

Incremental Generalized Hybrid A*

Sidharth Talia¹, Oren Salzman², and Siddhartha Srinivasa¹

Abstract—We address the problem of efficiently organizing search over very large trees, which arises in many applications ranging from autonomous driving to aerial vehicles. Here, we are motivated by off-road autonomy, where real-time planning is essential. Classical approaches use graphs of motion primitives and exploit dominance to mitigate the curse of dimensionality and prune expansions efficiently. However, for complex dynamics, repeatedly solving two-point boundary-value problems makes graph construction too slow for fast kinodynamic planning. Hybrid A* (HA*) addressed this challenge by searching over a tree of motion primitives and introducing approximate pruning using a grid-based dominance check. However, choosing the grid resolution is difficult: too coarse risks failure, while too fine leads to excessive expansions and slow planning. We propose Incremental Generalized Hybrid A* (IGHA*), an anytime tree-search framework that dynamically organizes vertex expansions without rigid pruning. IGH A* provably matches or outperforms HA*. For both on-road kinematic and off-road kinodynamic planning queries for a car-like robot, variants of IGH A* use $6\times$ fewer expansions to the best solution compared to an optimized version of HA* (HA*M, an internal baseline). In simulated off-road experiments in a high-fidelity simulator, IGH A* outperforms HA*M when both are used in the loop with a model predictive controller. We demonstrate real-time performance both in simulation and on a small-scale off-road vehicle, enabling fast, robust planning under complex dynamics. Website: <https://personalrobotics.github.io/IGHAStar/>

Index Terms—Motion and Path Planning; Field Robots; Autonomous Vehicle Navigation

I. INTRODUCTION

WE focus on the problem of efficiently organizing search over very large trees, a key challenge for motion planning with complex dynamics and costs often encountered in autonomous driving [1] and aerial vehicles [2]. In this paper, we are particularly motivated by the problem of planning for autonomous driving in unstructured, off-road environments where the autonomous vehicle (AV) only has access to a body-centric local map [3], [4]. This problem in general exists for more challenging aspects of on-road driving as well, which lack rules or structure, such as driving in parking lots, construction zones, and rural areas.

Manuscript received: August, 7, 2025; Revised October, 20, 2025; Accepted November, 22, 2025. This paper was recommended for publication by Editor Olivier Stasse upon evaluation of the Associate Editor and Reviewers' comments. This work was (partially) funded by the National Science Foundation NRI (#2132848) and CHS (#2007011), DARPA RACER (#HR0011-21-C-0171), the Office of Naval Research (#N00014-17-1-2617-P00004 and #2022-016-01 UW), Amazon, Technion Autonomous Systems Program (TASP), and the United States-Israel Binational Science Foundation (BSF) grant (#2021643). ¹University of Washington, Seattle, USA. {sidhtalia, siddh}@cs.washington.edu, ²Technion-Israel Institute of Technology, Haifa, Israel. osalzman@cs.technion.ac.il. Digital Object Identifier (DOI): see top of this page.

This problem has traditionally been cast as a search over a graph of primitive motions, with many advances in constructing good graphs [5], [6] and in searching efficiently over them [7], [8]. A key reason for the efficiency is the notion of *dominance* that arises from the Bellman condition [9]; when a new vertex lands on top of an existing vertex, only the one with a lower cost-to-come survives, and the other is pruned. Thus, given a branching factor b and path length d , far fewer than b^d vertices are practically expanded. Unfortunately, applications where the environment affects the dynamics of the system, as is the case for off-road autonomy [3], [10], will require repeatedly solving the two-point boundary-value problem (BVP) to build the graph [11], which can be too slow for kinodynamic planning for fast moving systems with computational constraints [12], [13].

Dolgov et al. [1] addressed this challenge by proposing the Hybrid A*(HA*) algorithm, which commits to search over a tree of motion primitives instead. Their key insight was to introduce the notion of *approximate dominance*, where an expanded tree vertex only had to fall within a grid cell of a previously-searched vertex for dominance to be checked. Hybrid A* soon became the de facto standard for planning for unstructured driving [14], [15]. However, the resolution of this grid is hard to tune. A low resolution over-prunes (Fig.1(a)), failing to find a path, whereas a high resolution over-expands, resulting in slow planning times (Fig.1(e)). Effectively, there is a coupling between vertex discovery and approximate dominance.

Our key insight is that we can reframe search over trees not as rigidly pruning the large tree like HA* proposes, but as dynamically organizing vertex expansion via an *anytime* planning framework which breaks this coupling. Our framework, which we term Incremental Generalized Hybrid A* (IGH A*), enables automatically changing the resolution at which vertex expansions occur. We show how different variants of IGH A* organize search based on the structural assumptions they make. In addition, we prove that IGH A* always does as well as, if not better than HA*. We evaluate IGH A* on planning queries from the Moving AI benchmark [16] and BeamNG simulator [17], showing that where an internal baseline – HA*M, an optimized version of HA* that also increases its search resolution—finds the optimal path, all variants of IGH A* find the same path with $\approx 6\times$ fewer expansions. In simulated off-road autonomy experiments using the BeamNG [17] simulator, IGH A* outperforms HA*M in completion time under a fixed expansion budget when run in a closed-loop fashion with an MPC [3], running at ≈ 4 Hz. We also demonstrate real-time performance (≈ 3 -4 Hz) on a 1/10th scale off-road vehicle sharing compute with MPC and perception on an Nvidia Jetson Orin NX.




Fig. 1: HA* operating at different resolutions for the same planning query. Vertices found (Open list) are shown in gray, pruned vertices in light gray, invalid (collision) in red, the goal set is marked as the circular region on the bottom right, the path is shown in black, and the gray bar on top represents the number of expansions. When using HA*, we don't a priori know the resolution for the first or the best solution, as shown in (c), (d). As a result, if the discretization is too coarse, we can end up with no solution as shown in (a) and (b) respectively. If it is too fine, we expand more than necessary as shown in (e), where the solution quality does not improve compared to (d) despite $\approx 3\times$ the number of expansions in the former.

