

Relevant Region Exploration On General Cost-maps For Sampling-Based Motion Planning

Sagar Suhas Joshi¹ and Panagiotis Tsiotras²

Abstract—Asymptotically optimal sampling-based planners require an intelligent exploration strategy to accelerate convergence. After an initial solution is found, a necessary condition for improvement is to generate new samples in the so-called “Informed Set”. However, Informed Sampling can be ineffective in focusing search if the chosen heuristic fails to provide a good estimate of the solution cost. This work proposes an algorithm to sample the “Relevant Region” instead, which is a subset of the Informed Set. The Relevant Region utilizes cost-to-come information from the planner’s tree structure, reduces dependence on the heuristic, and further focuses the search. Benchmarking tests in uniform and general cost-space settings demonstrate the efficacy of Relevant Region sampling.

I. INTRODUCTION

In recent years, sampling-based motion planning (SBMP) algorithms have gained popularity due to their ability to handle high dimensional search spaces. Deterministic search methods such as A* do not scale well owing to the computational cost associated with the a priori discretization of the search space. Incremental SBMP algorithms such as RRT [1], on the other hand, avoid this computational overhead by generating random samples to build a connectivity tree online. However, the RRT algorithm only guarantees probabilistic completeness. Complementing the exploration module of RRTs with an exploitation module results in asymptotic optimality for these randomized methods i.e., planner converges to the optimal solution almost-surely as the number of samples tend to infinity. While the exploration module generates new samples and extends the connectivity tree, the exploitation module processes this tree to improve the current solution. The popular RRT* algorithm [2] locally rewires the tree for exploitation, while RRT[#] [3] utilizes dynamic programming to implement a “global rewiring” procedure. RRT[#] ensures optimal connection for each vertex in the current graph at the end of every iteration. Other algorithms such as BIT* [4] and FMT* [5] also utilize heuristics and dynamic programming to conduct an efficient search and ensure faster convergence compared to RRT*. The DRRT algorithm [6] uses gradient descent to optimize the location of samples in the tree structure.

Conventionally, SBMP algorithms have employed a uniform random exploration strategy. This results in an implicit Voronoi bias, leading to a rapid exploration of the search space. However, in order to achieve a more focused search, several improvements to the uniform sampling strategy have been suggested. These include local biasing [7], use of a heuristic-based quality measure [8], application of an information-theoretic framework [9], [10] and translating ideas from A* to the continuous domain [11]. OBRRT [12] uses obstacle information to guide the tree growth. MARRT




Fig. 1: Planning in a multiple obstacle environment with Relevant Region sampling (top) and Informed Sampling (bottom). Note that the Relevant Region focuses on two pertinent homotopy classes whereas the Informed Sampling generates uniform samples inside the ellipsoidal region.

[13] retracts samples onto the medial axis of the free space to obtain high clearance paths.

After an initial (sub-optimal) solution is found, exploration can be focused on a subset of the search-space that can potentially further improve the current solution. Doing so can lead to a dramatically faster convergence [14]. The Relevant Region, introduced in [15], leverages the current solution cost and the planner’s tree structure information to focus search. A selective vertex inclusion procedure [15] and a machine learning approach [16] has been proposed to generate new samples in the Relevant Region. However, these approaches fall into the category of rejection sampling methods, which do not scale well for high dimensional problems. The current work rigorously defines the Relevant Region set, analyzes its theoretical properties and presents a *generative* method to sample it. This work also extends the Relevant Region framework for planning on general cost-maps.

Gammell et al. [14] use an admissible heuristic and the current solution cost to define the Informed Set. This set includes all points that can potentially improve the current solution. Exploration outside the Informed Set is thus redun-

dant. For minimum-length planning problems in Euclidean spaces, the authors of [14] proposed an efficient method to generate samples in the L_2 -Informed Set. The Lebesgue measure of the L_2 -Informed set decreases as the solution improves, leading to a focused search. However, Informed Sampling effectively resorts to uniform random sampling if the Lebesgue measure of the Informed Set is comparable to that of the entire search space. This can happen if the heuristic estimate of the solution cost fails to provide a good enough approximation of the true solution cost. The proposed approach addresses these issues by utilizing cost-to-come information from the planner’s tree structure and reducing dependence on heuristics.

The Expansive space trees (EST) algorithm [17] proceeds by selecting a vertex (with probability inversely proportional to number of vertices in its neighborhood) and generates a new sample in its vicinity. Guided ESTs [18] add the A* cost and an exploration term to the vertex weights. The SBA* algorithm [19] incorporates a graph density and a constriction measure into the vertex weight. While the algorithm proposed in this work falls into the category of EST-like methods, a crucial difference is that it only expands vertices and generates new samples in the Relevant Region. Thus, in contrast to EST and its variants, the proposed algorithm avoids needless exploration.

The SBMP algorithms and the exploration strategies mentioned above are traditionally geared towards finding the (length) optimal path in uniform cost spaces. However, many applications require planning algorithms to find the optimal path with respect to a provided cost function. These include the problem of navigation on a rough terrain for a mobile robot (see Fig.2), safety critical path planning with clearance cost-map (example in Fig.3), human aware motion planning [20], and planning on energy landscapes [21]. The Transition-based RRT (T-RRT) algorithm [22] takes a user-defined cost function as an additional input and adds a transition test based on the Metropolis criterion to accept or reject potential new states. The transition test favors exploration of low-cost regions of space and leads to better quality paths. An enhanced, bi-directional version of T-RRT is presented in [23]. Berenson et al [24] combine gradient information within the T-RRT framework to address the issue of navigating cost-space chasms. Finally, Devaurs et al [25] combine the filtering properties of the transition test with the local rewiring procedure of RRT* to obtain the asymptotically optimal T-RRT* algorithm. While the transition test promotes exploration of low-cost regions, unlike the Informed and Relevant Region sets, it does not focus exploration on to a subset of the search-space based on the current solution. Secondly, the probabilistic rejection strategy of the transition test might not scale well to higher dimensional spaces, as the probability of generating a “good” sample that can pass the transition test may decrease rapidly. The proposed algorithm addresses these issues by employing a generative sampling approach. It utilizes heuristics, the current solution cost and the cost function information to effectively focus the search in general cost-space settings.

