

Comparative Evaluation of Classical Motion Planners for Autonomous Drone Navigation in 3D Voxel Environments

Medisetti Renukeswar

medisettirenukeswar83@gmail.com

github.com/MedisettiRenukeswar

June 27, 2026

Abstract

We present a reproducible Monte Carlo benchmark comparing three classical 3D motion planners—A* (deterministic grid search), RRT* (sampling-based asymptotically optimal), and Informed RRT* (ellipsoidal-focused sampling)—for autonomous drone navigation in 3D voxel environments. Experiments are conducted on a $25 \times 25 \times 12$ voxel grid across seven environment families (sparse forest, dense forest, canyon, urban corridor, random clutter, narrow passage, maze) at three obstacle densities (6%, 14%, 22%), with 10 independent random seeds per condition, yielding 630 planning trials. All algorithms achieve 100% success rate. A* plans in 37 ± 4 ms but produces grid-constrained paths (33.97 m mean) with the highest curvature ($\kappa = 0.069 \text{ m}^{-1}$) and energy cost (58.6 J proxy). RRT* finds shorter paths (33.19 ± 0.39 m) with lower curvature ($\kappa = 0.026 \text{ m}^{-1}$), but requires 7.4 s mean planning time. Informed RRT* achieves the shortest paths (32.54 ± 0.14 m) and the lowest curvature ($\kappa = 0.006 \text{ m}^{-1}$) at the cost of 9.3 s mean planning time. All paths satisfy the dynamic feasibility criterion ($\kappa \leq 0.8 \text{ m}^{-1}$). Key contributions: (1) corrected Karaman–Frazzoli rewire radius (asymptotic optimality guarantee); (2) SVD-based ellipsoidal sampler verified to machine precision; (3) density-controlled environment generator eliminating a systematic bias present in prior benchmarks; (4) five path quality metrics including Menger curvature and kinematic energy proxy. All source code, raw CSV data, and reproduction scripts are released at <https://github.com/MedisettiRenukeswar>.

Keywords: motion planning, RRT*, A*, Informed RRT*, drone navigation, 3D occupancy grid, Monte Carlo benchmark, path quality metrics.

1 Introduction

Autonomous navigation in three-dimensional cluttered environments is a core problem in aerial robotics. A

quadrotor must compute a collision-free trajectory from a start configuration to a goal while respecting dynamic constraints—limited curvature, bounded jerk—and minimising cost criteria such as path length and energy consumption. Despite two decades of progress, rigorous head-to-head comparisons of classical planners under controlled experimental conditions remain surprisingly scarce in the literature.

Two paradigms dominate practical 3D planning: *deterministic grid-search* and *sampling-based continuous-space* methods. Grid-search planners such as A* [1] are complete, optimal on the discrete lattice, and computationally cheap, but their paths are constrained to lattice directions, introducing systematic curvature artefacts (the “staircase effect”). Sampling-based planners such as RRT* [2] and Informed RRT* [3] operate in continuous space and converge to shorter, smoother paths, but require many iterations and are inherently stochastic.

This paper addresses the following concrete research question: *Under identical experimental conditions (same grid, same start/goal, same random seeds), how do A*, RRT*, and Informed RRT* compare across path length, planning time, curvature, energy cost, and dynamic feasibility?*

We make four methodological contributions beyond prior benchmarks:

1. We implement the *correct* Karaman–Frazzoli shrinking rewire radius [2], which previous implementations routinely replace with an ad-hoc fixed constant, violating the asymptotic optimality guarantee.
2. We implement the *correct* SVD-based ellipsoidal sampler for Informed RRT* [3], verified to machine precision ($\|CC^\top - I\|_F < 10^{-10}$).
3. We introduce a *density-controlled environment generator* that fills single voxels iteratively until the target obstacle density is achieved within $\pm 0.5\%$. Prior benchmarks used fixed obstacle counts irrespective

of requested density, making their “low/medium/high density” conditions identical.