II. RELATED WORK

Off-road environments often lack structured cues like lane markers or signs, and systems are typically deployed without prior maps [18], requiring real-time planning in higher-dimensional spaces. Methods such as RRT* [19] have been applied to off-road planning in kinematic space [20], while kinodynamic planning has mostly focused on systems with linearizable dynamics [21], often unsuitable for off-road vehicles that depend heavily on complex terrain interactions [3], [10]. Recent work [11] addresses this gap by explicitly constructing graphs with full kinodynamic models. However, all these approaches require repeatedly solving costly two-point BVPs [12], [13]. Here, we commit to efficient search over implicit trees generated purely via forward model propagation.

HA* [1] is popular for on-road navigation [14], [15] relying on the notion of pruning (approximately) *duplicate* states as introduced by Barraquand and Latombe [22]. Methods like Sparse-RRT [13] and SST [12] also prune states to accelerate search, using r-discs instead of grids. Other work [23], [24] defines duplicates via similarity of successor validity. In general, these methods are sensitive to the resolution (or equivalent hyperparameters) used for duplicate detection. Here, we propose a general framework to automatically adapt this hyperparameter on the fly in an anytime fashion, akin to how Anytime A*-like algorithms [25]–[27] adjust heuristic weights dynamically.

Methods like Multi-Resolution A* (MRA*) [28], [29] perform simultaneous searches at multiple resolutions, with discrete graphs defined at multiple spatial resolutions but fundamentally rely on exact state overlap across graphs for information sharing, which is not guaranteed (measure 0) with complex system dynamics in continuous space. In contrast, our approach, as HA* [1], commits to searching over the tree generated by the system's dynamics. Prior approaches [30], [31] adapt resolution *locally* (such as near obstacles) by relying on domain knowledge. Our method changes resolution globally and requires no such priors, making it domain-agnostic. Our work builds on the notion of the informed set from prior work [32]–[34], and is conceptually related to graph densification approaches [35], where the search is performed over a pre-defined set of graphs of increasing density. However, our trees are not predefined; instead, we generate sub-trees on the fly by forward-propagating the system dynamics.

III. COMPLEXITY OF SEARCHING OVER A TREE

We consider the problem of computing a near-optimal path over an implicitly-defined tree $T = (V, E)$, where each edge $(u, v) \in E$ has an associated positive cost given by a function $w : E \rightarrow \mathbb{R}^+$. The tree is rooted at a start vertex $v_s \in V$, and the objective is to reach any vertex in a goal set $V_g \subseteq V$ while minimizing the accumulated edge cost. The tree structure is defined implicitly by a successor function SUCC that returns the set of valid children for any given vertex. With a slight abuse of notation, we set $w(\pi)$ to be the cost of a path $\pi = \langle v_1, v_2, \dots, v_n \rangle$. Namely, $w(\pi) := \sum_{i=1}^{n-1} w(v_i, v_{i+1})$. Additionally, we define $w(\emptyset) = \infty$ to denote the cost of failure. We assume access to a heuristic function $h : V \rightarrow \mathbb{R}_{\geq 0}$ that is consistent and admissible, providing an estimate of the remaining cost to the goal. A baseline algorithm for solving this problem is A* search applied directly to the tree. The algorithm maintains a priority queue (Q_v) of vertices sorted by $v.f = v.g + h(v)$, where $v.g$ is the cost of the path from v_s to v , and $h(v)$ is the heuristic estimate to the goal. The algorithm begins by inserting v_s into the queue with $v_s.g = 0$, and repeatedly expands the node with the lowest f -value, generating all of its children via SUCC , until a goal vertex in V_g is reached. As this algorithm considers every child at each expansion, its worst-case complexity is exponential in the depth d of the tree, i.e., $O(b^d)$, where b is the maximum branching factor. Since no two vertices coincide in the tree, there is no opportunity to prune, and the true complexity worryingly approaches the worst case. This is effectively Alg. 1 without the approximate pruning block highlighted in red.

IV. HYBRID A* (HA*)

The Hybrid A* algorithm (HA*, Alg. 1) was introduced by Dolgov et al. [1] for solving the motion-planning problem for a kinematic car in urban environments by reducing it to a search over a large tree. Like the aforementioned baseline algorithm, HA* also expands vertices in order of increasing f -value. The key difference lies in how it reduces search effort by *approximate dominance* (highlighted in red in Alg. 1): when two vertices fall into the same *discretized-region* of the continuous space (a grid cell of resolution R in [1]), they are subject to a dominance check. The function $\hat{v} : V \times \mathbb{N} \rightarrow V$ maps a vertex v to the vertex that currently dominates v 's grid cell at resolution R . If v has a lower g -value than its dominant

Algorithm 1 Hybrid A* Search on a Tree

Input: Start vertex v_s , Goal vertex set V_g , Resolution R
Output: Least-cost path or $\emptyset$ if FAILURE

```

1: function HA*( $v_s, V_g, R$ )
2:    $Q_v \leftarrow \{v_s\}$  ▷ Priority:  $v.g + h(v)$ 
3:   while  $Q_v \neq \emptyset$  do
4:      $u \leftarrow Q_v.\text{POP}()$ 
5:     if  $u \in V_g$  then
6:       return EMITPATH()
7:     for all  $v \in \text{SUCC}(u)$  do
8:        $v.g \leftarrow u.g + w(u, v)$ 
9:       if  $v.g < \hat{v}(v, R).g$  then ▷ Approx. dominance
10:         $Q_v.\text{Remove}(\hat{v}(v, R))$ 
11:         $\hat{v}(v, R) \leftarrow v$ 
12:         $Q_v.\text{Insert}(v)$ 
13:   return  $\emptyset$ 

```

Algorithm 2 Multi-Resolution Hybrid A* (HA*M)

Input: start vertex v_s , goal set V_g , resolution sequence $\mathcal{R}$
Output: Least-cost path or $\emptyset$ if FAILURE

```

1:  $\hat{\pi} \leftarrow \emptyset$ 
2: for  $i = 0$  to  $N$  do
3:    $\pi \leftarrow \text{HA}^*(v_s, V_g, R_i)$ 
4:   if  $w(\pi) < w(\hat{\pi})$  then
5:      $\hat{\pi} \leftarrow \pi$ 
6: return  $\hat{\pi}$ 

```

vertex (Alg. 1 Line 9), it becomes the dominant vertex in this grid cell and is inserted into the queue, and the pre-existing dominant is discarded. By expanding the domain of dominance from a single vertex to a region, approximate pruning can greatly reduce the number of vertices expanded. Unfortunately, selecting the correct resolution can be extremely hard in practice. Fig. 1 demonstrates that the discretization resolution used by HA* can be either too low or too high for a given problem. In the case of off-road kinodynamic planning, the successor vertices may be farther or closer depending on the speed of the parent vertex, and the speed can change with each successor. The discretization resolution that may seem to work reasonably well (Fig. 1(c)) in one search query can easily result in anything between not finding a solution (Fig. 1(a)) to spending too many expansions for a solution (Fig. 1(e)).

