x g$ and $X \leftarrow X \cap f x g$, respectively. As is customary, the minimum of an empty set is taken to be infinity and a prolate hyperspheroid defined by an infinite transverse diameter is taken to have infinite measure.




Fig. 8. An example of how Informed RRT* uses the current solution to focus search to the L^2 informed set. After an unfocused search for an initial solution, (a), Informed RRT* prunes the graph of unnecessary states and redefines the search domain to the current L^2 informed set, (b). Samples are then generated directly from this informed set, avoiding those that are known to be unable to improve the current solution. This reduced search space increases the likelihood of finding an improved solution, which in turn further reduces the search space, (c). This results in an algorithm that focuses the search to the subproblem given by the current solution, shown enlarged in (d), even as the subproblem decreases with further improvement.

Alg. 6: Informed RRT*($x_{start} \in X_{free}; X_{goal} \subseteq X$)

```

1  $V \leftarrow f(x_{start}, g; E); T \leftarrow (V; E);$ 
2  $V_{sol'n} \leftarrow \emptyset;$ 
3 for  $i = 1 \dots q$  do
4    $c_i \leftarrow \min_{v \in V_{sol'n}} f(g_T(v, goal); g);$ 
5    $x_{rand} \leftarrow \text{Sample}(x_{start}; X_{goal}; c_i);$ 
6    $v_{nearest} \leftarrow \text{Nearest}(V; x_{rand});$ 
7    $x_{new} \leftarrow \text{Steer}(v_{nearest}; x_{rand});$ 
8   if  $\text{IsFree}(v_{nearest}; x_{new})$  then
9     if  $x_{new} \in X_{goal}$  then
10       $V_{sol'n} \leftarrow V_{sol'n} \cup x_{new};$ 
11       $V \leftarrow V \cup x_{new};$ 
12       $V_{near} \leftarrow \text{Near}(V; x_{new}; r_{rewire});$ 
13       $v_{min} \leftarrow v_{nearest};$ 
14      for all  $v_{near} \in V_{near}$  do
15         $c_{new} \leftarrow g_T(v_{near}) + c(v_{near}; x_{new});$ 
16        if  $c_{new} < g_T(v_{min}) + c(v_{min}; x_{new})$  then
17          if  $\text{IsFree}(v_{near}; x_{new})$  then
18             $v_{min} \leftarrow v_{near};$ 
19       $E \leftarrow E \cup f(v_{min}; x_{new});$ 
20      for all  $v_{near} \in V_{near}$  do
21         $c_{near} \leftarrow g_T(x_{new}) + c(x_{new}; v_{near});$ 
22        if  $c_{near} < g_T(v_{near})$  then
23          if  $\text{IsFree}(x_{new}; v_{near})$  then
24             $v_{parent} \leftarrow \text{Parent}(v_{near});$ 
25             $E \leftarrow E \cup f(v_{parent}; v_{near});$ 
26             $E \leftarrow E \cup f(x_{new}; v_{near});$ 
27       $\text{Prune}(V; E; c_i);$ 
28 return  $T;$ 

```

B. Graph Pruning (Alg. 7)

Graph pruning simplifies a tree by removing unnecessary vertices. Vertices are often removed if their heuristic values are larger than the current solution (i.e., they do not belong to informed set). While this identifies vertices that cannot provide

Alg. 7: Prune ($V \subseteq X; E \subseteq V \times V; c_i \in \mathbb{R}_{\geq 0}$)

```

1 repeat
2    $V_{prune} \leftarrow \emptyset;$ 
3   for all  $v \in V$  do
4     if  $\phi(v) > c_i$ ; and  $\exists w \in V; (v; w) \in E$ ;
5   until  $V_{prune} = \emptyset;$ 

```

a better solution, it is not a sufficient condition to remove them without negatively affecting the search. Their descendants may still be capable of providing better solutions (i.e., they may belong to the informed set; Fig. 9) in which case their removal would negatively affect performance by decreasing vertex density in the search domain (i.e., the informed set; Fig. 10b).

An admissible pruning method that does not remove vertices from the informed set is presented in Alg. 7. It iteratively removes leaves of the tree that cannot provide a better solution until no such leaves exist. This only removes vertices if they and all their descendants cannot belong to a better solution (i.e., it only removes vertices from outside the informed set; Fig. 9). This retains all possibly beneficial vertices regardless of their current connections and does not alter the vertex distribution in areas being searched (Fig. 10c).

C. The Rewiring Neighbourhood

RRT* almost-surely converges asymptotically to the optimum by incrementally rewiring the tree around new states. In the r -disc variant this is the set of states within a radius, r_{rewire} , of the new state,

$$r_{rewire} \triangleq \min \{ r_{RRT}, g \}; \quad (19)$$

where r_{RRT} is the maximum allowable edge length of the tree and r_{RRT} is a function of the problem measure and the number




Fig. 9. An illustration of Alg. 7 that shows the retained (black) and pruned (grey) vertices given the L^2 informed set (dashed grey line) defined by the current solution (i.e., they are not members of the L^2 informed set) and neither can their descendants. This pruning condition avoids removing promising vertices (e.g., v) simply because they are currently descendants of vertices outside the subset (e.g., u) and maintains the vertex distribution of the L^2 informed set (Fig. 10).

of vertices in the tree [6]. Specifically,

$$r_{\text{RRT}} > r_{\text{RRT}} ;$$

$$r_{\text{RRT}} := 2 \left(1 + \frac{1}{n} \frac{(X)}{n} \frac{\log(jVj)}{jVj} \right)^{\frac{1}{n}} ; \quad (20)$$

where (\cdot) is the Lebesgue measure of a set (e.g., the volume), n is the Lebesgue measure of an n -dimensional unit ball, i.e., (6), and jVj is the cardinality of a set.

The rewiring neighbourhood in the k -nearest variant is the k_{RRT} -closest states to the new state, where

$$k_{\text{RRT}} > k_{\text{RRT}} ;$$

$$k_{\text{RRT}} := e \left(1 + \frac{1}{n} \log(jVj) \right) ; \quad (21)$$

Informed RRT* searches a subset of the original planning problem. The rewiring requirements to maintain almost-sure asymptotic optimality in this shrinking domain will be a function of the number of vertices in the informed set, $V \setminus X_{\text{fb}}$, and its measure, (X_{fb}) . The L^2 informed set is not known in closed form (it is an intersection of a prolate hyperspheroid and free space) but its measure can be bounded from above by the minimum measure of the prolate hyperspheroid and the problem domain,

$$X_{\text{fb}} = \min f(X); (X_{\text{PHS}})g;$$

This updates (20) and (21) to

$$r_{\text{RRT}} @2 \left(1 + \frac{1}{n} \frac{\min f(X); (X_{\text{PHS}})g}{n} \right)^{\frac{1}{n}}$$

$$@ \frac{\log(V \setminus X_{\text{fb}})}{V \setminus X_{\text{fb}}} A A \quad (22)$$

and

$$k_{\text{RRT}} = e \left(1 + \frac{1}{n} \log(V \setminus X_{\text{fb}}) \right) ; \quad (23)$$

where (X_{PHS}) is a function of the current solution, i.e., (5).

These rewiring neighbourhoods will be smaller than (20) and (21) when they can contain fewer vertices (i.e., only those in the informed set) and/or a smaller problem measure (i.e., the measure of the informed set). Smaller rewiring neighbourhoods




Fig. 10. An illustration of pruning a graph found by RRT*, (a), with both inadmissible, (b), and admissible, (c), methods. By removing all vertices that cannot belong to a better solution, the inadmissible method may greedily remove descendant vertices that will later provide a better solution once the graph is improved. By only removing vertices that cannot improve a solution if neither can their descendants, the admissible method (Alg. 7) maintains a uniform sample density in the entire informed set.

reduce the computational cost of rewiring at each iteration and improves the real-time performance of Informed RRT* while maintaining almost-sure asymptotic optimality.

VI. RATES OF CONVERGENCE

Almost-sure asymptotic optimality provides no insight into the rate at which solutions are improved. Previous work has found probabilistic rates for PRM* [37] and Fast Marching Tree (FMT*) [16] and estimated the expected length of RRT* solutions as a function of computational time [37].

Performance can be quantified analytically by evaluating the rate at which the sequence of solution costs converges to the optimum. This rate can be classified as sublinear, linear, or superlinear (Definition 16).

Definition 16 (Rate of convergence). A sequence of numbers, $(a_i)_{i=1}^{\infty}$, that monotonically and asymptotically approaches a limit, a_1 , has a rate of convergence given by

$$:= \lim_{i \rightarrow \infty} \frac{|a_{i+1} - a_1|}{|a_i - a_1|}.$$

The sequence is said to converge linearly if the rate is in the range $0 < \alpha < 1$, superlinearly (i.e., faster than linear) when $\alpha = 0$, and sublinearly (i.e., slower than linear) when $\alpha = 1$.

The expected convergence rate of an algorithm depends on its tuning and the planning problem. General rates can be calculated for holonomic minimum-path-length problems for RRT* with and without sample rejection and Informed RRT* (Theorems 18–20) by first calculating sharp bounds on the expected next-iteration cost (Lemma 17).