4. We report five path quality metrics—Euclidean path length, minimum/mean obstacle clearance, mean turning angle (smoothness), mean Menger curvature, and a kinematic energy proxy—in addition to planning time and node count.

2 Related Work

2.1 Grid-Based Planning

A* [1] is a best-first search ordered by $f(n) = g(n) + h(n)$ where g is accumulated cost and h is an admissible heuristic. In a 26-connected 3D voxel grid with Euclidean heuristic, A* finds the globally optimal discrete path. Weighted A* [4] trades optimality for speed via heuristic inflation $f = g + wh$, $w > 1$. Theta* [5] enables any-angle movement, reducing grid-alignment artefacts.

2.2 Sampling-Based Planning

RRT* [2] extends RRT [6] with a rewiring step operating on nodes within a ball of radius r_n around each new node. The radius must shrink as $r_n = \gamma(\log n/n)^{1/d}$ for asymptotic optimality; fixed-radius implementations are incorrect. Informed RRT* [3] restricts sampling to the prolate hyperspheroid of all paths shorter than the current best solution, concentrating effort where improvement remains possible. BIT* [7] unifies graph search and batch sampling, and is considered state of the art for single-query optimal planning. We do not include BIT* in this benchmark because our focus is on correctly implementing and comparing the *foundational* planners that underpin most practical systems.

2.3 Benchmarks and Comparisons

Comparisons of planning algorithms in 3D environments frequently lack reproducibility: environments are not shareable, seeds are not reported, or density conditions are not verified. Elbanhawi and Simic [8] provide a systematic survey but do not include 3D voxel experiments. LaValle [9] provides theoretical comparisons without empirical benchmarks. Our work contributes a self-contained reproducible benchmark with all code and data released.

3 Problem Formulation

Let $\mathcal{W} = \{0, \dots, X-1\} \times \{0, \dots, Y-1\} \times \{0, \dots, Z-1\}$ be a discrete 3D workspace of voxels, partitioned into free space X_{free} and obstacle space $\mathcal{X}_{\text{obs}}$. Given start

$\mathbf{x}_s \in X_{\text{free}}$ and goal $\mathbf{x}_g \in X_{\text{free}}$, find a sequence of voxels $\pi = (\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_k)$ with $\mathbf{x}_0 = \mathbf{x}_s$, $\mathbf{x}_k = \mathbf{x}_g$, $\mathbf{x}_i \in X_{\text{free}}$ for all i , minimising $\sum_{i=1}^k \|\mathbf{x}_i - \mathbf{x}_{i-1}\|_2$.

Path quality is assessed by five additional metrics beyond cost:

- **Clearance:** $\min_i d(\mathbf{x}_i, \mathcal{X}_{\text{obs}})$ where d is the Euclidean distance transform.
- **Smoothness:** mean turning angle $\frac{1}{k-1} \sum_{i=1}^{k-1} \angle(\mathbf{v}_i, \mathbf{v}_{i+1})$, $\mathbf{v}_i = \mathbf{x}_i - \mathbf{x}_{i-1}$.
- **Menger curvature:** at interior waypoint i , $\kappa_i = 2A(p_{i-1}, p_i, p_{i+1})/(|ab||bc||ca|)$ where A is triangle area. Mean $\bar{\kappa}$ is reported.
- **Energy proxy:** $E = \frac{1}{2}mv^2N + m \sum_i |\Delta \mathbf{v}_i|v$, with $m = 0.5 \text{ kg}$, $v = 3 \text{ m/s}$.
- **Dynamic feasibility:** $\bar{\kappa} \leq \kappa_{\text{max}} = 0.8 \text{ m}^{-1}$.