V. INCREMENTAL GENERALIZED HYBRID A* (IGHA*)

A. Warmup: HA*M- Simple multi-resolution HA* search

A straightforward extension of HA* is to restart the forward search at successively finer *levels* of resolution (given by a sequence $\mathcal{R} = [R_0, R_2, \dots, R_N]$), an algorithm which we call HA*M (Alg. 2). At each *iteration*, HA*M runs a complete search from scratch at the current resolution level, discarding all previously expanded vertices. In practice, we also branch-and-bound on the current best cost $w(\hat{\pi})$ to improve efficiency. Specifically, when we pop a vertex v from the queue, we terminate the search if $v.f \geq w(\hat{\pi})$. While HA*M can eventually recover the first and best path, it is still computationally inefficient. States that were expanded and pruned at coarser resolutions are repeatedly rediscovered through fresh expansions at finer levels. Our key observation is that although pruning is valuable, prematurely pruning vertices can cut off useful vertices forever.




Fig. 2: IGH A* decouples vertex dominance and activation. ((a)) Red arrows denote breaking a loop. ((b)) Instantiations of SHIFT, ACTIVATE for IGH A*- $\bar{H}$, where $\bar{H}$ is a measure of the *hysteresis* in switching to a lower resolution

B. The IGH A* Framework

At the heart of Hybrid A* variants lies the interplay between *approximate dominance and activation*. Only dominant vertices in the forward search are ever allowed to *actively* participate in the search process, meaning dominance controls the tree expansion and therefore vertex discovery. However, outside of the forward search, this interplay can be broken. Our key insight is to generalize this framework by *decoupling vertex discovery and dominance*. This separation gives us a powerful new design tool: rather than relying solely on forward search to explore promising regions, we can *adaptively reuse, resurrect, or even inject* vertices from previous search efforts by activating them based on domain insight or resolution dynamics; a framework which we term IGH A* (Alg. 3).

For this decoupling, we introduce two modular mechanisms, SHIFT and ACTIVATE, and illustrate them through high-level control flow of IGH A* (Fig. 2(a)). ACTIVATE decides which vertices participate in the following iteration of the forward search. The planner activates vertices, enters the forward search, and expands vertices by order of priority while checking whether further search is warranted. SHIFT decides what the next search resolution will be and can continue the search or break the loop. Whenever the loop is broken, ACTIVATE again decides which vertices in Q_v will be activated for the next search iteration, and this goes on until $Q_v = \emptyset$. In the following sections (Sec. V-C, V-D), we illustrate how specific instantiations of SHIFT and ACTIVATE result in IGH A* variations, and use Fig. 3 to draw out the distinctions between them.

C. IGH A*- Multi-resolution search with information reuse

In this section, we dive into the details of our algorithm and design a straightforward but effective instantiation of IGH A* (Alg. 3). Here, during forward search (Line 5) dominant vertices are marked as active, non-dominant as inactive, and both are saved in Q_v (Lines 14, 18, and 19) for potential use in future resolutions with the invariant that only active vertices from Q_v can be expanded (Line 11). Note that $Q_v.\text{Active}$ refers to the partition of Q_v containing active vertices. The

Algorithm 3 Incremental Generalized Hybrid A* (IGHA*)

Input: Start vertex v_s , Goal set V_g , Resolution sequence $\mathcal{R}$
Output: Least-cost path or $\emptyset$ if FAILURE

```

1:  $Q_v \leftarrow \{v_s\}$ ,  $v_s.Active \leftarrow \text{true}$ 
2:  $\hat{\pi} \leftarrow \emptyset$ ,  $l \leftarrow 0$ ,  $l' \leftarrow 0$ 
3: while  $Q_v \neq \emptyset$  do  $\triangleright$  Iterations over the forward search
4:   ACTIVATE ()
5:   while  $Q_v.Active \neq \emptyset$  do
6:     if  $Q_v.Active.PEEK().f \geq w(\hat{\pi})$  then break
7:     if  $Q_v.Active.PEEK() \in V_g$  then
8:        $\hat{\pi} \leftarrow \text{EMITPATH}()$ 
9:       break
10:    if SHIFT() then break  $\triangleright$  Internally update  $l'$ 
11:     $u \leftarrow Q_v.Active.POP()$ 
12:    for all  $v \in \text{SUCC}(u)$  do
13:      if  $v.g < \hat{v}(v, R_l).g$  then
14:         $\hat{v}(v, R_l).Active \leftarrow \text{false}$ 
15:         $\hat{v}(v, R_l) \leftarrow v$ 
16:         $v.Active \leftarrow \text{true}$ 
17:      else
18:         $v.Active \leftarrow \text{false}$ 
19:       $Q_v.INSERT(v)$ 
20:    BOUND( $w(\hat{\pi})$ )  $\triangleright \forall v \in Q_v$  Remove  $v : v.f > w(\hat{\pi})$ 
21:    PROJECT( $R_l$ )  $\triangleright l \leftarrow l'$ , update  $R_l$ , and recompute  $\hat{v}$ 
22: return  $\hat{\pi}$ 

```

forward search concludes if either we've run out of vertices to expand, the branch and bound check fails, or a better path has been found (Lines 5, 6, 7 resp.). Then, we recheck the branch and bound, *project* all vertices to a finer resolution (Lines 20, 21 resp.), activate vertices newly dominant at this resolution (Fig. 2(b), ACTIVATE), and continue the search seamlessly like before at the finer resolution. Monotonically improving paths are emitted whenever they are discovered, as with HA*M, while requiring fewer expansions through vertex reuse.