Lemma 17 (Expected next-iteration cost of minimum-path-length planning). The expected value of the next solution to a minimum-path-length problem, $E[c_{i+1}]$, is bounded by

$$Pr \frac{nc_i^2 + c_{\min}^2}{(n+1)c_i} + (1 - p_f) c_i \leq E[c_{i+1}] \leq c_i ; \quad (24)$$

where c_i is the current solution cost, c_{\min} is the theoretical minimum solution cost, n is the state dimension of the planning problem, and $p_f = P(x_{\text{new}} \in X_f)$ is the probability of adding a state that is a member of the omniscient set (i.e., that can belong to a better solution). While not explicitly shown, the subset, X_f , and the probability of improving the solution, p_f , are generally functions of the current solution cost.

This lower bound is sharp over the set of all possible minimum-path-length planning problems and algorithm configurations and is exact for versions of RRT* with an infinite rewiring radius (i.e., $r = 1$, and $r_{\text{RRT}} = 1$) searching an obstacle-free environment without constraints.

Proof. The proof of Lemma 17 from the supplementary online material appears in Appendix C. \square

This result allows sharp bounds on the convergence rates of RRT* (with and without rejection sampling) and Informed RRT* to be calculated for any configuration or holonomic minimum-path-length planning problem. These bounds will be exact in problems without obstacles and constraints and with an infinite rewiring neighbourhood (i.e., $r = 1$, and $r_{\text{RRT}} = 1$) and show that RRT* always has sublinear convergence to the optimum (Theorem 18).

Theorem 18 (Sublinear convergence of RRT* in holonomic minimum-path-length planning). RRT* converges sublinearly towards the optimum of holonomic minimum-path-length planning problems,

$$E[\text{RRT}] = 1: \quad (25)$$

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. The proof of Theorem 18 follows directly from Lemma 17 when p_f is given by (7) and appears from the supplementary online material in Appendix D-A. \square

Rectangular rejection sampling improves the convergence rate of RRT*. This improvement is maximized by sampling a rectangle that tightly bounds the informed set (Fig. 6a). The resulting adaptive rectangular rejection sampling (e.g., [9]) allows RRT* to converge linearly in the absence of obstacles and constraints and with an infinite rewiring neighbourhood (Theorem 19).

Theorem 19 (Linear convergence of RRT* with adaptive rectangular rejection sampling in holonomic minimum-path-length planning). RRT* with adaptive rectangular rejection sampling converges at best linearly towards the optimum of holonomic minimum-path-length planning problems but factorially approaches sublinear convergence with increasing state dimension,

$$1 - \frac{\frac{n}{2}}{(n+1)2^{n-1}\Gamma(\frac{n}{2}+1)} = E[\text{Rect}] = 1: \quad (26)$$

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. The proof of Theorem 19 follows directly from Lemma 17 when p_f is calculated by substituting (10) in (7) and appears from the supplementary online material in Appendix D-B. \square

This convergence rate diminishes factorially (i.e., quickly) as state dimension increases due to the minimum-path-length curse of dimensionality. Informed RRT* avoids this limitation




Fig. 11. An illustration of the lower-bounds on linearity, $E[1 - \mu^*]$, of RRT* with rejection sampling and Informed RRT* (Corollary 21). As predicted by Theorems 19 and 20, the convergence rates bounds diverge as state dimensions increase, with rejection sampling factorially approaching sublinear convergence.

with direct informed sampling. It also converges linearly in the absence of obstacles and constraints and with an infinite rewiring neighbourhood but has a weaker dependence on state dimension (Theorem 20).

Theorem 20 (Linear convergence of Informed RRT* in holonomic minimum-path-length planning). Informed RRT* converges at best linearly towards the optimum of holonomic minimum-path-length planning problems,

$$\frac{n-1}{n+1} = E[\text{Inf}] = 1: \quad (27)$$

where the lower-bound occurs exactly with an infinite rewiring neighbourhood in the absence of obstacles and constraints.

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. The proof of Theorem 20 follows directly from Lemma 17 when $p_f = 1$ and appears from the supplementary online material in Appendix D-C. \square

Theorems 18–20 result in the following corollary regarding the relative convergence rates of the algorithms.

Corollary 21 (The faster convergence of Informed RRT* in holonomic minimum-path-length planning). The best-case convergence rate of Informed RRT*, $E[\text{Inf}]$, is always better than that of RRT*, with or without rejection sampling in holonomic minimum-path-length planning,

$$8n - 2; \frac{n-1}{n+1} = E[\text{Inf}] = E[\text{Rect}] = E[\text{RRT}] = 1:$$

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. The proof follows immediately from the lower bounds in (25), (26), and (27). It is illustrated in Fig. 11. \square

A. Experimental Validation and Extension

Convergence rates are investigated experimentally for infinite, constant finite, and decreasing finite rewiring radii. To isolate the effects of the rewiring parameters, Informed RRT* was run on obstacle- and constraint-free problems in \mathbb{R}^2 , \mathbb{R}^4 , and \mathbb{R}^8 for 10^4 trials of each configuration. Each trial started from the

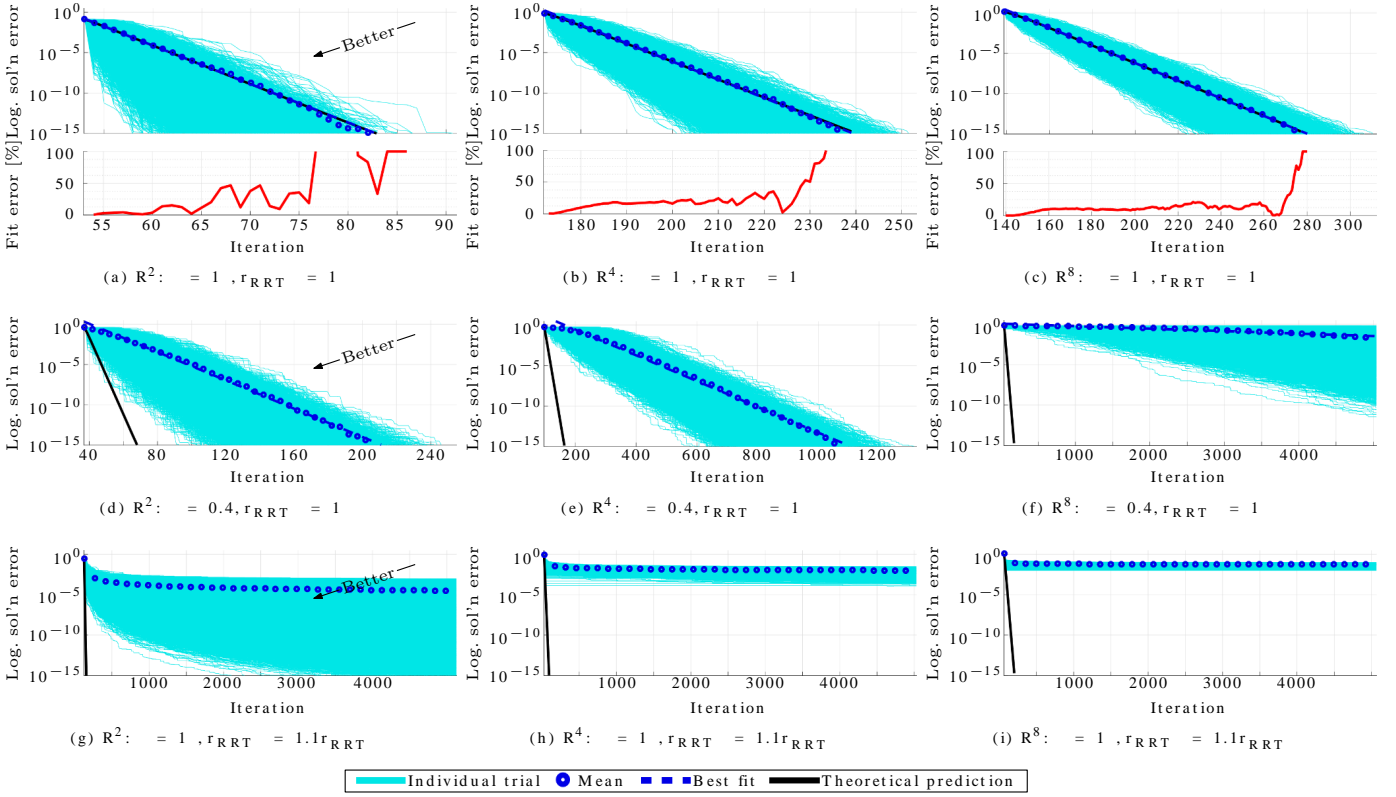


Fig. 12. Experimental validation and extension of Lemma 17 and Theorem 20 in R^2 , R^4 and R^8 . Informed RRT* was run from a common initial solution 10^4 times in R^2 , R^4 and R^8 with different pseudo-random seeds. The error relative to the known optimum, $\log(c_i - c^*)$, was plotted for each instance at each iteration (cyan lines) along with the mean error (blue circles), a line of best fit (blue dashed line), and the lower-bound error predicted by Lemma 17 (black line). The difference between the predicted lower bound and the mean errors (red lines), $(c_{\text{mean};i} - c_{\text{theory};i} = (c_{\text{mean};i} - c^*))$, with infinite rewiring neighbourhoods, (a)–(c), confirms experimentally that convergence is linear (Theorem 20). The mean error for a finite but constant rewiring neighbourhood, (d)–(f), shows experimentally that convergence is slower but possibly still linear. The mean error for a finite and decreasing rewiring neighbourhood, (g)–(i), shows experimentally that it is slower and sublinear. The results of (d)–(i) motivate further research on the effects of the RRT* rewiring neighbourhood.

same initial solution but used different pseudo-random seeds to search for improvements. The logarithmic error relative to the known optimum, $\log(c_i - c^*)$, and the resulting mean were calculated at each iteration of each trial and used to validate Theorem 20 and illustrate the effects of rewiring parameters on the convergence rate.