In the following sections, the path planning problem




Fig. 2: Planning on a terrain cost-map with the proposed sampling strategy. Here, white regions represent rough (high cost) areas and the blacks signify smooth sections.

on general cost-maps is formally defined, followed by a comparison between the Informed and the Relevant Region sets. A technique to generate samples in the Relevant Region is then proposed, followed by benchmarking results.

II. PROBLEM DEFINITION

A. Path Planning Problem

Consider the search space $\mathcal{X}$, which is assumed to be a subset of $\mathbb{R}^d$, where d is a positive integer such that $d \geq 2$. Let $\mathcal{X}_{\text{obs}}$ denote the obstacle space and $\mathcal{X}_{\text{free}} = \text{cl}(\mathcal{X} \setminus \mathcal{X}_{\text{obs}})$ denote the free space. Here, $\text{cl}(A)$ represents closure of the set $A \subset \mathbb{R}^d$. Let $\mathcal{M}(A)$ denote the Lebesgue measure of the set $A \subset \mathbb{R}^d$. Let $\mathbf{x}_s \in \mathcal{X}_{\text{free}}$ denote the initial state and $\mathcal{X}_{\text{goal}} \subset \mathcal{X}_{\text{free}}$ represent the goal region. Let $C : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ denote a continuous state cost function that imposes a “cost-map” over the search space. The path-cost from $\mathbf{x}_1 \in \mathcal{X}$ to $\mathbf{x}_2 \in \mathcal{X}$ along a path $\pi : [0, 1] \rightarrow \mathcal{X}$ with $\pi(0) = \mathbf{x}_1$, $\pi(1) = \mathbf{x}_2$ is given by

$$d_\pi(\mathbf{x}_1, \mathbf{x}_2) = \int_0^1 C(\pi(s)) \left\| \frac{d\pi(s)}{ds} \right\|_2 ds. \quad (1)$$

Equation (1) represents the integral of cost (IC) along a path as a measure of the path quality. Please see [22] for a discussion on different cost criteria. If the path is a straight line, i.e., $\pi(s) = \mathbf{x}_1 + (\mathbf{x}_2 - \mathbf{x}_1)s$, $s \in [0, 1]$, then the IC cost is given by

$$d_\ell(\mathbf{x}_1, \mathbf{x}_2) = \|\mathbf{x}_2 - \mathbf{x}_1\|_2 \int_0^1 C(\mathbf{x}_1 + (\mathbf{x}_2 - \mathbf{x}_1)s) ds. \quad (2)$$

Note that if $C(\mathbf{x}) = 1$ for all $\mathbf{x} \in \mathcal{X}$, then the IC cost in (2) reduces to the familiar Euclidean distance $\|\mathbf{x}_2 - \mathbf{x}_1\|_2$. Let Π denote the set of paths from $\mathbf{x}_s$ to $\mathcal{X}_{\text{goal}}$. Thus, given $(\mathcal{X}, \mathcal{X}_{\text{obs}}, \mathbf{x}_s, \mathcal{X}_{\text{goal}}, C)$, the optimal path planning problem is one of finding the minimum-cost, feasible path $\pi^* \in \Pi$.

$$\begin{aligned} & \arg \min_{\pi \in \Pi} d_\pi(\mathbf{x}_s, \mathbf{x}_g) \\ & \text{subject to: } \pi(0) = \mathbf{x}_s, \pi(1) = \mathbf{x}_g \in \mathcal{X}_{\text{goal}}, \\ & \pi(s) \in \mathcal{X}_{\text{free}}, \quad s \in [0, 1]. \end{aligned} \quad (3)$$

Consider the graph $\mathcal{G} = (V, E)$ with a finite set of vertices $V \subset \mathcal{X}_{\text{free}}$ and edges $E \subseteq V \times V$. A spanning tree $\mathcal{T} = (V_s, E_s)$ is embedded in $\mathcal{G}$ such that $V_s = V$ and

$E_s = \{(\mathbf{u}, \mathbf{v}) \in E \mid \mathbf{v} = \text{parent}(\mathbf{u})\}$. Here the function $\text{parent} : V \rightarrow V$ represents the mapping from a vertex to its unique parent vertex. By definition, $\text{parent}(\mathbf{x}_s) = \mathbf{x}_s$. SBMP algorithms numerically integrate (2) to calculate the edge-cost $d_\ell(\mathbf{v}, \mathbf{u})$ for any edge $(\mathbf{u}, \mathbf{v}) \in E$. This work considers cost functions with $C(\mathbf{x}) \geq 1$ for all $\mathbf{x} \in \mathcal{X}$, so that the edge-cost between any two vertices is at least the Euclidean distance between these two points. Given the spanning tree $\mathcal{T}$ in $\mathcal{G}$, the function $g_{\mathcal{T}} : V \rightarrow \mathbb{R}_{\geq 0}$ provides the cost-to-come value for any $\mathbf{v} \in V$, i.e., it is the sum of the edge-costs along the path from $\mathbf{v}$ to $\mathbf{x}_s$ in $\mathcal{T}$. A consistent heuristic function on $\mathcal{X}$ (such as the Euclidean distance or L_2 -norm) is defined as: $h : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$. The function h always gives an under-estimate of the path-cost between any two points in the search space, and obeys the triangle inequality. The SBMP algorithms solve the planning problem (3) by drawing random samples from $\mathcal{X}$ and by incorporating the collision-free ones in $\mathcal{G}$. An efficient sampling strategy must generate samples so as to find an initial solution or improve the current one. The exploration problem in SBMP is one of finding such a sampling strategy to yield faster convergence.

B. The Informed Set

Let c_i be the cost of the best solution found by the planning algorithm after i iterations. The Informed Set [14] is defined as

$$\mathcal{X}_{\text{inf}} = \{\mathbf{x} \in \mathcal{X} \mid h(\mathbf{x}_s, \mathbf{x}) + h(\mathbf{x}, \mathbf{x}_g) < c_i\}. \quad (4)$$

Note that $\mathcal{X}_{\text{inf}}$ uses a heuristic approximation of both the cost-to-come $h(\mathbf{x}_s, \mathbf{x})$ and the cost-to-go $h(\mathbf{x}, \mathbf{x}_g)$ to get an (under)estimate of the *solution cost* constrained to pass through any $\mathbf{x} \in \mathcal{X}$. Generating new samples in $\mathcal{X}_{\text{inf}}$ is thus a necessary condition for improving the current solution. An algorithm for direct sampling of the L_2 -Informed Set is given in [14].