D. IGH A* - $\bar{H}$ - Dynamically SHIFT resolution up and down

We have deliberately skipped over SHIFT until now; although monotonically increasing resolution is often a fine strategy, consider a scenario where a higher resolution was only needed to pass through a *single bottleneck* (SB) (Fig. 3(a)). Once vertices are through the bottleneck, we would observe dominance not only at the current but also lower resolutions. Staying at a higher resolution when a lower one would suffice results in an influx of vertices that would have been deactivated otherwise, creating more work. Instead of committing to the higher resolution (Fig. 3(b)), this information of *dominance at a lower resolution* can be used as an indicator to SHIFT to a lower resolution as soon as we pop an active vertex that indicates this (Fig. 3(d)) and save on expansions.

However, if there are *multiple bottlenecks* (MB) to pass through that the search process has not encountered yet (Fig. 3(f)), switching back at the first sign of low resolution dominance (Fig. 3(i)) can also work against us; we switch to a low resolution, cannot anymore fit through the next bottleneck, expand into the empty regions around the bottleneck, wasting effort, and once we've expanded as much as we can, finally switch to a higher resolution. Predicting imminent bottlenecks in general configuration spaces may not always be possible.

While we may have some intuition based on domain knowledge, we don't *exactly* know which situation we are in a priori. A safe bet is to incorporate some *hysteresis* until we've seen some evidence before switching to a lower resolution. This is the essence of IGH A* - $\bar{H}$'s SHIFT function (Fig. 2(b)) — the hysteresis threshold $\bar{H}$ allows us to interpolate between the two extremes of $\bar{H} = \infty$, which monotonically increases resolution, and $\bar{H} = 0$, which switches to a lower resolution as soon as it can (See Fig. 3(j), 3(e) and caption), effectively allowing us to encode a structural assumption into the algorithm through SHIFT and ACTIVATE. Note that we omit certain implementation details for simplicity, see *Implementation* in Sec. VIII.

E. Questions to address

Our work motivates the following key questions:

- 1) Are the SHIFT and ACTIVATE methods in Fig. 2(b) the only possible implementations? In Sec. VI we discuss the theoretical properties of ACTIVATE and SHIFT that ensure convergence.
- 2) Have we replaced one parameter (R in HA*) with another ($\bar{H}$ in IGH A* - $\bar{H}$)? We show in Sec. VII that although certain $\bar{H}$ outperform others, *all* of them significantly outperform HA*M. Thus, if one is allergic to hyperparameters, one can safely set $\bar{H} = \infty$ and ignore SHIFT.
- 3) How does IGH A* - $\bar{H}$ compare to HA*M in practice for off-road autonomy? We detail these experiments in Sec. VII-C).

VI. THEORETICAL ANALYSIS

Definition VI.1 (SHIFT). A SHIFT method dictates the search's resolution in the following iteration of the forward search. A SHIFT method can terminate the forward search by returning false after the first iteration and after at least one vertex from $Q_v.Active$ has been expanded in the current iteration. We term a SHIFT method *naive* if it only uses information in Q_v .

Definition VI.2 (ACTIVATE). An ACTIVATE method sets the active/inactive status of the vertices in Q_v after the first iteration of the forward search. We term an ACTIVATE method *naive* if it only uses information in Q_v .

Definition VI.3 (Rule). A combination of SHIFT and ACTIVATE methods is considered a rule. A rule R initializes the forward search of IGH A*, and results in the generation of the sub-trees $T_i^R = (V_i^R, E_i^R)$, $i \in \mathbb{Z}_{\geq 0}$, where i represents the iteration. Two rules A, B are considered *distinct* if for at least one planning query, for at least one iteration i , $T_i^A \neq T_i^B$. Combining a naive SHIFT and a naive ACTIVATE method results in a *naive* rule. A *practical* rule, for any iteration, does not activate the entire Q_v . A valid rule is such that $|Q_v.Active| \geq 1$ before starting forward search when $Q_v \neq \emptyset$.

Assumptions:

- A1 The heuristic h is admissible. I.e., $\forall v \in V$, $h(v) \leq h^*(v)$, where $h^*(v)$ is the lowest cost from v to V_g .
- A2 The max. branching factor at all vertices is b with $b < \infty$.
- A3 $w : E \rightarrow \mathbb{R}_{\geq \epsilon}$, $\epsilon > 0$.




Fig. 3: Here, the resolution is too low for HA* to find a solution for either the Single-Bottleneck (SB) scenario ((a)), or the Multi-Bottleneck (MB) scenario ((f)). In contrast to Fig. 1, for figures corresponding to IGHA*, translucent blue vertices are *deactivated*. IGHA* variants like IGHA*^{-∞} only increase resolution and spend more expansions after escaping a bottleneck ((b)) when a lower resolution could suffice IGHA*⁻⁰ ((d)), resulting in fewer expansions to the first path for IGHA*⁻⁰ ((e)). However, variants like IGHA*⁻⁰ can prematurely switch back in situations where staying at the higher resolution is necessary and spend extra expansions ((i)) by comparison to IGHA*^{-∞} ((g)). In both scenarios, an intermediate value for $\bar{H}$ can provide intermediate behavior; doing better than IGHA*^{-∞} in SB ((c)), and better than IGHA*⁻⁰ in MB ((h)). This interpolation is also visible through the rank plots corresponding to either scenario. In the **rank plot**, cells show the percentage of times (darker = higher) a planner achieved a rank for first-path; lower ranks (left) = faster solutions.

Observation VI.1. For a given planning query and starting resolution, HA* and IGHA* expand the same vertices at iteration 0 for any rule.

Observation VI.2. $w(\hat{\pi})$ strictly decreases with IGHA* iterations for any rule when assumption A1 holds.

Observation VI.3. From Def. VI.1, VI.3, at least one vertex is expanded each iteration as long as $Q_v \neq \emptyset$.