The experimental results for an infinite rewiring neighbourhood (i.e., $\alpha = 1$ and $r_{\text{RRT}} = 1$) show excellent agreement with the theoretical predictions in Theorem 20 (Figs. 12a–c). The mean solution cost converges linearly towards the optimum and closely matches the lower-bound predicted by Lemma 17.

The experimental results for a constant finite rewiring neighbourhood (i.e., $\alpha = 0.4$ and $r_{\text{RRT}} = 1$) show that the convergence rate is lower than predicted by Theorem 20 (Figs. 12d–f). The convergence rate appears to be initially nonlinear but then become linear. It is hypothesized that this is related to the density of samples relative to the maximum edge length as reflected by α in Theorem 6.

The experimental results for a decreasing finite rewiring neighbourhood (i.e., $\alpha = 1$ and $r_{\text{RRT}} = 1.1r_{\text{RRT}}$) show that the convergence rate appears to be sublinear (Figs. 12g–i). It is hypothesized that this is a result of the rewiring neighbourhood shrinking ‘too’ fast relative to the sample density.

These experiments suggest that further research is necessary to study the tradeoff between per-iteration cost and the

number of iterations needed to find a solution. While a shrinking rewiring neighbourhood limits the number of rewirings, the apparent resulting sublinear convergence would require significantly more iterations to find high-quality solutions. Alternatively, while linear convergence needs fewer iterations to find equivalent solutions, the required constant radius would allow the number of rewirings to increase indefinitely.

VII. EXPERIMENTS

Informed RRT* was evaluated on simulated problems in R^2 , R^4 , and R^8 (Sections VII-A and VII-B) and for HERB (Section VII-C) using OMPL². It was compared to the original RRT* and versions that focus the search with graph pruning (e.g., Alg. 7), heuristic rejection on x_{new} , heuristic rejection on x_{rand} , and all three techniques combined.

All planners used the same tuning parameters and the ordered rewiring technique presented in [38]. Planners used a goal-sampling bias of 5% and an RRT* radius of $r_{\text{RRT}} = 2r_{\text{RRT}}$. The maximum edge length was selected experimentally to reduce the time required to find an initial solution on a training problem, with values of $\alpha = 0.3, 0.5, 0.9$, and 1.3 used in R^2 ,

²The experiments were run on a laptop with 16 GB of RAM and an Intel i7-4810MQ processor. The abstract experiments were run in Ubuntu 12.04 (64-bit) with Boost 1.58, while the HERB experiments were run in Ubuntu 14.04 (64-bit).

R^4 , R^8 , and on HERB (R^{14}), respectively. Available planning time was limited for each state dimension to 3, 30, 150, and 600 seconds, respectively. Planners with heuristics used the L^2 norm as estimates of cost-to-come and cost-to-go while those with graph pruning delayed its application until solution cost changed by more than 5%.

These experiments were designed to investigate admissible methods to focus search. More advanced extensions of RRT* were not considered as they commonly include some combination of the investigated techniques.

A. Toy Problems

Two separate experiments were run in R^2 , R^4 , and R^8 on randomized variants of the toy problem depicted in Fig. 13a to investigate the effects of obstacles on convergence.

The problem consists of a (hyper)cube of width 1 with a single start and goal located at $[-0.5; 0; \dots; 0]^T$ and $[0.5; 0; \dots; 0]^T$, respectively. A single (hyper)cube obstacle of width $w \in U[0.25; 0.5]$ sits between the start and goal in the centre of the problem domain.

The first experiment investigates finding near-optimal solutions in the presence of obstacles. The time required for each planner to find a solution within various fractions of the known optimum, c , was recorded over 100 trials with different pseudo-random seeds for maps of width $l = 2$. The percentage of trials that found a solution within the target tolerance of the optimum and the median time necessary to do so are presented for each planner in Figs. 14a–c. Trials that did not find a suitable solution were treated as having infinite time for the purpose of calculating the median. The results show that Informed RRT* performs equivalently to rejection sampling algorithms in low state dimensions but outperforms all existing techniques in higher dimensions.

The second experiment investigates finding near-optimal solutions in large planning problems. The time required for each planner to find a near-optimal solution was recorded over 100 trials with different pseudo-random seeds for maps of increasing width, l . Planners sought a solution better than $1.01c$, $1.05c$, and $1.15c$ in R^2 , R^4 , and R^8 , respectively. The percentage of trials that found a sufficiently near-optimal solution and the median time necessary to do so are presented for each planner in Figs. 14d–f. Trials that did not find a suitable solution were treated as having infinite time for the purpose of calculating the median. The results show that Informed RRT* outperforms all existing techniques in large-domain planning problems and that the difference increases in higher state dimensions.

These experiments show that increasing problem size and state dimension decreases the ability of nondirect sampling methods to find near-optimal solutions, as predicted by (8). Informed RRT* limits these effects and outperforms existing techniques by efficiently focusing its search to the L^2 informed set using direct informed sampling.

B. Worlds with Many Homotopy Classes

The algorithms were tested on more complicated problems with many homotopy classes in R^2 , R^4 , and R^8 . The worlds consisted of a (hyper)cube of width $l = 4$ with the start and goal

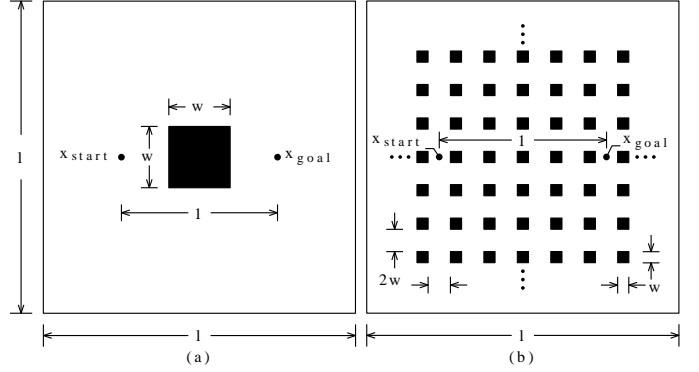


Fig. 13. Illustrations of the planning problems used for Sections VII-A and VII-B to study performance relative to a known optimum, the effect of map width, l , and performance in problems with many homotopy classes. The width of the obstacle in (a) is a random variable uniformly distributed over the range $[0.25; 0.5]$. The regularly spaced obstacles in (b) are chosen in to scale efficiently to high dimensions and their width is such that the start and goal states are 5 ‘columns’ apart.

located at $[-0.5; 0; \dots; 0]^T$ and $[0.5; 0; \dots; 0]^T$, respectively. The problem domain was filled with a regular pattern of axis-aligned (hyper)cube obstacles with a width such that the start and goal were 5 ‘columns’ apart (Fig. 13b).

The planners were tested with 100 different pseudo-random seeds on each world and state dimension. The solution cost of each planner was recorded every 1 millisecond by a separate thread and the median was calculated from the 100 trials by interpolating each trial at a period of 1 millisecond. The absence of a solution was considered an infinite cost for the purpose of calculating the median.

The results are presented in Figs. 14g–i, where the percent of trials solved and the median solution cost are plotted versus run time. They demonstrate how Informed RRT* has better real-time convergence towards the optimum than existing techniques, especially in higher state dimensions.

C. Motion Planning for HERB

Informed RRT* was demonstrated on a high-dimensional problem using HERB, a 14-DOF mobile manipulation platform [11]. Poses were defined for the two arms to create a sequence of three planning problems (Fig. 15) inspired by [39]. The objective of these problems was to find the minimum path length through a 14-dimensional search space with strict limits (each joint has no more than -2 radians of travel). While path length is not a common cost function for manipulation, these experiments illustrate that direct informed sampling is beneficial in high-dimensional problem domains even with strict search limits.

RRT*, RRT* with pruning and rejection, and Informed RRT* were each run for 50 trials on each problem of the cycle. The resulting median path lengths are presented in Fig. 16. Trials that did not find a solution were considered to have infinite length for the purpose of calculating the median. This only occurred for the problem from (a) to (b), where the planners found a solution on 94% of the trials.