C. Relevant Region

Consider the set of *relevant vertices* defined as

$$V_{\text{rel}} = \{\mathbf{v} \in V \mid g_{\mathcal{T}}(\mathbf{v}) + h(\mathbf{v}, \mathbf{x}_g) < c_i\}. \quad (5)$$

Let $\epsilon > 0$ ball around a relevant vertex $\mathbf{v} \in V_{\text{rel}}$ be defined as

$$\mathcal{B}^\epsilon(\mathbf{v}) = \{\mathbf{x} \in \mathcal{X} \mid \|\mathbf{x} - \mathbf{v}\|_2 < \epsilon, \mathbf{v} \in V_{\text{rel}}\}. \quad (6)$$

Consider the estimate of the solution cost constrained to pass through $\mathbf{x} \in \mathcal{B}^\epsilon(\mathbf{v})$

$$\hat{f}_{\mathbf{v}}(\mathbf{x}) = d_\ell(\mathbf{v}, \mathbf{x}) + g_{\mathcal{T}}(\mathbf{v}) + h(\mathbf{x}, \mathbf{x}_g). \quad (7)$$

The *Relevant Set* around $\mathbf{v} \in V_{\text{rel}}$ is defined as

$$\mathcal{B}_{\text{rel}}^\epsilon(\mathbf{v}) = \{\mathbf{x} \in \mathcal{B}^\epsilon(\mathbf{v}) \mid \hat{f}_{\mathbf{v}}(\mathbf{x}) < c_i\}. \quad (8)$$

Using (5), (8), the *Relevant Region* is defined as the union of the relevant sets around all relevant vertices

$$\mathcal{X}_{\text{rel}}^\epsilon = \bigcup_{\mathbf{v} \in V_{\text{rel}}} \mathcal{B}_{\text{rel}}^\epsilon(\mathbf{v}). \quad (9)$$

Note that, in contrast to $\mathcal{X}_{\text{inf}}$ which uses the heuristic estimate $h(\mathbf{x}_s, \mathbf{x})$ of the cost-to-come, $\mathcal{X}_{\text{rel}}^\epsilon$ uses $d_\ell(\mathbf{v}, \mathbf{x}) + g_{\mathcal{T}}(\mathbf{v})$




Fig. 3: Planning on a “potential-field” like cost-map. The objective is to reach the goal state while avoiding the two danger (white) regions.

from (7). This approximation considers the cost-function information (see (2)), the structure of $\mathcal{T}$, and hence the topology of $\mathcal{X}_{\text{free}}$. While the L_2 -norm is still a consistent heuristic for cost-maps with $C(\mathbf{x}) \geq 1$ for all $\mathbf{x} \in \mathcal{X}$, it does not take into account C or $\mathcal{X}_{\text{obs}}$. It may provide a poor estimate of the solution cost, leading to $\mathcal{M}(\mathcal{X}_{\text{inf}}) \approx \mathcal{M}(\mathcal{X})$. Informed Sampling effectively resorts to uniform random sampling in this case. The set $\mathcal{X}_{\text{rel}}^\epsilon$ alleviates this dependence on a heuristic. The value of ϵ , which controls the size of the Relevant Set, is taken to be slightly greater than the step-size parameter η (in our benchmarking simulations, we used $\epsilon = 1.5\eta$). The step-size parameter η in SBMP controls the maximum edge length in $\mathcal{G}$ [14]. Note that a very small value of ϵ would hinder exploration, while a large value of ϵ may provide a poor estimate of the cost-to-come in (7), as the edge $(\mathbf{x}, \mathbf{v})$ may not be feasible. The following theorem proves that for any $\epsilon > 0$, $\mathcal{B}_{\text{rel}}^\epsilon(\mathbf{v})$ is not a singleton.

Theorem 1. *For every $\mathbf{v} \in V_{\text{rel}}$, there exists $\delta > 0$, such that, for all $\mathbf{x} \in \mathcal{B}^\delta(\mathbf{v})$, it follows that $\hat{f}_{\mathbf{v}}(\mathbf{x}) < c_i$.*

Proof. Note from (7) that for each given $\mathbf{v} \in V_{\text{rel}}$ the function $\hat{f}_{\mathbf{v}}$ is continuous in $\mathbf{x}$ since both $d_\ell(\cdot, \mathbf{v})$, $h(\cdot, \mathbf{x}_g)$ are continuous. Also, $\hat{f}_{\mathbf{v}}(\mathbf{v}) = g_{\mathcal{T}}(\mathbf{v}) + h(\mathbf{v}, \mathbf{x}_g) < c_i$ since $\mathbf{v} \in V_{\text{rel}}$. Since $\hat{f}_{\mathbf{v}}$ is continuous at $\mathbf{v}$, it follows that for any $\zeta > 0$ there exists $\delta > 0$ such that $\mathbf{x} \in \mathcal{B}^\delta(\mathbf{v})$ implies that $|\hat{f}_{\mathbf{v}}(\mathbf{x}) - \hat{f}_{\mathbf{v}}(\mathbf{v})| < \zeta$. Choosing $\zeta = c_i - \hat{f}_{\mathbf{v}}(\mathbf{v}) > 0$ one then obtains that for all $\mathbf{x} \in \mathcal{B}^\delta(\mathbf{v})$ we have that $|\hat{f}_{\mathbf{v}}(\mathbf{x}) - \hat{f}_{\mathbf{v}}(\mathbf{v})| < c_i - \hat{f}_{\mathbf{v}}(\mathbf{v})$ and hence $\hat{f}_{\mathbf{v}}(\mathbf{x}) < c_i$. $\square$

Corollary 2. *Let $\mathbf{v} \in V_{\text{rel}}$. For every $\epsilon > 0$ there exists $\delta > 0$ such $\mathcal{B}^\delta(\mathbf{v}) \subset \mathcal{B}_{\text{rel}}^\epsilon(\mathbf{v})$.*