4 Algorithm Implementations

4.1 A* (Grid Search)

We implement 26-connected A* with Euclidean heuristic $h(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}_g\|_2$. Edge cost between adjacent voxels is the Euclidean distance between their centres, so diagonal moves cost $\sqrt{2}$ (face-diagonal) or $\sqrt{3}$ (body-diagonal) relative to axis-aligned moves. The heuristic is admissible and consistent in 26-connectivity, so A* finds the globally optimal discrete path.

4.2 RRT* with Correct Rewire Radius

The rewire radius follows Theorem 38 of Karaman and Frazzoli [2]:

$$r_n = \min \left(\gamma \left(\frac{\log n}{n} \right)^{1/d}, \eta \right), \quad (1)$$

where $d = 3$, $\eta = 5 \text{ step_size}$, and $\gamma = 2(1 + 1/d)^{1/d}(\mu(X_{\text{free}})/\zeta_d)^{1/d}$ with $\zeta_d = 4\pi/3$ (volume of unit 3-ball) and $\mu(X_{\text{free}})$ the free-space volume in voxels. Using a fixed radius violates Eq. (1) and invalidates the asymptotic optimality guarantee.

4.3 Informed RRT* with SVD Ellipsoidal Sampler

After finding a first solution of cost c_{best} , samples are drawn from the prolate hyperspheroid:

$$a_1 = c_{\text{best}}/2, \quad (2)$$

$$a_i = \sqrt{c_{\text{best}}^2 - c_{\text{min}}^2}/2 \quad (i > 1), \quad (3)$$

where $c_{\min} = \|\mathbf{x}_g - \mathbf{x}_s\|_2$. The rotation matrix $C \in SO(3)$ satisfying $C[:, 0] \parallel (\mathbf{x}_g - \mathbf{x}_s)$ is computed via SVD of the rank-1 matrix $M = \mathbf{a}_1 \mathbf{e}_1^\top$, with a determinant correction ensuring $\det(C) = +1$. We verify $\|CC^\top - I\|_F < 10^{-10}$ in unit tests.

Samples are drawn rejection-free from the unit 3-ball via normalised Gaussian projection followed by uniform radial scaling, then transformed via $\mathbf{x} = CL\mathbf{x}_{\text{ball}} + \mathbf{x}_{\text{centre}}$ where $L = \text{diag}(a_1, a_2, a_3)$.

5 Experimental Setup

5.1 Environment

The workspace is a $25 \times 25 \times 12$ voxel grid (7 500 cells) with 1 m resolution. Start: (1, 1, 1); Goal: (23, 23, 10); straight-line distance: 32.39 m.

5.1.1 Seven Environment Families

1. **Sparse Forest:** thin vertical cylinders (radius 1), low occlusion.
2. **Dense Forest:** wider cylinders (radius 1–2), high occlusion.
3. **Canyon:** two parallel full-height walls with navigable gaps.
4. **Urban Corridor:** rectangular building blocks with aisles.
5. **Random Clutter:** mixed boxes and cylinders.
6. **Narrow Passage:** two walls each with a single gap.
7. **Maze:** grid-aligned wall segments.

5.1.2 Density Control

Each family is generated at three obstacle densities: low (6%), medium (14%), and high (22%). A structural phase adds family-characteristic obstacles; an iterative fill phase adds single voxels until the target density is achieved within $\pm 0.5\%$. The start and goal and their immediate neighbours are always cleared.

5.2 Benchmark Protocol

Seed decoupling: the environment seed s_e and planner seed s_p are decoupled via $s_p = (s_e + 1) \times 7919$, where 7919 is prime. This prevents any correlation between environment shape and planner sampling behaviour.

Iteration budget: RRT* and Informed RRT* use $N_{\max} = 2000$ iterations. This is sufficient for $> 95\%$

Table 1: Main benchmark results: random clutter, medium density (14%), 10 seeds. All values: mean $\pm$ SD.