RRT* with and without pruning and rejection sampling both fail to improve the initial solutions on all three planning

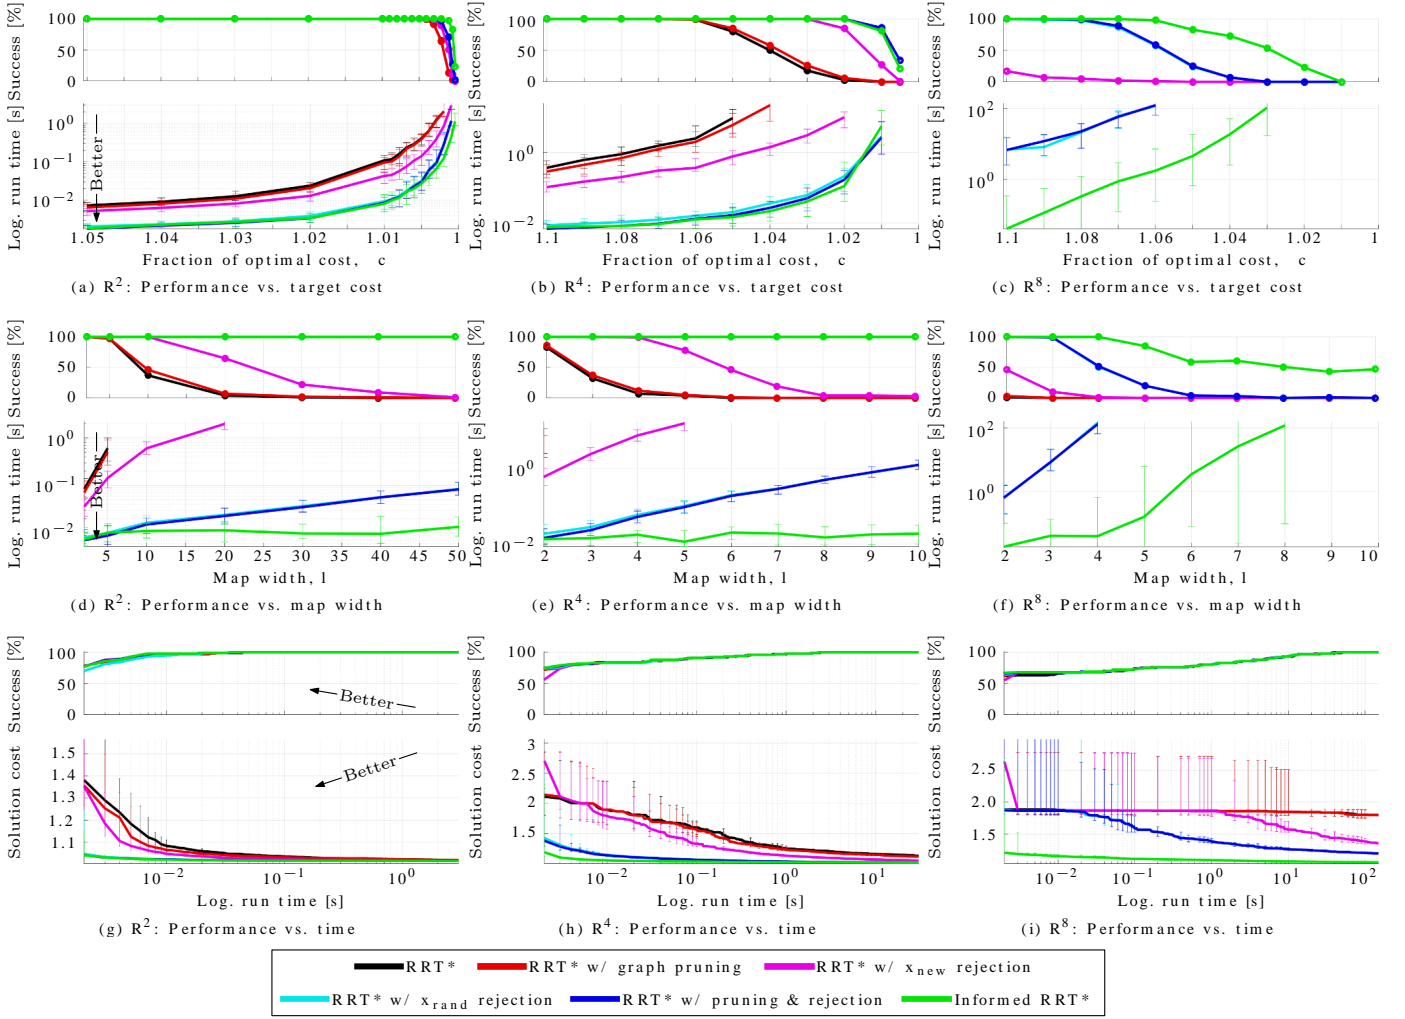


Fig. 14. Results for the experiments described in Sections VII-A and VII-B. Each planner was run 100 different times in R^2 , R^4 , and R^8 on each problem for 3, 30, and 150 seconds, respectively. The percentage of trials that found the desired solution are plotted above and the median performance is plotted below for each experiment. Unsuccessful trials were assigned an infinite value for the purpose of calculating the median and the error bars denote a nonparametric 99% confidence interval on the median value. The times required to find different near-optimal solutions, $c_i < c$, for the problem illustrated in Fig. 13a with $l = 2$ are presented in (a)–(c). The times required to find a solution within a fraction of the known optimum ($1:01c$, $1:05c$, and $1:15c$, respectively) for the problem illustrated in Fig. 13a for various map widths are presented in (d)–(f). Solution cost is plotted versus run time for the problem illustrated in Fig. 13b in (g)–(i). Taken together, these experiments demonstrate the benefits of direct informed sampling even in large or high-dimensional problems, with a high number of obstacles, and many homotopy classes.

problems but Informed RRT* is able to improve the path length by 3:9%, 7:9%, and 28:2%, respectively. The improvement for (a) to (b) is not statistically significant but (b) to (c) and (c) to (d) demonstrate the benefits of considering the relative sizes of the informed set and problem domain in high state dimensions.

VIII. DISCUSSION & CONCLUSION

RRT* almost-surely converges asymptotically to the optimum by asymptotically finding the optimal paths to every state in the problem domain. This is inefficient in single-query scenarios as, once a solution is found, searches only need to consider states that can belong to a better solution (i.e., the omniscient set; Definition 3, Lemma 4). Previous work has focused search to estimates of this set (i.e., informed sets; Definition 7) but has not used these estimates to analyze performance. This paper proves that for holonomic problems the probability of

sampling an admissible informed set provides an upper bound on the probability of improving a solution (Theorem 13).

A popular admissible heuristic for problems seeking to minimize path length is the L^2 norm (i.e., Euclidean distance). This paper shows that existing techniques to exploit it are insufficient. The majority of approaches either reduce the ability to find solutions in other homotopy classes (i.e., reduce recall; Definition 9) or fail to account for the reduction of the L^2 informed set in response to solution improvement (i.e., have decreasing precision; Definition 8). Even existing adaptive techniques that address these problems (e.g., [9]) fail to account for its factorial decrease in measure with state dimension (i.e., the minimum-path-length curse of dimensionality; Theorem 14).

This paper presents a method to avoid these limitations through direct sampling of the L^2 informed set (Algs. 1–5; Section IV). This approach generates uniformly distributed samples in the informed set regardless of its size relative to




Fig. 15. A motion planning problem for HERB inspired by [39]. Planners must find a collision-free path between each pair of subsequent poses, e.g., (a) to (b). HERB’s 14 DOFs and large number of potential self-collisions make this a nontrivial planning problem for RRT*. The planners were given 600 seconds for each phase of the planning problem and the results are presented in Fig 16.

the problem domain or the state dimension (i.e., it has 100% recall and high precision). This paper presents Informed RRT* as a demonstration of how these techniques can be used in sampling-based planning (Algs. 6 and 7; Section V).

Informed RRT* considers all homotopy classes that could provide a better solution (i.e., 100% recall), unlike sample biasing techniques. It is effective regardless of the relative size of the informed set or the state dimension, unlike sample rejection or graph pruning. When the heuristic does not provide any information (e.g., small planning problems and/or large informed sets) it is identical to RRT*.

This paper also uses the shape of the L^2 informed set to analyze the theoretical performance of RRT* on minimum-path-length problems (Section VI) by bounding the expected solution cost (Lemma 17) and convergence rates (Theorems 18–20). The bounds are sharp over the set of all (Lemma 17) or all holonomic (Theorems 18–20) minimum-path-length planning problems and algorithm configurations with the lower bounds exact for an infinite rewiring radius in the absence of obstacles and constraints. These results prove that RRT* converges sublinearly (i.e., slower than linear) for all configurations and holonomic minimum-path-length problems and that focused variants (e.g., Informed RRT*) can have linear convergence.

This analysis is extended experimentally to different configurations. The results confirm the theoretical findings and suggest that obstacle- and constraint-free convergence remains linear when the rewiring radius is constant but becomes sublinear when it decreases in the manner proposed by [6]. As previous analysis of this radius has focused on per-iteration complexity, we believe this result motivates future research into the trade off between per-iteration cost and convergence rate.

The practical advantages of Informed RRT* are shown on a variety of planning problems (Section VII). These experiments demonstrate how its theoretical convergence rate corresponds to better performance on real planning problems. The amount of improvement depends on how efficiently the L^2 informed set decreases the search domain and may be limited in small problem domains and/or long circuitous solutions (e.g., the small/low-dimensional problems in Section VII-A and the first problem of Section VII-C). The design of Alg. 3 assures that in these situations Informed RRT* performs no worse than other methods to exploit the L^2 heuristic (e.g., rejection sampling).