Theorem 3. *For any $\epsilon > 0$, the Relevant Region $\mathcal{X}_{\text{rel}}^\epsilon$ is a subset of the Informed Set $\mathcal{X}_{\text{inf}}$.*

Proof. Let $\mathbf{x} \in \mathcal{X}_{\text{rel}}^\epsilon$. Then there exists $\mathbf{v} \in V_{\text{rel}}$, so that $\mathbf{x} \in \mathcal{B}_{\text{rel}}^\epsilon(\mathbf{v})$, and hence $d_\ell(\mathbf{x}, \mathbf{v}) + g_{\mathcal{T}}(\mathbf{v}) + h(\mathbf{x}, \mathbf{x}_g) < c_i$. Since the heuristic function is consistent, $h(\mathbf{x}, \mathbf{v}) < d_\ell(\mathbf{x}, \mathbf{v})$ and $h(\mathbf{v}, \mathbf{x}_s) < g_{\mathcal{T}}(\mathbf{v})$. Using the triangle inequality, it follows that, $h(\mathbf{x}_s, \mathbf{x}) < h(\mathbf{x}, \mathbf{v}) + h(\mathbf{v}, \mathbf{x}_s)$. Combining the above inequalities yields $h(\mathbf{x}_s, \mathbf{x}) + h(\mathbf{x}, \mathbf{x}_g) < d_\ell(\mathbf{x}, \mathbf{v}) + g_{\mathcal{T}}(\mathbf{v}) + h(\mathbf{x}, \mathbf{x}_g) < c_i$. Hence, $\mathbf{x} \in \mathcal{X}_{\text{inf}}$. It follows that $\mathcal{X}_{\text{rel}}^\epsilon \subset \mathcal{X}_{\text{inf}}$. $\square$

Theorem 3 implies that generating samples in $\mathcal{X}_{\text{rel}}^\epsilon$ does not lead to redundant exploration outside $\mathcal{X}_{\text{inf}}$. However,

note that sampling in $\mathcal{X}_{\text{rel}}^\epsilon$ is not a necessary condition for improving the current solution, i.e., there may be points $\mathbf{x} \in \mathcal{X}_{\text{inf}}$ such that $\mathbf{x} \notin \mathcal{X}_{\text{rel}}^\epsilon$ which may improve the current solution. Relevant Region sampling is thus utilized in conjunction with Informed/Uniform Sampling. As shown in the numerical examples later on, this interplay of exploration by Informed Sampling, combined with focusing properties of Relevant Region, leads to accelerated convergence.

III. SAMPLING IN THE RELEVANT REGION

Since $\mathcal{X}_{\text{rel}}^\epsilon$ depends on $\mathcal{T}$, a direct sampling strategy is not possible. Hence, the proposed sampling strategy proceeds by first selecting a relevant vertex $\mathbf{v}_p \in V_{\text{rel}}$, sampling a random direction $\hat{\mathbf{e}}$, $\|\hat{\mathbf{e}}\|_2 = 1$ and finding the maximum magnitude of travel $\gamma_{\text{rel}} > 0$ along $\hat{\mathbf{e}}$, so that for all $\gamma \in (0, \gamma_{\text{rel}})$ the new sample $\mathbf{x} = \mathbf{v}_p + \gamma\hat{\mathbf{e}} \in \mathcal{B}_{\text{rel}}^\epsilon(\mathbf{v}_p)$. Please see Fig. 4. Note that Theorem 1 guarantees the existence of γ_{rel} . Concretely, the following optimization problem needs to be solved:

$$\begin{aligned} & \sup_{\gamma \in (0, \epsilon)} \gamma, \\ & \text{subject to: } \hat{f}_{\mathbf{v}_p}(\mathbf{v}_p + \gamma\hat{\mathbf{e}}) < c_i. \end{aligned} \quad (10)$$

A. Case 1: Uniform Cost-Map

Consider the problem (10) with $C(\mathbf{x}) = 1$ for all $\mathbf{x} \in \mathcal{X}$. Using the L_2 -norm heuristic in (7), the inequality in (10) yields,

$$\hat{f}_{\mathbf{v}_p}(\mathbf{v}_p + \gamma\hat{\mathbf{e}}) = \gamma + g_{\mathcal{T}}(\mathbf{v}_p) + \|\mathbf{v}_p + \gamma\hat{\mathbf{e}} - \mathbf{x}_g\|_2 < c_i. \quad (11)$$

Rearrange the terms in (11) to obtain

$$\|\mathbf{v}_p + \gamma\hat{\mathbf{e}} - \mathbf{x}_g\|_2 < c_i - g_{\mathcal{T}}(\mathbf{v}_p) - \gamma. \quad (12)$$

To ensure that the RHS in (12) is positive, choose

$$\gamma < c_i - g_{\mathcal{T}}(\mathbf{v}_p). \quad (13)$$

Let $\mathbf{x}_{pg} = \mathbf{v}_p - \mathbf{x}_g$ and $g_{gp} = c_i - g_{\mathcal{T}}(\mathbf{v}_p)$. Also note that $\mathbf{x}_{pg}^T \mathbf{x}_{pg} = h^2(\mathbf{v}_p, \mathbf{x}_g)$ and $\mathbf{x}_{pg}^T \hat{\mathbf{e}} = h(\mathbf{v}_p, \mathbf{x}_g) \cos \theta$, where θ is the angle between the vectors $\mathbf{x}_{pg}$ and $\hat{\mathbf{e}}$. Squaring both sides in (12) yields,

$$\begin{aligned} h^2(\mathbf{v}_p, \mathbf{x}_g) + 2\gamma \mathbf{x}_{pg}^T \hat{\mathbf{e}} + \gamma^2 &< g_{gp}^2 - 2\gamma g_{gp} + \gamma^2 \\ \text{and hence } \gamma &< \frac{g_{gp}^2 - h^2(\mathbf{v}_p, \mathbf{x}_g)}{2(\mathbf{x}_{pg}^T \hat{\mathbf{e}} + g_{gp})}. \end{aligned}$$

Define the RHS in the above inequality as

$$\gamma_{\text{uni}} = \frac{(c_i - g_{\mathcal{T}}(\mathbf{v}_p))^2 - h^2(\mathbf{v}_p, \mathbf{x}_g)}{2[h(\mathbf{v}_p, \mathbf{x}_g) \cos \theta + (c_i - g_{\mathcal{T}}(\mathbf{v}_p))]} \quad (14)$$