Theorem VI.1. When at least one solution to a planning query has been found, IGHA* will terminate with finite expansions when assumptions A1-A3 hold.

Proof: Suppose the first path is found at iteration i with cost $w(\hat{\pi})_i$. For any iteration $j > i$, given assumption A1 and the invariant $Q_v.\text{Active.PEEK}().f < w(\hat{\pi})_i$, let $V' : \{v \in V | v.g < w(\hat{\pi})_i\}$ represent the set of all vertices that Q_v could have. Given A2, A3, the maximum possible depth (from v_s) is $d < w(\hat{\pi})_i/\epsilon$ and so $|V'| \leq O(b^d)$, meaning $|V'|$ is upper bounded by a finite number. Following Obs. VI.3, the algorithm must terminate within finite expansions. $\square$

Observation VI.4. From Obs. VI.2, VI.3, and Thm. VI.1, when at least one solution to a planning query has been found, and assumptions A1-A3 hold, IGHA* will find the best-cost path in $T = (V, E)$ at or before termination.

Discretization can cause different trees to be generated depending on which vertices in Q_v are expanded. No naive rule reliably finds the first or best path with minimal expansions, as Q_v lacks information about which subset yields which tree.


Theorem VI.2. For any *naive* and *practical* rule A there exists another *distinct* naive and practical rule B for which there exists a scenario where B finds the first solution with fewer expansions than A .

Proof: Throughout the proof and with a slight abuse of notation, we denote IGHA*- R as the instantiation of IGHA* using rule R . We use ξ_i^R to represent the number of expansions for IGHA*- R at iteration i , and set $\Xi_i^R = \sum_{k=0}^i \xi_k^R$. Finally, we use $\xi_i^R(u, v)$, $u, v \in V_i^R$ to represent expansions used to reach v from u at iteration i . Given some scenario, we call vertex $v_b \in V$ a *bottleneck vertex* if v_b belongs to every path between v_s and some vertex in V_g and that v_b is one expansion away from V_g . Now consider a scenario containing a bottleneck vertex v_b where IGHA*- A does not find the solution at iterations 0 and 1 and does not unfreeze v_b – possible as Q_v does not contain information about which vertex will lead to V_g . As V_g is one expansion away from v_b , we have that $\xi^R(v_b, V_g) = 1$ for any rule R at any iteration. Define rule B using the same SHIFT method as rule A but an ACTIVATE method that activates only $v_b \in Q_v^A, v_b \notin Q_v^A.\text{Active}$ at the beginning of iteration 1. Such a rule is possible as rule A is naive and practical, and B is distinct from A .

At iteration 1, IGHA*- B finds a solution, $\Xi_1^B = \xi_0^B + 1$. At iteration $j > 1$, IGHA*- A finds a solution. $\Xi_j^A = \xi_0^A + \xi_1^A \dots \xi_j^A(v_s, v_b) + 1$. Following Obs. VI.1, and VI.3 we have that $\xi_0^A = \xi_0^B$, and that $\xi_1^A + \dots \xi_j^A(v_s, v_b) \geq 1$. Therefore $\Xi_j^A \geq \Xi_1^B + 1$ and thus $\Xi_j^A > \Xi_1^B$. $\square$

Corollary VI.1. For any *naive* and *practical* rule A there exists another *distinct* naive and practical rule B for which there exists a scenario where B finds the best solution with fewer expansions than A .

Proof: In Thm. VI.2, it is possible to have a scenario where the first solution is the best solution. Then, using the same mechanism as in Thm. VI.2, the Corollary can be proven. $\square$



(a) SB IGHA^{*}-∞ (b) SB IGHA^{*}-0 (c) MB IGHA^{*}-∞ (d) MB IGHA^{*}-0
Fig. 4: In some queries of the SB environment, IGHA^{*}-0 can have more expansions than IGHA^{*}-∞ (Fig. 3(e)). Examples (a) and (b) show that this happens when the low resolution of IGHA^{*}-0 prevents it from finding a straight-line path to the goal. In some queries of the SB environment, IGHA^{*}-∞ can have more expansions than IGHA^{*}-0 (j)). Examples (c) and (d) show that this can happen when the tree produced by IGHA^{*}-0 contains a path at a lower resolution that passes through the two other gaps.

VII. EXPERIMENTS AND RESULTS

In this section, IGHA^{*} and IGHA^{*}- $\bar{H}$ (Sec. V-D) are used interchangeably. We use HA^{*}M as a baseline (Sec. V-A). We use 95% confidence intervals.

A. $\mathbb{R}^2$ Environments

To show that the $\bar{H}$ discussed in Sec. V-D encodes a structural assumption into IGHA^{*}, we evaluate on 500 automatically generated $\mathbb{R}^2$ environments, with Fig. 3(a) showing one of the queries for the *Single-Bottleneck* environment (SB), and Fig. 3(f) showing one from the *Multi-Bottleneck* environments (MB). We vary the location of the gaps (whitespace) in the walls (black) and vary v_s and V_g in the top left and bottom right corners of the map, respectively. The planning queries are generated such that HA^{*} would not find the solution at the starting resolution. We compare IGHA^{*}-∞ against IGHA^{*}-0. We evaluate the total expansions needed to find the first path.

Hypotheses:

H1: In SB, IGHA^{*}-0 finds the first solution with fewer expansions than IGHA^{*}-∞ more often than not.

H2: In MB, IGHA^{*}-∞ finds the first solution with fewer expansions than IGHA^{*}-0 more often than not.


Comparison	Ratio	Speed Up
MB IGHA [*] -∞ / IGHA [*] -0	0.70 ± 0.04	3.52 ± 0.26
SB IGHA [*] -0 / IGHA [*] -∞	0.84 ± 0.03	1.92 ± 0.07
Kinematic IGHA [*] -∞ / IGHA [*] -0	0.67 ± 0.06	4.10 ± 0.63
Kinodynamic IGHA [*] -0 / IGHA [*] -∞	0.60 ± 0.06	3.17 ± 0.52

TABLE I: The **Ratio** is the ratio of the total queries in which A beats B (A/B) to the first solution. The **Speed Up** shows how much faster A is relative to B when A beats B.