Designing these experiments highlighted the relationship between the maximum edge length, ϵ , and algorithm performance. This user-selected value not only affected the time required to




Fig. 16. Median path length results from the motion planning problems depicted in Fig. 15. Planners found a solution between each pose in every trial after 600 seconds other than the transition from (a) to (b), where solutions were only found in 94% of the 50 trials. For the purpose of calculating the median, these unsolved trials were assigned an infinite cost. Error bars denote a nonparametric 99% confidence interval on the median value. The results show that even in the presence of strict state-space limits, Informed RRT* can outperform rejection sampling in high-dimensional problems.

find an initial solution but, as a result of (19), also the quality of the solution found in finite time. Specifically, large values of ϵ appeared to decrease the difference between algorithms; however, also resulted in order of magnitude increases in the time required to find initial solutions. When coupled with the results of Section VII, this result should further motivate more research into the effects of the RRT* tuning parameters, ϵ and r_{RRT} , on real-time performance. Given that anytime improvement of a solution is a major feature of RRT*, we tuned ϵ for these experiments to minimize the initial-solution time on a series of independent test problems.

We believe that defining precise and admissible informed sets is a fundamental challenge of using anytime almost-surely asymptotically optimal planners in real-world applications. The L^2 informed set is a sharp, uniformly admissible estimate of the omniscient set for problems seeking to minimize path length, even in the presence of constraints, and is exact in the absence of obstacles and constraints. This suggests that any informed set that is more precise must either (i) exploit additional information about the problem domain (e.g., obstacles, constraints), and/or (ii) be inadmissible for some minimum-path-length planning problems. Finding ways to define new admissible heuristics from additional problem-specific information could potentially allow focused search algorithms to converge linearly in the presence of obstacles and/or constraints.

We ultimately believe that heuristics are a key component of successful planning algorithms. To this end, we are currently investigating methods to extend heuristics to entire sampling-based searches, similar to how A* [2] extends Dijkstra’s algorithm [1]. We accomplish this in BIT* [34–36] by extending the ideas presented in this paper to batches of randomly generated samples. These samples are limited to informed sets and searched in order of potential solution quality. Information on OMPL implementations of both Informed RRT* and BIT* are available at <http://asrl.utias.utoronto.ca/code>.

APPENDIX A PROOFS OF LEMMAS 4 AND 5

This section restates and proves Lemmas 4 and 5.

A. Proof of Lemma 4

Lemma 4 (The necessity of adding states in the omniscient set). Adding a state from the omniscient set, $x_{\text{new}} \in X_f$, is a necessary condition for RRT* to improve the current solution, c_i ,

$$c_{i+1} < c_i \Rightarrow x_{\text{new}} \in X_f :$$

This condition is necessary but not sufficient to improve the solution as the ability of states in X_f to provide better solutions at any iteration depends on the structure of the tree (i.e., its optimality).

Proof. At the end of iteration $i + 1$, the cost of the best solution found by RRT* will be the minimum of the previous best solution, c_i , and the best cost of any new or newly improved solutions, c_{new} ,

$$c_{i+1} = \min \{c_i, c_{\text{new}}\} : \quad (28)$$

Each iteration of RRT* only adds connections to or from the newly added state, x_{new} , and therefore all new or modified paths pass through this new state. The cost of any of these new paths that extend to the goal region will be bounded from below by the cost of the optimal solution of a path through x_{new} ,

$$c_{\text{new}} = f(x_{\text{new}}) : \quad (29)$$

Lemma 4 is now proven by contradiction. Assume that RRT* has a solution with cost c_i after iteration i and that it is improved at iteration $i + 1$ by adding a state not in the omniscient set, $c_{i+1} < c_i$; $x_{\text{new}} \notin X_f$. By (1), the costs of solutions through any $x_{\text{new}} \notin X_f$ are bounded from below by the current solution,

$$f(x_{\text{new}}) \geq c_i ;$$

which by (29) is also a bound on the cost of any new or modified solutions,

$$c_{\text{new}} = f(x_{\text{new}}) \geq c_i :$$

By (28), the cost of the best solution found by RRT* at the end of iteration $i + 1$ must therefore be c_i . This contradicts the assumption that the solution was improved by a state not in the omniscient set and proves Lemma 4. \square

B. Proof of Lemma 5

Lemma 5 (The necessity of sampling states in the omniscient set in holonomic planning). Sampling the omniscient set, $x_{\text{rand}} \in X_f$, is a necessary condition for RRT* to improve the current solution to a holonomic problem, c_i , after an initial iterations,

$$8i \quad c_{i+1} < c_i \Rightarrow x_{\text{rand}} \in X_f ;$$

for any sample distribution that maintains a nonzero probability over the entire omniscient set.

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints with appropriate assumptions.

Proof. In RRT (and therefore RRT*), the distribution of vertices in the graph approaches the sample distribution as

the number of iterations approach infinity [?]. In the limit, all reachable regions of the problem domain with a nonzero sampling probability will therefore be sampled and the number of vertices in these regions will increase indefinitely with the number of iterations. This ever increasing number of vertices means that the worst-case distance between any state in a sampled subset and the nearest vertex in the graph will decrease indefinitely and monotonically.

Lemma 5 is now proven by contradiction. Assume that by iteration there are a sufficient number and distribution of vertices in the tree such that all possible states in X_f are no further than from a vertex,

$$\forall x \in X_f ; \exists v \in V \text{ s.t. } \|x - v\|_2 < \epsilon ; \quad (30)$$

and that RRT* has a solution with cost c_i after iteration i . Now assume that RRT* improves the solution at iteration $i + 1$ without sampling the omniscient set, $c_{i+1} < c_i$; $x_{\text{rand}} \notin X_f$.

As improving a solution requires adding a state from the omniscient set, $x_{\text{new}} \in X_f$, (Lemma 4) this implies that the state added to the graph is not the randomly sampled state, $x_{\text{new}} \neq x_{\text{rand}}$. These two states are related in holonomic planning by expansion constraints, (2) and (3), that find a new state as near as possible to x_{rand} and no further than from the nearest vertex in the tree.

The triangle inequality implies that the nearest vertex to the sample, v_{nearest} , is also the nearest vertex to the proposed new state,

$$\begin{aligned} v_{\text{nearest}} &:= \arg \min_{v \in V} \|x_{\text{rand}} - v\|_2 \\ &= \arg \min_{v \in V} \|x_{\text{new}} - v\|_2 ; \end{aligned}$$

which from (30) is bounded in its distance from x_{new} by

$$\|x_{\text{new}} - v_{\text{nearest}}\|_2 < \epsilon ; \quad (31)$$

Due to (3), the relationship in (31) is only satisfied in holonomic planning when $x_{\text{new}} = x_{\text{rand}}$. As by assumption the random sample is not a member of the omniscient set, $x_{\text{rand}} \notin X_f$, then therefore neither is the newly added state, $x_{\text{new}} \notin X_f$, and by Lemma 4 the solution is not improved, $c_{i+1} = c_i$. This contradicts the assumption that the solution was improved by sampling a state not in the omniscient set and proves Lemma 5. \square

APPENDIX B PROOF OF LEMMA 15

This section restates Lemma 15 as presented by [31] along with a full proof as presented in [14].

Lemma 15 (The uniform distribution of samples transformed into a hyperellipsoid from a unit n -ball. Originally Lemma 1 in [31]). If the random points distributed in a hyperellipsoid are generated from the random points uniformly distributed in a hypersphere through a linear invertible nonorthogonal transformation, then the random points distributed in the hyperellipsoid are also uniformly distributed.

Proof. Let the sets $X_{\text{ball}} \subset \mathbb{R}^n$ and $X_{\text{ellipse}} \subset \mathbb{R}^n$ be the unit n -dimensional ball and a n -dimensional hyperellipsoid with radii r_j $g_{j=1}^n$, respectively, having measures of

$$\begin{aligned} \mu(X_{\text{ball}}) &= \frac{\pi^{n/2}}{\Gamma(n/2 + 1)} \\ \mu(X_{\text{ellipse}}) &= \frac{\pi^{n/2}}{\Gamma(n/2 + 1)} \prod_{j=1}^n r_j \end{aligned}$$

Let $p_{\text{ball}}(\cdot)$ be the probability density function of samples drawn uniformly from the unit n -ball such that,

$$p_{\text{ball}}(x) := \begin{cases} \frac{1}{\mu(X_{\text{ball}})} & \text{if } x \in X_{\text{ball}} \\ 0 & \text{otherwise;} \end{cases} \quad (32)$$

Let $T(\cdot)$ be an invertible transformation from the unit n -ball to a hyperellipsoid such that,

$$\begin{aligned} T: X_{\text{ball}} &\rightarrow X_{\text{ellipse}}; \\ T^{-1}: X_{\text{ellipse}} &\rightarrow X_{\text{ball}}; \end{aligned}$$

By definition, the probability density function in the hyperellipsoid, $p_{\text{ellipse}}(\cdot)$, resulting from applying this transformation to samples distributed in the unit n -ball is then

$$p_{\text{ellipse}}(x) := p_{\text{ball}}(T^{-1}(x)) \det \frac{dT^{-1}}{dx} \quad (33)$$

The proposed transformation in (12) has the inverse

$$T^{-1}(x_{\text{ellipse}}) = L^{-1}(x_{\text{ellipse}} - x_{\text{centre}});$$

and the Jacobian

$$\frac{dT^{-1}}{dx_{\text{ellipse}}} = L^{-1}. \quad (34)$$

Substituting (34) and (32) into (33) gives,

$$p_{\text{ellipse}}(x) := \begin{cases} \frac{1}{\mu(X_{\text{ellipse}})} \det(L^{-1}) & \text{if } x \in X_{\text{ellipse}} \\ 0 & \text{otherwise;} \end{cases} \quad (35)$$

using the fact that $T^{-1}(x) \in X_{\text{ball}} \iff x \in X_{\text{ellipse}}$. As $p_{\text{ellipse}}(\cdot)$ is constant for all $x_{\text{ellipse}} \in X_{\text{ellipse}}$, this proves that using (12) to transform uniformly distributed samples in the unit n -ball results in a uniform distribution over the hyperellipsoid and proves Lemma 15.