Note that $\gamma_{\text{uni}} > 0$ for $\mathbf{v}_p \in V_{\text{rel}}$, and attains its maximum value $\bar{\gamma}_{\text{uni}}$ at $\theta = \pi$, in which case,

$$\bar{\gamma}_{\text{uni}} = (c_i - g_{\mathcal{T}}(\mathbf{v}_p) + h(\mathbf{v}_p, \mathbf{x}_g))/2,$$

and also, $\bar{\gamma}_{\text{uni}} < c_i - g_{\mathcal{T}}(\mathbf{v}_p)$ for $\mathbf{v}_p \in V_{\text{rel}}$, satisfying (13). Thus, the solution to problem (10) for uniform cost-map is

$$\gamma_{\text{rel}} = \min(\gamma_{\text{uni}}, \epsilon). \quad (15)$$




Fig. 4: A schematic for Relevant Region sampling.

B. Case 2: General Cost-Maps

Now consider the problem (10) with $C(\mathbf{x}) > 1$ for all $\mathbf{x} \in \mathcal{X}$. The following inequality needs to be solved for γ ,

$$\gamma \int_0^1 C(\mathbf{v}_p + \gamma\hat{\mathbf{e}}s) ds + g_{\mathcal{T}}(\mathbf{v}_p) + h(\mathbf{v}_p + \gamma\hat{\mathbf{e}}, \mathbf{x}_g) < c_i. \quad (16)$$

Often, C may not have a tractable closed-form expression and hence the planner has access only to the value of C at any point in the search space. In order to avoid a computationally expensive procedure to solve (16), we let

$$d_\ell(\mathbf{v}_p, \mathbf{v}_p + \gamma\hat{\mathbf{e}}) = \gamma \int_0^1 C(\mathbf{v}_p + \gamma\hat{\mathbf{e}}s) ds \approx \gamma C(\mathbf{v}_p) \quad (17)$$

Note that (17) uses a zeroth-order approximation of the integrand to estimate the integral. Higher order approximations are possible, but these will result in a computationally more involved process to find γ (see below). It follows from (17) that

$$\gamma C(\mathbf{v}_p) + g_{\mathcal{T}}(\mathbf{v}_p) + \|\mathbf{v}_p + \gamma\hat{\mathbf{e}} - \mathbf{x}_g\|_2 < c_i, \quad (18)$$

$$\text{or, } \|\mathbf{v}_p + \gamma\hat{\mathbf{e}} - \mathbf{x}_g\|_2 < c_i - g_{\mathcal{T}}(\mathbf{v}_p) - \gamma C(\mathbf{v}_p). \quad (19)$$

To ensure that the RHS of (19) is positive, choose

$$\gamma < (c_i - g_{\mathcal{T}}(\mathbf{v}_p))/C(\mathbf{v}_p). \quad (20)$$

Let again $\mathbf{x}_{pg} = \mathbf{v}_p - \mathbf{x}_g$, $g_{gp} = c_i - g_{\mathcal{T}}(\mathbf{v}_p)$, $\mathbf{x}_{pg}^T \mathbf{x}_{pg} = h^2(\mathbf{v}_p, \mathbf{x}_g)$ and $\mathbf{x}_{pg}^T \hat{\mathbf{e}} = h(\mathbf{v}_p, \mathbf{x}_g) \cos \theta$, where θ is the angle between the vectors $\mathbf{x}_{pg}$ and $\hat{\mathbf{e}}$. Squaring both sides in (19) and simplifying yields,

$$\gamma^2(C^2(\mathbf{v}_p) - 1) - 2\gamma(g_{gp}C(\mathbf{v}_p) + \mathbf{x}_{pg}^T \hat{\mathbf{e}}) + g_{gp}^2 - h^2(\mathbf{v}_p, \mathbf{x}_g) > 0. \quad (21)$$

Let γ_1, γ_2 be the roots of the quadratic equation corresponding to inequality (21), and assume $\gamma_2 > \gamma_1$.

$$\begin{aligned} \gamma_2 &= \frac{g_{gp}C(\mathbf{v}_p) + h(\mathbf{v}_p, \mathbf{x}_g) \cos \theta + \sqrt{\Delta}}{(C^2(\mathbf{v}_p) - 1)} \\ \gamma_1 &= \frac{g_{gp}C(\mathbf{v}_p) + h(\mathbf{v}_p, \mathbf{x}_g) \cos \theta - \sqrt{\Delta}}{(C^2(\mathbf{v}_p) - 1)} \\ \Delta &= (g_{gp}C(\mathbf{v}_p) + h(\mathbf{v}_p, \mathbf{x}_g) \cos \theta)^2 \\ &\quad - (C^2(\mathbf{v}_p) - 1)(g_{gp}^2 - h^2(\mathbf{v}_p, \mathbf{x}_g)). \end{aligned} \quad (22)$$

The maximum and minimum values of the radicand Δ are obtained at $\theta = 0$ and $\theta = \pi$, respectively, where

$$(g_{gp} - h(\mathbf{v}_p, \mathbf{x}_g)C(\mathbf{v}_p))^2 \leq \Delta \leq (g_{gp} + h(\mathbf{v}_p, \mathbf{x}_g)C(\mathbf{v}_p))^2. \quad (23)$$

Hence, $\gamma_1, \gamma_2 \in \mathbb{R}_{\geq 0}$ for $\mathbf{v}_p \in V_{\text{rel}}$. Then (21) yields,

$$(\gamma - \gamma_1)(\gamma - \gamma_2) > 0. \text{ equivalently, } \gamma > \gamma_2 \text{ or } \gamma < \gamma_1. \quad (24)$$