Results: Table I shows data in the first and second rows that **confirms hypotheses H1 and H2**, showing that when one beats the other, there is a 2 – 3× speedup. We have shown examples of these scenarios in Fig. 3, showing additionally how intermediate values of $\bar{H}$ operate, and we show counterexamples and how they come up in Fig. 4.

B. $\mathbb{R}^3, \mathbb{R}^4$ Environments

Here, we intend to show that the best path performance of IGHA^{*} variants is not hypersensitive to the choice of the hysteresis threshold hyperparameter. We consider two types of environments: urban city maps (from the moving AI benchmark [16], example shown in Fig. 5(a)) with a kinematic car [36] ($\mathbb{R}^3$), and off-road maps (from BeamNG [17], example



(a) Urban environments [16] (b) Off road environments [17]
Fig. 5: Notice that in Table I, IGHA^{*}-∞ beats IGHA^{*}-0 in the kinematic environments and IGHA^{*}-0 beats IGHA^{*}-∞ in the off-road/kinodynamic environments. It goes to show that the urban city environments ((a)) exhibit more clutter, similar to MB, and off-road environments ((b)) appear sparser as in SB.




Fig. 6: When HA^{*}M terminates before reaching the expansion limit, IGHA^{*} variants generally terminate faster. The solid lines represent the speed-up in terms of expansions, whereas the dotted lines represent the wall-clock speed-up for our C++/CUDA implementation. Observe that the speed up in expansions (solid lines) is not hypersensitive to the hysteresis value. Expansions take $\approx 50\mu s$ on a desktop (Intel i9-12900k CPU/Nvidia RTX3080 GPU), whereas projecting each vertex takes $\approx 0.3\mu s$. When $\bar{H}$ is small, switchbacks occur more frequently, leading to more time spent projecting and activating vertices.

shown in Fig. 5(b)) with a kinodynamic car [37], assuming no side-slip, with its state ($\mathbb{R}^4$) defined by its SE(2) pose and body frame forward velocity. The other components of the state (roll, pitch, and so on) are unique for a given SE(2) pose, velocity, and elevation map. We generate 500 queries for both environments with the hyperparameters (related to successor, goal set, starting resolution) set such that HA^{*} can find a solution at the starting resolution to reflect a practical setting. We use variants of IGHA^{*} with the $\bar{H}$ drawn from $[0, 10, 50, 250, 1250, \infty]$ and compare them to HA^{*}M. We use $\bar{H} = 0, \infty$ as references, and chose the intermediate values for $\bar{H}$ based on the average expansions used by HA^{*}M to find the first solution for both environments, which is ~ 1500 . All methods are limited to 100,000 expansions for both environments. Note that here we are running the planner *open loop*.

Hypothesis H3: IGHA^{*}-n variants use fewer expansions to terminate on average compared to HA^{*}M, when HA^{*}M terminates before the expansion limit.

Result: Fig. 6 confirms H3 for both the kinematic and kinodynamic environments, showing that in general the IGHA^{*}-n variants are $\sim 6\times$ faster than HA^{*}M. See Table I and the caption of Fig. 5 for observations on the structure in the kinematic and kinodynamic off-road planning problems.

C. Closed Loop Experiments for Off-Road Autonomy

We run the kinodynamic planner in the loop with MPPI [38], using a real-time C++/CUDA implementation with GPU-




Fig. 7: Across off-road scenarios (Fig. 5(b)), IGHA* enables faster goal-reaching, outperforming HA*M, MPC, SST, and KPIECE ((a)). Representative example shown in (b), where the heightmap is grayscale (lighter = higher), with purple as obstacles, green as the plan (brighter = faster), black for MPC rollouts, and a white circle marking V_g . IGHA* (top left) yields lower-cost plans than HA*M, SST, and KPIECE, (top center, top right, bottom left resp.). Both IGHA* and HA*M outperform MPC (bottom center) in avoiding local minima (global snapshot shown bottom right). Here, IGHA* uses $\approx 0.8\text{MB}$ (c), 200 samples for the setup corresponding to (a), well within the capacity of embedded systems such as the Jetson Orin NX 16GB. Real-world tests ((d)) in an outdoor dirt bike park demonstrate online planning on a small-scale platform (left) at 3–4 Hz under tight computational limits, using a $16 \times 16\text{m}^2$ body-centric map (right).

accelerated edge expansion and evaluation. Fig. 6 shows that the wall-clock speedup (dotted grey lines) peaks near the middle, which aids in tuning for the hysteresis value. To avoid perception and state estimation confounds, we run experiments in BeamNG [17], accessing ground-truth state and body-centric maps over a $100 \times 100\text{m}^2$ area at 25 Hz. We compare IGHA*-100 (or IGHA*), HA*M, SST [12] and KPIECE [39] from OMPL [40], and a baseline MPC. The MPC uses a 1.5s horizon with 1024 rollouts when used with the planner, but 5s with 8192 rollouts when run as a baseline for it to stay competitive. We tune the resolution for SST to avoid overpruning at standstill, using the same resolution for KPIECE for consistency, and default values for other hyperparameters. Hyperparameters for map size, MPC horizon, control frequency, and so on follow off-road autonomy norms [3], [4], with a planning budget of 5000 expansions, replanning at ≈ 4 Hz. We evaluate 30 start-goal pairs ($\approx 250\text{m}$ apart), each run 4 times with a 120s timeout; success requires reaching the goal within time. If the goal lies outside the map, it’s projected to the edge.

Hypothesis H4: IGHA* has a higher success rate compared to the baselines.

Result: Fig. 7(a) shows that IGHA* generally achieves higher success rates than the baselines, and reaches the same success rates within a much shorter time frame. Fig. 7(b) illustrates this: IGHA* yields lower-cost paths than the baselines under a fixed budget.