For hyperellipsoids whose axes are orthogonal (e.g., a prolate hyperspheroid), (35) can be expressed in a more familiar and intuitive form. Using (13) for $T(\cdot)$ and the orthogonality of rotation matrices makes (35)

$$p_{\text{ellipse}}(x) := \begin{cases} \frac{1}{\mu(X_{\text{ellipse}})} \det(L^{-1}) & \text{if } x \in X_{\text{ellipse}} \\ 0 & \text{otherwise;} \end{cases} \quad (36)$$

where $L^0 = \text{diag}(r_1; r_2; \dots; r_n)$ is a diagonal matrix which then simplifies (36) to

$$p_{\text{ellipse}}(x) := \begin{cases} \frac{1}{\mu(X_{\text{ellipse}})} \prod_{j=1}^n \frac{1}{r_j} & \text{if } x \in X_{\text{ellipse}} \\ 0 & \text{otherwise;} \end{cases} \quad (37)$$

since the determinant is a linear operator, all rotation matrices have a unity determinant, $\det(C_{\text{we}}) = 1$, and the determinant of a diagonal matrix is the product of its diagonal entries. As expected, (37) is the inverse of the volume of an n -dimensional hyperellipsoid with radii r_j $g_{j=1}^n$. \square

APPENDIX C PROOF OF LEMMA 17

This section restates and proves Lemma 17, which is used in support of Theorems 18–20. An earlier version of this proof appeared in [13].

Lemma 17 (Expected next-iteration cost of minimum-path-length planning). The expected value of the next solution to a minimum-path-length problem, $E[c_{i+1}]$, is bounded by

$$p_f \frac{nc_i^2 + c_{\min}^2}{(n+1)c_i} + (1-p_f)c_i \leq E[c_{i+1}] \leq c_i; \quad (24 \text{ redux})$$

where c_i is the current solution cost, c_{\min} is the theoretical minimum solution cost, n is the state dimension of the planning problem, and $p_f = P(x_{\text{new}} \in X_f)$ is the probability of adding a state that is a member of the omniscient set (i.e., that can belong to a better solution). While not explicitly shown, the subset, X_f , and the probability of improving the solution, p_f , are generally functions of the current solution cost.

This lower bound is sharp over the set of all possible minimum-path-length planning problems and algorithm configurations and is exact for versions of RRT* with an infinite rewiring radius (i.e., $r = 1$, and $r_{\text{RRT}} = 1$) searching an obstacle-free environment without constraints.

Proof. Proof of the upper bound is trivial. RRT* only accepts new solutions that improve its existing solution, assuring that the cost monotonically decreases,

$$c_{i+1} \leq c_i; \quad (38)$$

Proof of the lower bound comes from finding an exact expression for the expected value of the solution cost found in the absence of obstacles and constraints with an infinite rewiring neighbourhood.

The expected solution cost of RRT* depends on the probability of sampling the omniscient set,

$$\begin{aligned} E[c_{i+1}] &= p_f E[c_{i+1} | x_{\text{new}} \in X_f] \\ &\quad + (1-p_f) E[c_{i+1} | x_{\text{new}} \notin X_f]; \end{aligned} \quad (39)$$

where $p_f = P(x_{\text{new}} \in X_f)$. Adding a state from the omniscient set, X_f , is a necessary condition to improve the solution (Lemma 4) and any other state will not change the solution cost, $E[c_{i+1} | x_{\text{new}} \notin X_f] = c_i$. This simplifies (39) to

$$E[c_{i+1}] = p_f E[c_{i+1} | x_{\text{new}} \in X_f] + (1-p_f)c_i; \quad (40)$$

The costs of solutions found by adding states inside the omniscient are bounded from below by the optimal path through the newly added state,

$$E[c_{i+1} | x_{\text{new}} \in X_f] \geq E[f(x_{\text{new}}) | x_{\text{new}} \in X_f]; \quad (41)$$

where $f(x)$ is the cost of the optimal path from the start to the goal constrained to pass through a state, x . With a uniform sample distribution over X_f the right-hand side of (41) becomes

$$E[f(x_{\text{new}}) | x_{\text{new}} \in X_f] = \frac{1}{\mu(X_f)} \int_{X_f} f(x_{\text{new}}) dV;$$

When RRT* uses an infinite rewiring radius it attempts connections between every new state and the start and goal.

In the absence of obstacles and constraints these paths will be feasible and represent the optimal solutions using the state. This makes the expected value of this best-case configuration of RRT* equivalent to the expected optimal solution cost in the absence of obstacles,

$$E[c_{i+1} j x_{\text{new}} 2 X_f] = E[f(x_{\text{new}}) j x_{\text{new}} 2 X_f]: \quad (42)$$

The lower bound provided by (41) is therefore sharp over the set of all possible planning problems and algorithm configurations.

In this absence of obstacles and constraints, the optimal solution using any state is given by (4) and the omniscient set is the prolate hyperspheroid, $X_f = X_{\text{p}} = X_{\text{PHS}}$. The measure of the omniscient set, $(X_f) = \text{PHS}$, is given by (5). This allows (42) to be written as

$$E[c_{i+1} j x_{\text{new}} 2 X_f] = \frac{1}{\text{PHS}} \int_{X_{\text{PHS}}} (kx - x_{\text{start}} k_2 + kx_{\text{goal}} - x k_2) dV: \quad (43)$$

The prolate hyperspheroidal coordinates, $;$ $1;::; n n 2$,

$$x_1 = a \cosh \cos ;$$

$$x_2 = a \sinh \sin \cos 1;$$

$$x_3 = a \sinh \sin \sin 1 \cos 2;$$

$$\vdots$$

$$x_{n-1} = a \sinh \sin \sin 1 \sin 2 :: \sin n n 3 \cos n n 2;$$

$$x_n = a \sinh \sin \sin 1 \sin 2 :: \sin n n 3 \sin n n 2;$$

and the parameterization $a = 0.5c_{\text{min}}$, simplifies (4) to

$$f(x) = c_{\text{min}} \cosh \quad (44)$$

Substituting (44) and the prolate hyperspheroidal differential volume,

$$dV = a^n \sinh^2 + \sin^2 \sinh^{n-2} \sin^{n-2} \sin^{n-3} 1 :: \sin \psi_1 3 d d d 1 :: d \psi_n 2;$$

into (43) results in

$$E[c_{i+1} j x_{\text{new}} 2 X_f] = \frac{c_{\text{min}}^{n+1}}{2^n \text{PHS}} \int_0^1 \int_0^1 \int_0^1 \int_0^1 \dots \int_0^1 (\sinh^2 + \sin^2 \sinh^{n-2} \cosh \sin^{n-2} \sin^{n-3} 1 :: \sin n -3 d d d 1 :: d n -2; \quad (45)$$

where the integration limit for ψ_1 is derived from (44) as

$$\cosh \psi_1 := \frac{c_i}{c_{\text{min}}}: \quad (46)$$

Integrating (45) requires applying a series of identities, first

$$(n-1) \int_0^1 \int_0^1 \int_0^1 \dots \int_0^1 \sinh^{n-3} 1 :: \sin n -3 d d 1 :: d n -2;$$

simplifies (45) to

$$E[c_{i+1} j x_{\text{new}} 2 X_f] = \frac{(n-1) c_{\text{min}}^{n+1} \int_0^1 \int_0^1 \dots \int_0^1 (\sinh^2 + \sin^2 \sinh^{n-2} \cosh \sin^{n-2} d d :: \quad (47)$$

Next, the definite integral of the product of powers of sin and cos,

$$\int_0^1 \int_0^1 \dots \int_0^1 \sin^{2m-1} \cos^{2n-1} d d :: B(m;n);$$

where $B(;)$ is the beta function,

$$B(m;n) := \int_0^1 t^{m-1} (1-t)^{n-1} dt;$$

is used to evaluate the integral over ψ_1 in (47), giving

$$E[c_{i+1} j x_{\text{new}} 2 X_f] = \frac{(n-1) c_{\text{min}}^{n+1} \int_0^1 \int_0^1 \dots \int_0^1 \sin^{n-1} \cosh d + B\left(\frac{n+1}{2}; \frac{1}{2}\right) \int_0^1 \int_0^1 \dots \int_0^1 \sinh^{n-2} \cosh d :: \quad (48)$$