Algorithm 1: Sampling Algorithm

```

1  $V \leftarrow \{\mathbf{x}_s\}; E \leftarrow \phi; \mathcal{G} \leftarrow (V, E);$ 
2 for  $i = 1 : N$  do
3    $c_i \leftarrow \min_{\mathbf{v} \in V_{\text{goal}}} g_{\mathcal{T}}(\mathbf{v});$ 
4    $u_{\text{rand}} \sim \mathcal{U}(0, 1);$ 
5   if  $u_{\text{rand}} < p_{\text{rel}}$  and  $c_i < \infty$  then
6      $\mathbf{v}_p \leftarrow \text{chooseVertex}(V_{\text{rel}});$ 
7      $\hat{\mathbf{e}} \leftarrow \text{generateDirection}();$ 
8      $\gamma_{\text{rel}} \leftarrow \text{RelevantStepLimit}(\mathbf{v}_p, \hat{\mathbf{e}});$ 
9      $u_{\text{rand}} \sim \mathcal{U}(0, 1);$ 
10     $\mathbf{x}_{\text{rand}} \leftarrow \mathbf{v}_p + (u_{\text{rand}})^{\frac{1}{d}} \gamma_{\text{rel}} \hat{\mathbf{e}};$ 
11  else
12     $\mathbf{x}_{\text{rand}} \leftarrow \text{InformedSampling}();$ 
13   $\mathbf{x}_{\text{new}} \leftarrow \text{Extend}(\mathbf{x}_{\text{rand}});$ 
14   $\text{Exploitation}(\mathcal{G});$ 
15 return  $\mathcal{G}$ 

```

Consider the larger root γ_2 from (22). The minimum value of γ_2 is attained when $\theta = \pi$, so that

$$\gamma_2 \geq \frac{g_{\text{gp}}C(\mathbf{v}_p) - h(\mathbf{v}_p, \mathbf{x}_g) + |g_{\text{gp}} - h(\mathbf{v}_p, \mathbf{x}_g)C(\mathbf{v}_p)|}{(C^2(\mathbf{v}_p) - 1)}. \quad (25)$$

Define the RHS in (25) as $\bar{\gamma}_2$. Simplifying yields,

$$\bar{\gamma}_2 = \begin{cases} \frac{g_{\text{gp}} + h(\mathbf{v}_p, \mathbf{x}_g)}{C(\mathbf{v}_p) + 1}, & g_{\text{gp}} < h(\mathbf{v}_p, \mathbf{x}_g)C(\mathbf{v}_p), \\ \frac{g_{\text{gp}} - h(\mathbf{v}_p, \mathbf{x}_g)}{C(\mathbf{v}_p) - 1}, & g_{\text{gp}} > h(\mathbf{v}_p, \mathbf{x}_g)C(\mathbf{v}_p). \end{cases} \quad (26)$$

Note that $\bar{\gamma}_2 > g_{\text{gp}}/C(\mathbf{v}_p)$. This implies $\gamma_2 > g_{\text{gp}}/C(\mathbf{v}_p)$, violating (20). Thus, $\gamma > \gamma_2$ is an infeasible solution of (18). Next, consider γ_1 . Differentiating with respect to θ , the extrema are obtained at $\theta = 0, \pi$. Calculating the second derivative yields, $\gamma_1''(\theta = 0) > 0$ and $\gamma_1''(\theta = \pi) < 0$. The maximum value of γ_1 obtained at $\theta = \pi$ is given by

$$\bar{\gamma}_1 = \begin{cases} \frac{g_{\text{gp}} + h(\mathbf{v}_p, \mathbf{x}_g)}{C(\mathbf{v}_p) + 1}, & g_{\text{gp}} > h(\mathbf{v}_p, \mathbf{x}_g)C(\mathbf{v}_p), \\ \frac{g_{\text{gp}} - h(\mathbf{v}_p, \mathbf{x}_g)}{C(\mathbf{v}_p) - 1}, & g_{\text{gp}} < h(\mathbf{v}_p, \mathbf{x}_g)C(\mathbf{v}_p). \end{cases} \quad (27)$$

Now, $\bar{\gamma}_1 < g_{\text{gp}}/C(\mathbf{v}_p)$. This implies $\gamma_1 < g_{\text{gp}}/C(\mathbf{v}_p)$. It follows that $\gamma < \gamma_1$ satisfies (20). Thus, the solution to problem (10) with the approximation in (18) is

$$\gamma_{\text{rel}} = \min(\gamma_1, \epsilon). \quad (28)$$

For the special case when $\Delta = 0$ and $\gamma_1 = \gamma_2 = \gamma_c$, inequality (21) simplifies to $(\gamma - \gamma_c)^2 > 0$. Considering (20) yields $\gamma_{\text{rel}} = \min(g_{\text{gp}}/C(\mathbf{v}_p), \epsilon)$. Note that if $C(\mathbf{v}_p) = 1$, then inequality (16) reduces to (11) and the analysis for uniform cost-maps is applicable.

IV. PROPOSED ALGORITHM

The outline of the proposed algorithm is given in Algorithm 1. The procedure initializes a vertex at the start state $\mathbf{x}_s$. At every iteration, the current best solution cost c_i is updated (line 3). If a sub-optimal solution exists (c_i is finite), with probability p_{rel} (line 5), Relevant Region sampling is employed to generate a new random sample $\mathbf{x}_{\text{rand}}$. Otherwise, conventional Informed Sampling is used. Relevant Region sampling consists of first choosing a relevant vertex $\mathbf{v}_p$, generating a random direction $\hat{\mathbf{e}}$ and calculating the maximum




Fig. 5: Planning for 7 DOF Panda Arm in the joint space from the start state (left) to a given joint goal state (right).




Fig. 6: Percentage of successful trials (where planner found a feasible solution) with different sampling strategies.

magnitude of travel along $\hat{\mathbf{e}}$ (line 6-8). If $C(\mathbf{v}_p) = 1$, then (15) is used for obtaining γ_{rel} along $\hat{\mathbf{e}}$, else (28) is used. The exponent $1/d$ (line 10) biases the travel magnitude towards γ_{rel} and promotes exploration. After $\mathbf{x}_{\text{rand}}$ is generated, conventional SBMP modules incorporate a new vertex $\mathbf{x}_{\text{new}}$ in $\mathcal{G}$ (line 13). These include: a) finding the nearest neighbor $\mathbf{x}_{\text{nearest}}$ to $\mathbf{x}_{\text{rand}}$ in $\mathcal{G}$; b) local steering from $\mathbf{x}_{\text{nearest}}$ in the direction of $\mathbf{x}_{\text{rand}}$ to obtain $\mathbf{x}_{\text{new}}$; c) ensuring feasibility of edge-connections in the neighborhood of $\mathbf{x}_{\text{new}}$. This is followed by the exploitation module (local/global rewiring, etc). The chooseVertex module selects a relevant vertex to be expanded from the set V_{rel} . Similar to the procedure in Guided-ESTs [18] a weight $q_{\mathbf{v}}$ is allocated for each $\mathbf{v} \in V_{\text{rel}}$.