Deployment: The kinodynamic planner runs at $\approx 20\text{k}$, 4k, and 2k expansions/sec on our desktop, Nvidia Orin AGX (32GB), and Orin NX (16GB), respectively, yielding a ≈ 4 Hz planning rate in simulation on desktop, (vertex-mem. req. shown in Fig 7(c)). We deploy on a small-scale off-road autonomy platform with all modules (MPC, perception, planning) running onboard on an Orin NX at ≈ 2 m/s in real off-road terrain (Fig. 7(d); see [41] for details about the platform, MPC, and the perception system). On-robot, we target a 3–4 Hz update rate with a budget of 500 expansions (actual: $4.4 \pm 0.3\text{Hz}$, variations due to intermittent resource contention and early termination by the planner).

VIII. DISCUSSION

IGHA* represents a shift from a purely expansion-driven planner to a more general *vertex reactivation framework*, capable of leveraging both prior computation and resolution-aware

reasoning. As long as the implementation of alternative SHIFT and ACTIVATE methods fit within the definitions in Sec. VI, the guarantees hold. Notice that HA*M can be constructed using IGHA*, by using a SHIFT method that always increases resolution but never breaks the loop, and an ACTIVATE method that resets Q_v with v_s . While a large number of variations of IGHA* are possible, the focus of this work is to show the value that even a straightforward variant can bring. Further, note that in our deployment, we make no domain-specific modifications to Alg. 3; deploying the algorithm on a different robot, or for higher-dimensional environments, is a matter of choosing the appropriate successor, heuristic, and goal set definitions – computational burden, due to the curse of dimensionality in higher dimensional problems, is expected to change across the board.

Remark: in the instantiation in Sec. V-D, the rule used is *naive*; it only depends on information directly available in Q_v ; the $\bar{H}$ we use for our closed loop experiments (Sec. VII-C) is informed by our domain knowledge about off-road autonomy. However, for different domains, one could learn the appropriate *average $\bar{H}$* values or predict them on the fly, while technically still retaining the guarantees in Sec. VI. While we use a grid-based discretization approach to perform duplicate detection, other methods of discretization, such as use of r-disks as in [12], [13] or successor validity as in [23], [24] for duplicate detection may also be used as long as they satisfy the assumptions in Sec. VI and expose a hyperparameter equivalent to the discretization resolution used here.

Implementation: The pseudocode shown here omits certain mechanisms for simplicity. We only maintain Q_v . Active as a queue, and the rest as an unordered list so as to not incur the cost of sorting inactive vertices at every step of forward search. We handle the edge cases that prevent SHIFT from being called at least once each iteration. The vertex attributes, such as g , f values, and the lowest dominant resolution are set implicitly during vertex expansion.

REFERENCES

- [1] D. Dolgov, S. Thrun, M. Montemerlo, and J. Diebel, “Path Planning for Autonomous Vehicles in Unknown Semi-Structured Environments,” *International Journal of Robotics Research*, vol. 29, no. 5, pp. 485–501, 2010.
- [2] J. Tordesillas, B. T. Lopez, J. Carter, J. Ware, and J. P. How, “Real-time Planning with Multi-fidelity Models for Agile Flights in unknown Environments,” in *IEEE International Conference on Robotics and Automation*. IEEE, 2019, pp. 725–731.