The identity,

$$B(m+1;n) = \frac{m}{m+n} B(m;n);$$

and the recursive nature of the n -dimensional unit ball,

$$\int_0^1 \int_0^1 \dots \int_0^1 B\left(\frac{n+1}{2}; \frac{1}{2}\right) \int_0^1 \dots \int_0^1 ::$$

simplifies (48) to

$$E[c_{i+1} j x_{\text{new}} 2 X_f] = \frac{c_{\text{min}}^{n+1} \int_0^1 \int_0^1 \dots \int_0^1 \sinh^n \cosh d + (n-1) \int_0^1 \int_0^1 \dots \int_0^1 \sinh^{n-2} \cosh d :: \quad (49)$$

The indefinite integral,

$$\int \sinh^m \cosh d = \frac{\sinh^{m+1}}{m+1};$$

is then used to evaluate (49), giving

$$E[c_{i+1} j x_{\text{new}} 2 X_f] = \frac{c_{\text{min}}^{n+1} \int_0^1 \int_0^1 \dots \int_0^1 \sinh^{n+1} \psi_1 + \sinh^{n-1} \psi_1 :: \quad (50)$$

Using (5) to expand the measure PHS in (50) cancels the measure of the unit n -ball, giving

$$E[c_{i+1} j x_{\text{new}} 2 X_f] = \frac{c_{\text{min}}^{n+1}}{c_i (c_i^2 - c_{\text{min}}^2)^{\frac{n-1}{2}}} \int_0^1 \int_0^1 \dots \int_0^1 \sinh^{n+1} \psi_1 + \sinh^{n-1} \psi_1 :: \quad (51)$$

Using the relationship

$$\cosh \psi_1 = b() \sinh \psi_1 = \frac{p}{b^2 - 1};$$

some algebraic manipulation, and (46) finally simplifies (51) to

$$E [c_{i+1} | x_{\text{new}} \in X_f] = \frac{nc_i^2 + c_{\min}^2}{(n+1)c_i};$$

an exact value for the best-case expected solution cost of RRT*. This result allows (40) to be written as the sharp bound,

$$E [c_{i+1}] \leq p_f \frac{nc_i^2 + c_{\min}^2}{(n+1)c_i} + (1 - p_f) c_i;$$

which when combined with (38) proves Lemma 17. \square

APPENDIX D

PROOFS OF THEOREMS 18–20

This section restates and proves Theorems 18–20.

A. Proof of Theorem 18

Theorem 18 (Sublinear convergence of RRT* in holonomic minimum-path-length planning). RRT* converges sublinearly towards the optimum of holonomic minimum-path-length planning problems,

$$E [RRT] = 1; \quad (25 \text{ redux})$$

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. The expected rate of convergence (Definition 16) of RRT* is

$$E [RRT] = E \lim_{i \rightarrow \infty} \frac{c_i - c}{c_{i-1} - c}; \quad (52)$$

since $8i; c_i \rightarrow c$. As RRT* almost-surely converges asymptotically to the optimum, this sequence also almost-surely converges to a finite value, $0 \leq RRT \leq 1$,

$$P \lim_{i \rightarrow \infty} \frac{c_i - c}{c_{i-1} - c} = RRT = 1:$$

By Lebesgue's dominated convergence theorem this allows the expectation operator to be brought inside the limit of (52), giving

$$E [RRT] = \lim_{i \rightarrow \infty} \frac{E [c_i] - c}{c_{i-1} - c}; \quad (53)$$

since c_i is the only random variable at iteration i .

Lemma 17 provides sharp bounds for the expected solution cost at any iteration, $E [c_i]$, with the lower-bound corresponding to an infinite rewiring radius in the absence of obstacles and constraints. Substituting this lower bound and that $c = c_{\min}$ in the absence of obstacles into (53) and simplifying gives an expression for the expected best-case convergence rate,

$$E [RRT] = 1 + \frac{1}{(n+1)} \lim_{i \rightarrow \infty} \frac{p_f (c_{\min}^2 - c_{i-1}^2)}{(c_{i-1}^2 - c_{\min} c_{i-1})};$$

such that $E [RRT] = E [RRT]$ is a sharp bound over all possible planning problems and algorithm configurations. Applying l'Hôpital's rule [?] with respect to c_{i-1} gives

$$E [RRT] = 1 + \frac{1}{(n+1)} \lim_{i \rightarrow \infty} \frac{\frac{\partial p_f}{\partial c_{i-1}} (c_{\min}^2 - c_{i-1}^2) - 2p_f c_{i-1}}{(2c_{i-1} - c_{\min})}; \quad (54)$$

As iterations go to infinity the probability of adding a sample in X_f becomes the probability of sampling it (Lemma 5). The lower bound from Lemma 17 is for an obstacle- and constraint-free problem and therefore the informed set is the omniscient set, $X_f = X_b$, and the probability of sampling it is given by (7) with a partial derivative of

$$\frac{\partial p_f}{\partial c_{i-1}} = \frac{\frac{n}{2}}{2^n \Gamma(\frac{n}{2} + 1)} (X_{\text{samp}})^{\frac{n}{2}} (c_{i-1}^2 - c_{\min}^2)^{\frac{n-1}{2}} \left(1 + \frac{(n-1)c_{i-1}^2}{(c_{i-1}^2 - c_{\min}^2)} \right);$$

Almost-sure convergence to c_{\min} implies $\lim_{i \rightarrow \infty} \frac{1}{i!} c_{i-1} = c_{\min}$ and therefore $\lim_{i \rightarrow \infty} \frac{1}{i!} p_f = 0$ and $\lim_{i \rightarrow \infty} \frac{\partial p_f}{\partial c_{i-1}} = 0$, making (54),

$$E [RRT] = 1;$$

As by definition the expected rate of convergence of RRT* is bounded by,

$$E [RRT] \leq E [RRT] = 1;$$

this result proves Theorem 18. \square

B. Proof of Theorem 19

Theorem 19 (Linear convergence of RRT* with adaptive rectangular rejection sampling in holonomic minimum-path-length planning). RRT* with adaptive rectangular rejection sampling converges at best linearly towards the optimum of holonomic minimum-path-length planning problems but factorially approaches sublinear convergence with increasing state dimension,

$$1 - \frac{\frac{n}{2}}{(n+1)2^{n-1}\Gamma(\frac{n}{2} + 1)} \leq E [Rect] \leq 1; \quad (26 \text{ redux})$$

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. Proof of Theorem 19 follows that of Theorem 18 but with the probability of adding a new state from X_f instead calculated from (7) using (10), as

$$p_f = \frac{\frac{n}{2}}{2^n \Gamma(\frac{n}{2} + 1)}; \quad (55)$$

As $\frac{\partial p_f}{\partial c_{i-1}} = 0$, (54) becomes

$$E [Rect] = 1 - \frac{1}{(n+1)} \lim_{i \rightarrow \infty} \frac{2p_f c_{i-1}}{(2c_{i-1} - c_{\min})}; \quad (56)$$

Noting that almost-sure convergence to c_{\min} implies $\lim_{i \rightarrow \infty} \frac{1}{i!} c_{i-1} = c_{\min}$ and substituting (55) into (56) results in

$$E [Rect] = 1 - \frac{2p_f}{(n+1)} \left(1 - \frac{\frac{n}{2}}{(n+1)2^{n-1}\Gamma(\frac{n}{2} + 1)} \right);$$

As by definition the expected rate of convergence of RRT* with rectangular rejection sampling is bounded by,

$$E[Rect] = E[Rect] - 1;$$

This result proves Theorem 19 with sharp bounds over all possible holonomic planning problems and algorithm configurations. \square

C. Proof of Theorem 20

Theorem 20 (Linear convergence of Informed RRT* in holonomic minimum-path-length planning). Informed RRT* converges at best linearly towards the optimum of holonomic minimum-path-length planning problems,

$$\frac{n-1}{n+1} \leq E[Inf] \leq 1; \quad (27 \text{ redux})$$

where the lower-bound occurs exactly with an infinite rewiring neighbourhood in the absence of obstacles and constraints.

For simplicity, this statement is limited to holonomic planning but it can be extended to specific constraints by expanding Lemma 5.

Proof. Proof of Theorem 20 follows that of Theorem 18 but with a unity probability of adding a new state from X_f . From (54), the convergence rate of Informed RRT* is then,

$$E[Inf] = 1 - \frac{1}{(n+1)!} \lim_{i \rightarrow \infty} \frac{2c_{i-1}}{(2c_{i-1} - c_{\min})};$$

As almost-sure convergence to c_{\min} implies $\lim_{i \rightarrow \infty} c_{i-1} = c_{\min}$, this gives,

$$E[Inf] = \frac{n-1}{n+1};$$

As by definition the expected rate of convergence of RRT* with rectangular rejection sampling is bounded by,

$$E[Inf] = E[Inf] - 1;$$

This result proves Theorem 20 with sharp bounds over all possible holonomic planning problems and algorithm configurations. \square

ACKNOWLEDGMENT

The authors would like to thank the editorial board and reviewers for their helpful comments. We would also like to thank Laszlo-Peter Berczi for discussions on precision and recall, Christopher Dellin and Michael Koval for discussions and insight on the sampling of arbitrarily overlapping shapes, Jennifer King and Clinton Liddick for help running the HERB experiments, and Paul Newman for providing the time to finish this manuscript at the University of Oxford.