$$q_{\mathbf{v}} = \lambda_1 p_{\mathbf{v}} + \lambda_2 d_{\mathbf{v}} + \lambda_3 (g_{\mathcal{T}}(\mathbf{v}) + h(\mathbf{v}, \mathbf{x}_g)) / c_i. \quad (29)$$

Here, $p_{\mathbf{v}}$ represents the number of times $\mathbf{v}$ has been selected in the past. This penalizes multiple selections and the exploration of the region around a particular vertex. The second term, $d_{\mathbf{v}}$ is the number of edges connected to $\mathbf{v}$. It promotes sampling in relatively unexplored regions. The last term $0 < (g_{\mathcal{T}}(\mathbf{v}) + h(\mathbf{v}, \mathbf{x}_g)) / c_i < 1$ is the estimate of the solution cost through $\mathbf{v}$, normalized by the current best cost. This prioritizes exploration of regions with low solution cost estimates. The parameters $(\lambda_1, \lambda_2, \lambda_3) > 0$ modulate the behavior of the selection algorithm. A large value of λ_3 leads to a greedy focus on low solution cost areas, whereas increasing λ_1, λ_2 promotes exploration. A binary heap is used to update and sort V_{rel} according to the weight in (29). A relevant vertex $\mathbf{v}_p$ is selected by choosing randomly from the top n_q elements in the sorted list. This injects randomness in the selection process and promotes desirable exploration.




Fig. 7: Convergence plots for different sampling methods in various test environments. Solid lines indicate the average value and the standard deviation is shaded. Error bar indicate the upper and lower quartiles.

V. NUMERICAL EXPERIMENTS

The performance of the proposed sampling method was benchmarked against direct Informed Sampling [14] and the third variant of adaptive rejection sampling (described in [15]) in uniform cost-space environments (length-optimal planning). For all experiments, the exploration strategies were paired with RRT[#]'s dynamic programming based global rewiring for exploitation. In general cost-map environments, benchmarking was done against Informed Sampling and T-RRT[#] (combining conventional RRT[#] with the transition-test described in [22]) with different initial temperatures T_{init} . All the algorithms were implemented in C++ using the popular OMPL framework [26], and the tests were run using OMPL's standardized benchmarking tools [27]. Please see, https://github.gatech.edu/DCSL/relevant_region. A 64-bit desktop PC with 64 GB RAM and an Intel Xeon(R) Processor running Ubuntu 16.04 OS was used. The data was recorded over 100 trials for all the cases. The proposed algorithm used the following parameter values: $\epsilon = 1.5\eta$, $(\lambda_1, \lambda_2, \lambda_3) = (10, 5, 100)$, $n_q = 10$. A goal bias of 5% was used in all sampling methods. A description of the different environments is provided below.

A. Uniform Cost-Map Cases

Multiple Obstacle World: This environment is illustrated in Fig. 1. The 2D environment was extended to $\mathbb{R}^4$ and $\mathbb{R}^6$ by imparting a length of 2 units symmetrically to all of the obstacles. A step-size of $\eta = 0.6$ and $\eta = 1.2$ was used in $\mathbb{R}^4$ and $\mathbb{R}^6$ respectively.

Panda Arm: A planning problem for Panda Arm (by Franka Emika) is illustrated in Fig. 5. The objective was to find a minimum length path in a 7-dimensional configuration (joint) space with joint limits ($\mathbb{R}^7$). These limits and collision checking module were implemented using MoveIt! [28]. The step-size was set to $\eta = 0.7$ for this example.

B. General Cost-Map Cases

Terrain Map: A 2D terrain map shown in Fig. 2 consists of rough, high-cost white areas and the easily navigable black regions. The step-size was set to $\eta = 0.3$ for this example.

Potential Cost-Map: The environment in Fig. 3 emulates the problem of finding the shortest path while staying away from danger areas (white regions). The cost function is defined as

$$C(\mathbf{x}) = 1 + 9(e^{-\frac{\|\mathbf{x}_1^d - \mathbf{x}\|_2^2}{5}} + e^{-\frac{\|\mathbf{x}_2^d - \mathbf{x}\|_2^2}{5}}). \quad (30)$$

Here, $\mathbf{x}_1^d, \mathbf{x}_2^d$ are the center points of the danger regions. A step-size of $\eta = 0.6$ and $\eta = 1.5$ was used in $\mathbb{R}^4$ and $\mathbb{R}^6$ version of the environment respectively.

VI. CONCLUSION

This work proposes a novel algorithm to sample the Relevant Region set, a subset of the Informed Set, for SBMP. The Relevant Region set considers the topology of $\mathcal{X}_{\text{free}}$, reduces the dependence on heuristics, and effectively focuses the search to accelerate convergence. Numerical experiments validate the utility of Relevant Region sampling in conjunction with Informed/Uniform Sampling. The proposed method leads to faster convergence in all cases (see Fig. 7). This is observed especially in higher dimensional problem instances. Transition-test based exploration is more effective than purely Uniform/Informed Sampling for planning on general cost-maps. However, the tendency to (probabilistically) reject samples may hinder exploration in some cases. This can be seen in the terrain cost-map (Fig. 2) which is similar to the cost-space chasms scenario described in [24]. As conveyed in Fig. 6, the transition-test based exploration fails to find a feasible solution in roughly 40% of total trials, whereas the proposed method finds a solution in all trials and also accelerates the convergence.

This work presents many avenues for future research. The proposed method inherits an additional computational overhead to maintain V_{rel} . However, this can be alleviated by leveraging ideas from sparse tree planners [29], [30] to maintain a sparse set of vertices in V_{rel} . The cost function's gradient information (if available) can be used to bias the search. Data from past iterations can also be used to infer the nature of cost-map for intelligent exploration.

Acknowledgment: This work has been supported by NSF award IIS-1617630.