- [3] T. Han, A. Liu, A. Li, A. Spitzer, G. Shi, and B. Boots, "Model Predictive Control for Aggressive Driving Over Uneven Terrain," *Robotics: Science and Systems*, 2023.
- [4] X. Meng, N. Hatch, A. Lambert, A. Li, N. Wagener, M. Schmittle, J. Lee, W. Yuan, Z. Chen, S. Deng, G. Okopal, D. Fox, B. Boots, and A. Shaban, "Terrainnet: Visual Modeling of Complex Terrain for High-speed, Off-road Navigation," in *Robotics: Science and Systems*, 2023.
- [5] R. A. Knepper and A. Kelly, "High Performance State Lattice Planning Using Heuristic Look-up Tables," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Citeseer, 2006, pp. 3375–3380.
- [6] M. Pivtoraiko and A. Kelly, "Efficient Constrained Path Planning via Search in State Lattices," in *International Symposium on Artificial Intelligence, Robotics, and Automation in Space*. Munich Germany, 2005, pp. 1–7.
- [7] D. Ferguson and A. Stentz, "Field D*: An Interpolation-based Path Planner and Replanner," in *International Symposium on Robotics Research*. Springer, 2007, pp. 239–253.
- [8] B. Hedegaard, E. Fahnestock, J. Arkin, A. Menon, and T. M. Howard, "Discrete Optimization of Adaptive State Lattices for Iterative Motion Planning on Unmanned Ground Vehicles," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2021, pp. 5764–5771.
- [9] R. Bellman, "Dynamic Programming," *Science*, vol. 153, no. 3731, pp. 34–37, 1966.
- [10] H. Lee, T. Kim, J. Mun, and W. Lee, "Learning Terrain-Aware Kinodynamic Model for Autonomous Off-road Rally Driving with Model Predictive Path Integral Control," *IEEE Robotics and Automation Letters*, 2023.
- [11] E. R. Damm, J. M. Gregory, E. S. Lancaster, F. A. Sanchez, D. M. Sahu, and T. M. Howard, "Terrain-Aware Kinodynamic Planning with Efficiently Adaptive State Lattices for Mobile Robot Navigation in Off-road Environments," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2023, pp. 9918–9925.
- [12] Y. Li, Z. Littlefield, and K. E. Bekris, "Sparse Methods for Efficient Asymptotically Optimal Kinodynamic Planning," in *Algorithmic Foundations of Robotics XI: Selected Contributions of the Eleventh International Workshop on the Algorithmic Foundations of Robotics*. Springer, 2015, pp. 263–282.
- [13] —, "Asymptotically Optimal Sampling-based Kinodynamic Planning," *International Journal of Robotics Research*, vol. 35, no. 5, pp. 528–564, 2016.
- [14] K. Kurzer, "Path Planning in Unstructured Environments : A Real-time Hybrid A* Implementation for Fast and deterministic Path Generation for the Kth Research Concept Vehicle," 2016.
- [15] W. Sheng, B. Li, and X. Zhong, "Autonomous Parking Trajectory Planning with tiny Passages : A Combination of Multistage Hybrid A* Algorithm and Numerical Optimal Control," *IEEE Access*, vol. 9, pp. 102 801–102 810, 2021.
- [16] N. Sturtevant, "Benchmarks for Grid-Based Pathfinding," *Transactions on Computational Intelligence and AI in Games*, vol. 4, no. 2, pp. 144 – 148, 2012. [Online]. Available: <http://web.cs.du.edu/~sturtevant/papers/benchmarks.pdf>
- [17] BeamNG GmbH, "BeamNG.tech," 2022. [Online]. Available: <https://www.beamng.tech/>
- [18] A. Thoesen and H. Marvi, "Planetary Surface Mobility and Exploration: A Review," *Current Robotics Reports*, vol. 2, no. 3, pp. 239–249, 2021.
- [19] S. Karaman and E. Frazzoli, "Sampling-based Algorithms for Optimal Motion Planning," *International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.
- [20] Y. Ji, Y. Tanaka, Y. Tamura, M. Kimura, A. Umemura, Y. Kaneshima, H. Murakami, A. Yamashita, and H. Asama, "Adaptive Motion Planning Based on Vehicle Characteristics and Regulations for Off-road UGVs," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 1, pp. 599–611, 2018.
- [21] D. J. Webb and J. Van Den Berg, "Kinodynamic RRT*: Asymptotically Optimal Motion Planning for Robots with Linear Dynamics," in *IEEE International Conference on Robotics and Automation*. IEEE, 2013, pp. 5054–5061.
- [22] J. Barraquand and J.-C. Latombe, "Nonholonomic Multibody Mobile Robots : Controllability and Motion Planning in the Presence of Obstacles," *Algorithmica*, vol. 10, no. 2, pp. 121–155, 1993.
- [23] W. Du, S.-K. Kim, O. Salzman, and M. Likhachev, "Escaping Local Minima in Search-Based Planning using Soft Duplicate Detection," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019, pp. 2365–2371.
- [24] M. Fu, K. Solovey, O. Salzman, and R. Alterovitz, "Toward Certifiable Optimal Motion Planning for Medical Steerable Needles," *International Journal of Robotics Research*, vol. 42, no. 10, pp. 798–826, 2023.
- [25] M. Likhachev, G. J. Gordon, and S. Thrun, "ARA*: Anytime A* with Provable Bounds on Sub-optimality," *Advances in Neural Information Processing Systems*, vol. 16, 2003.
- [26] J. Van Den Berg, R. Shah, A. Huang, and K. Goldberg, "Anytime Nonparametric A*," in *AAAI Conference on Artificial Intelligence*, vol. 25, no. 1, 2011, pp. 105–111.
- [27] H. Zhang, O. Salzman, A. Felner, C. H. Ulloa, and S. Koenig, "A-A*pex: Efficient anytime approximate multi-objective search," in *International Symposium on Combinatorial Search*. AAAI Press, 2024, pp. 179–187.
- [28] W. Du, F. Islam, and M. Likhachev, "Multi-Resolution A*," in *Proceedings of the International Symposium on Combinatorial Search*, vol. 11, no. 1, 2020, pp. 29–37.
- [29] D. M. Saxena, T. Kusnur, and M. Likhachev, "AMRA*: Anytime Multi-Resolution Multi-heuristic A*," in *IEEE International Conference on Robotics and Automation*. IEEE, 2022, pp. 3371–3377.
- [30] A. Yahja, A. Stentz, S. Singh, and B. L. Brumitt, "Framed-Quadtree Path Planning for Mobile Robots Operating in Sparse Environments," in *IEEE International Conference on Robotics and Automation*, vol. 1. IEEE, 1998, pp. 650–655.
- [31] F. M. Garcia, M. Kapadia, and N. I. Badler, "GPU-based Dynamic Search on Adaptive Resolution Grids," in *IEEE International Conference on Robotics and Automation*. IEEE, 2014, pp. 1631–1638.
- [32] J. D. Gammell, T. D. Barfoot, and S. S. Srinivasa, "Batch Informed Trees (BIT*): Informed Asymptotically Optimal Anytime Search," *International Journal of Robotics Research*, vol. 39, no. 5, pp. 543–567, 2020.
- [33] R. Scalise, A. Mandalika, B. Hou, S. Choudhury, and S. S. Srinivasa, "GuILD : Guided Incremental local Densification for Accelerated Sampling-Based Motion Planning," in *IEEE International Conference on Robotics and Automation*. IEEE, 2023, pp. 10212–10218.
- [34] D. Yi, R. Thakker, C. Gulino, O. Salzman, and S. S. Srinivasa, "Generalizing informed sampling for asymptotically-optimal sampling-based kinodynamic planning via markov chain monte carlo," in *IEEE International Conference on Robotics and Automation*. IEEE, 2018, pp. 7063–7070.
- [35] S. Choudhury, O. Salzman, S. Choudhury, and S. S. Srinivasa, "Densification Strategies for Anytime Motion Planning over large Dense Roadmaps," in *IEEE International Conference on Robotics and Automation*. IEEE, 2017, pp. 3770–3777.
- [36] J. Kong, M. Pfeiffer, G. Schildbach, and F. Borrelli, "Kinematic and Dynamic Vehicle Models for Autonomous Driving Control Design," in *2015 IEEE Intelligent Vehicles Symposium*. IEEE, 2015, pp. 1094–1099.
- [37] T. Fork, H. E. Tseng, and F. Borrelli, "Models for Ground Vehicle Control on Nonplanar Surfaces," *Vehicle System Dynamics*, pp. 1–25, 2023.
- [38] G. Williams, A. Aldrich, and E. A. Theodorou, "Model Predictive Path Integral Control: From Theory to Parallel Computation," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [39] I. A. Sucan and L. E. Kavraki, "A Sampling-Based Tree Planner for Systems with Complex Dynamics," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 116–131, 2011.
- [40] 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, December 2012, <https://ompl.kavrakilab.org>.
- [41] S. Talia, M. Schmittle, A. Lambert, A. Spitzer, C. Mavrogiannis, and S. S. Srinivasa, "Demonstrating Hound: A Low-cost Research Platform for High-speed Off-road Underactuated Nonholonomic Driving," *Robotics: Science and Systems*, 2024.