REFERENCES

- [1] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
- [2] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, Jul. 1968.
- [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, R. Kindel, J.-C. Latombe, and S. Rock, "Randomized kinodynamic motion planning with moving obstacles," *The International Journal of Robotics Research*, vol. 21, no. 3, pp. 233–255, 2002.
- [5] 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.
- [6] 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.
- [7] D. Ferguson and A. Stentz, "Anytime RRTs," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2006, pp. 5369–5375.
- [8] B. Akgun and M. Stilman, "Sampling heuristics for optimal motion planning in high dimensions," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011, pp. 2640–2645.
- [9] M. Otte and N. Correll, "C-FOREST: Parallel shortest path planning with superlinear speedup," *IEEE Transactions on Robotics (T-RO)*, vol. 29, no. 3, pp. 798–806, Jun. 2013.
- [10] 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, Dec. 2012, library available from <http://ompl.kavrakilab.org>.
- [11] S. Srinivasa, D. Berenson, M. Cakmak, A. Collet Romea, M. Dogar, A. Dragan, R. A. Knepper, T. D. Niemueller, K. Strabala, J. M. Vandeweghe, and J. Ziegler, "HERB 2.0: Lessons learned from developing a mobile manipulator for the home," *Proceedings of the IEEE*, vol. 100, no. 8, pp. 1–19, Jul. 2012.
- [12] 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 *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Chicago, Illinois, USA, 14–18 Sep. 2014, pp. 2997–3004.
- [13] —, "On recursive random prolate hyperspheroids," *Autonomous Space Robotics Lab*, University of Toronto, Tech. Rep. TR-2014-JDG002, Mar. 2014, arXiv:1403.7664 [math.ST].
- [14] J. D. Gammell and T. D. Barfoot, "The probability density function of a transformation-based hyperellipsoid sampling technique," *Autonomous Space Robotics Lab*, University of Toronto, Tech. Rep. TR-2014-JDG004, Apr. 2014, arXiv:1404.1347 [math.ST].
- [15] J. D. Gammell, T. D. Barfoot, and S. S. Srinivasa, "Informed sampling for asymptotically optimal path planning (consolidated version)," *Autonomous Space Robotics Lab*, University of Toronto, Tech. Rep., Jun. 2017, arXiv:1706.06454 [cs.RO].
- [16] 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 (IJRR)*, vol. 34, no. 7, pp. 883–921, 2015.
- [17] J. Nasir, F. Islam, U. Malik, Y. Ayaz, O. Hasan, M. Khan, and M. S. Muhammad, “RRT*-SMART: A rapid convergence implementation of RRT*,” *International Journal of Advanced Robotic Systems*, vol. 10, p. 299, 2013.
- [18] S. Kiesel, E. Burns, and W. Ruml, “Abstraction-guided sampling for motion planning,” in *Proceedings of the Fifth Annual Symposium on Combinatorial Search (SoCS)*, 2012.
- [19] D. Kim, J. Lee, and S.-E. Yoon, “Cloud RRT*: Sampling cloud based RRT*,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 2519–2526.
- [20] H. Choset and J. Burdick, “Sensor based planning, part I: The generalized voronoi graph,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, vol. 2, May 1995, pp. 1649–1655.
- [21] S. Karaman, M. R. Walter, A. Perez, E. Frazzoli, and S. Teller, “Anytime motion planning using the RRT*,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 1478–1483.
- [22] O. Arslan and P. Tsiotras, “Use of relaxation methods in sampling-based algorithms for optimal motion planning,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2013.
- [23] —, “Dynamic programming guided exploration for sampling-based motion planning algorithms,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2015, pp. 4819–4826.
- [24] C. Urmson and R. Simmons, “Approaches for heuristically biasing RRT growth,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 2, 2003, pp. 1178–1183.
- [25] R. Alterovitz, S. Patil, and A. Derbakova, “Rapidly-exploring roadmaps: Weighing exploration vs. refinement in optimal motion planning,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2011, pp. 3706–3712.
- [26] Y. X. Shan, B. J. Li, J. Zhou, and Y. Zhang, “An approach to speed up RRT*,” in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2014, pp. 594–598.
- [27] O. Salzman and D. Halperin, “Asymptotically near-optimal RRT for fast, high-quality, motion planning,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 4680–4685.
- [28] D. Devaurs, T. Siméon, and J. Cortés, “Efficient sampling-based approaches to optimal path planning in complex cost spaces,” in *Algorithmic Foundations of Robotics XI*, ser. Springer Tracts in Advanced Robotics, H. L. Akin, N. M. Amato, V. Isler, and A. F. van der Stappen, Eds. Springer International Publishing, 2015, vol. 107, pp. 143–159.
- [29] M. Otte and E. Frazzoli, “RRT^X: Asymptotically optimal single-query sampling-based motion planning with quick replanning,” *The International Journal of Robotics Research (IJRR)*, vol. 35, no. 7, pp. 797–822, 2016.
- [30] L. Euler, “De progressionibus transcendentibus seu quarum termini generales algebrae dari nequeunt,” *Commentarii academiae scientiarum Petropolitanae*, vol. 5, pp. 36–57, 1738.
- [31] H. Sun and M. Farooq, “Note on the generation of random points uniformly distributed in hyper-ellipsoids,” in *Proceedings of the Fifth International Conference on Information Fusion*, vol. 1, 2002, pp. 489–496.
- [32] G. Wahba, “A least squares estimate of satellite attitude,” *Society for Industrial and Applied Mathematics (SIAM) Review*, vol. 7, p. 409, 1965.
- [33] A. H. J. de Ruiter and J. R. Forbes, “On the solution of Wahba’s problem on SO(n),” *The Journal of the Astronautical Sciences*, vol. 60, no. 1, pp. 1–31, 2013.
- [34] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, “Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 3067–3074.
- [35] J. D. Gammell, “Informed anytime search for continuous planning problems,” Ph.D. dissertation, University of Toronto, Feb. 2017.
- [36] J. D. Gammell, T. D. Barfoot, and S. S. Srinivasa, “Informed asymptotically optimal anytime search,” *The International Journal of Robotics Research (IJRR)*, 2018, manuscript #IJR-17-2980, in revision, arXiv: 1707.01888 [cs.RO].
- [37] A. Dobson, G. V. Moustakides, and K. E. Bekris, “Geometric probability results for bounding path quality in sampling-based roadmaps after finite computation,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4180–4186.
- [38] A. Perez, S. Karaman, A. Shkolnik, E. Frazzoli, S. Teller, and M. R. Walter, “Asymptotically-optimal path planning for manipulation using incremental sampling-based algorithms,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2011, pp. 4307–4313.
- [39] J. Morali and V. Willis, “Y.M.C.A.” *Cruisin’*, Nov. 13 1978, Casablanca Records NBLP 7118.
- [40] J. D. Gammell, T. D. Barfoot, and S. S. Srinivasa, “Informed sampling for asymptotically optimal path planning,” *IEEE Transactions on Robotics (T-RO)*, 2018.



Jonathan D. Gammell is a Departmental Lecturer in Robotics with the Oxford Robotics Institute (ORI) at the University of Oxford, Oxford, United Kingdom. He performed this work at the Autonomous Space Robotics Lab at the University of Toronto, Toronto, Canada during his Ph. D. degree. He is interested in developing conceptually sound approaches to the fundamental problems of real-world autonomous robotics.



Timothy D. Barfoot is a Professor at the University of Toronto Institute for Aerospace Studies (UTIAS). He holds the Canada Research Chair (Tier II) in Autonomous Space Robotics and works in the area of guidance, navigation, and control of mobile robots in a variety of applications. He is interested in developing methods to allow mobile robots to operate over long periods of time in large-scale, unstructured, three-dimensional environments, using rich onboard sensing (e.g., cameras and laser rangefinders) and computation.

Dr. Barfoot took up his position at UTIAS in May 2007, after spending four years at MDA Space Missions, where he developed autonomous vehicle navigation technologies for both planetary rovers and terrestrial applications such as underground mining. He sits on the editorial boards of the International Journal of Robotics Research and the Journal of Field Robotics. He recently served as the General Chair of Field and Service Robotics (FSR) 2015, which was held in Toronto.



Siddhartha S. Srinivasa is the Boeing Endowed Professor at the School of Computer Science and Engineering, University of Washington. He works on robotic manipulation, with the goal of enabling robots to perform complex manipulation tasks under uncertainty and clutter, with and around people. To this end, he founded the Personal Robotics Lab in 2005. Dr. Srinivasa is also passionate about building end-to-end systems (HERB, ADA, HRP3, CHIMP, Andy, among others) that integrate perception, planning, and control in the real world. Understanding

the interplay between system components has helped produce state of the art algorithms for robotic manipulation, motion planning, object recognition and pose estimation (MOPED), and dense 3D modelling (CHISEL, now used by Google Project Tango), and mathematical models for Human-Robot Collaboration.