REFERENCES

- [1] 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.
- [2] 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, June 2011.
- [3] O. Arslan and P. Tsiotras, "Use of relaxation methods in sampling-based algorithms for optimal motion planning," in *IEEE International Conference on Robotics and Automation*, Karlsruhe, Germany, May 6–10 2013, pp. 2421–2428.
- [4] 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 *IEEE International Conference on Robotics and Automation*, Seattle, WA, May, 25–30 2015, pp. 3067–3074.
- [5] L. Janson, E. Schmerling, A. Clark, and M. Pavone, "Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 883–921, May 2015.
- [6] F. Hauer and P. Tsiotras, "Deformable rapidly-exploring random trees," in *Robotics: Science and Systems*, Cambridge, MA, July 12–16 2017.
- [7] B. Akgun and M. Stilman, "Sampling heuristics for optimal motion planning in high dimensions," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Francisco, CA, Sept. 25–30 2011, pp. 2640–2645.
- [8] C. Urnson and R. Simmons, "Approaches for heuristically biasing RRT growth," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 2, Las Vegas, NV, Oct. 27–31 2003, pp. 1178–1183.
- [9] B. Burns and O. Brock, "Toward optimal configuration space sampling," in *Robotics: Science and Systems*, Cambridge, MA, June 8–11 2005, pp. 105–112.
- [10] —, "Single-query entropy-guided path planning," in *IEEE International Conference on Robotics and Automation*, Barcelona, Spain, 4 2005, pp. 2124–2129.
- [11] R. Diankov and J. Kuffner, "Randomized statistical path planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, San Diego, CA, Oct. 29–31 2007, pp. 1–6.
- [12] S. Rodriguez, X. Tang, J.-M. Lien, and N. M. Amato, "An obstacle-based rapidly-exploring random tree," in *IEEE International Conference on Robotics and Automation*, Orlando, FL, May 15–19 2006, pp. 895–900.
- [13] J. Denny, E. Greco, S. Thomas, and N. M. Amato, "MARRT: Medial axis biased rapidly-exploring random trees," in *IEEE International Conference on Robotics and Automation*, Hong Kong, China, June 1–7 2014, pp. 90–97.
- [14] J. D. Gammell, T. D. Barfoot, and S. S. Srinivasa, "Informed sampling for asymptotically optimal path planning," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 966–984, Aug. 2018.
- [15] O. Arslan and P. Tsiotras, "Dynamic programming guided exploration for sampling-based motion planning algorithms," in *IEEE International Conference on Robotics and Automation*, Seattle, WA, May 26–30 2015, pp. 4819–4826.
- [16] —, "Machine learning guided exploration for sampling-based motion planning algorithms," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, Hamburg, Germany, Sept. 28–Oct. 2, 2015, pp. 2646–2652.
- [17] D. Hsu, J.-C. Latombe, and R. Motwani, "Path planning in expansive configuration spaces," in *IEEE International Conference on Robotics and Automation*, vol. 3, Albuquerque, NM, April 25–29 1997, pp. 2719–2726.
- [18] J. M. Phillips, N. Bedrossian, and L. E. Kavraki, "Guided expansive spaces trees: a search strategy for motion-and cost-constrained state spaces," in *IEEE International Conference on Robotics and Automation*, New Orleans, LA, April 26–30 2004, pp. 3968–3973.
- [19] S. M. Persson and I. Sharf, "Sampling-based A* algorithm for robot path-planning," *The International Journal of Robotics Research*, vol. 33, no. 13, pp. 1683–1708, 10 2014.
- [20] J. Mainprice, E. A. Sisbot, L. Jaillet, J. Cortés, R. Alami, and T. Siméon, "Planning human-aware motions using a sampling-based costmap planner," in *IEEE International Conference on Robotics and Automation*, Shanghai, China, May 9–13 2011, pp. 5012–5017.
- [21] L. Jaillet, F. J. Corcho, J.-J. Pérez, and J. Cortés, "Randomized tree construction algorithm to explore energy landscapes," *Journal of Computational Chemistry*, vol. 32, no. 16, pp. 3464–3474, 2011.
- [22] L. Jaillet, J. Cortés, and T. Siméon, "Sampling-based path planning on configuration-space costmaps," *IEEE Transactions on Robotics*, vol. 26, no. 4, pp. 635–646, 8 2010.
- [23] D. Devaurs, T. Siméon, and J. Cortés, "Enhancing the transition-based RRT to deal with complex cost spaces," in *IEEE International Conference on Robotics and Automation*, Karlsruhe, Germany, May 6–10 2013, pp. 4120–4125.
- [24] D. Berenson, T. Siméon, and S. S. Srinivasa, "Addressing cost-space chasms in manipulation planning," in *IEEE International Conference on Robotics and Automation*, Shanghai, China, May 9–13 2011, pp. 4561–4568.
- [25] D. Devaurs, T. Siméon, and J. Cortés, "Optimal path planning in complex cost spaces with sampling-based algorithms," *IEEE Transactions on Automation Science and Engineering*, vol. 13, no. 2, pp. 415–424, 2015.
- [26] I. A. Sucan, M. Moll, and L. E. Kavraki, "The open motion planning library," *IEEE Robotics & Automation Magazine*, vol. 19, no. 4, pp. 72–82, Dec. 2012.
- [27] M. Moll, I. A. Sucan, and L. E. Kavraki, "Benchmarking motion planning algorithms: An extensible infrastructure for analysis and visualization," *IEEE Robotics & Automation Magazine*, vol. 22, no. 3, pp. 96–102, 2015.
- [28] S. Chitta, I. Sucan, and S. Cousins, "Moveit![ROS topics]," *IEEE Robotics & Automation Magazine*, vol. 19, no. 1, pp. 18–19, 2012.
- [29] A. Dobson and K. E. Bekris, "Sparse roadmap spanners for asymptotically near-optimal motion planning," *The International Journal of Robotics Research*, vol. 33, no. 1, pp. 18–47, 2014.
- [30] Y. Li, Z. Littlefield, and K. E. Bekris, "Sparse methods for efficient asymptotically optimal kinodynamic planning," in *Algorithmic Foundations of Robotics XI*. Springer, 2015, pp. 263–282.