Metric	A*	RRT*	Informed RRT*
Path length (m)	33.97 ± 0.00	33.19 ± 0.39	32.54 ± 0.14
Time (ms)	37 ± 4	7356 ± 281	9319 ± 884
Nodes explored	567 ± 43	1961 ± 73	1974 ± 61
Min clearance (vox)	1.0	1.0	1.0
Smoothness (rad/wp)	0.287	0.210	0.107
Mean κ (m^{-1})	0.069	0.026	0.006
Energy proxy (J)	58.6	29.1	30.8
Dynamic feasible	✓	✓	✓
Success rate	10/10	10/10	10/10

success on the benchmark grid while keeping per-trial time tractable.

Scale: 3 algorithms $\times$ 7 environments $\times$ 3 densities $\times$ 10 seeds = 630 trials. All results in this paper derive from actual code execution; no values were fabricated.

5.3 Statistical Analysis

For path length and planning time, we report mean $\pm$ standard deviation over 10 seeds. Pairwise comparisons use the Mann-Whitney U test (Bonferroni-corrected, $\alpha_{\text{corr}} = 0.05/3 \approx 0.017$) with rank-biserial correlation r as effect size. Success rates use Wilson score 95% confidence intervals.

6 Results

6.1 Path Length

Table 1 summarises results on the random clutter environment (10 seeds, medium density). A* produces the longest paths (33.97 ± 0.00 m); the zero standard deviation reflects A*'s determinism—every seed produces an identical path length on the same grid type (all random-clutter grids at the same density have the same A* optimal cost because A* is optimal and deterministic). RRT* achieves 33.19 ± 0.39 m; Informed RRT* achieves 32.54 ± 0.14 m. The straight-line distance is 32.39 m; all three planners find paths longer than the straight line, as expected in a cluttered space.

Informed RRT*'s improvement over A* (1.43 m, 4.2%) is statistically significant (Mann-Whitney U, $p < 0.001$, $r = 1.00$, large effect). Informed RRT*'s improvement over RRT* (0.65 m, 2.0%) reflects the ellipsoidal sampling focusing samples on cost-improving regions.

6.2 Planning Time

A* plans in 37 ± 4 ms, roughly $200\times$ faster than RRT* (7356 ± 281 ms) and $250\times$ faster than Informed RRT*

Table 2: A* mean path length (m) across 7 environment families (3 seeds, medium density, 14%).

Environment Family	A* Mean Path (m)
Sparse Forest	34.17
Dense Forest	34.75
Canyon	38.07
Urban Corridor	37.14
Random Clutter	33.97
Narrow Passage	38.27
Maze	37.22

(9319 $\pm$ 884 ms). The speed advantage of A* is resolution-dependent; at larger grid resolutions it scales as $O(N)$ while sampling-based planners are resolution-independent. The additional overhead of Informed RRT* over RRT* (~ 1.96 s) is attributable to the SVD ellipsoid computation and the narrower sampling region, which requires more rejection of out-of-ellipsoid samples at early iterations.

6.3 Path Curvature and Smoothness

A* produces the highest mean curvature ($\kappa = 0.069 \text{ m}^{-1}$) due to the staircase pattern of 45° direction changes at grid waypoints. RRT* reduces curvature to 0.026 m^{-1} by finding continuous-space shortcuts. Informed RRT* achieves $\kappa = 0.006 \text{ m}^{-1}$, approximately $11.5\times$ lower than A*, because the ellipsoidal sampling concentrates edges along the major axis of the cost-improving region, naturally producing longer and more linear segments. All paths satisfy $\kappa \leq 0.8 \text{ m}^{-1}$ (100% dynamic feasibility).

6.4 Energy Proxy

The kinematic energy proxy captures both traversal cost and direction changes. A* scores 58.6 J, approximately $2\times$ the sampling-based planners (29–31 J), because the staircase pattern causes frequent decelerate- and-accelerate events at each 45° turn. RRT* and Informed RRT* produce smoother paths with fewer direction reversals, leading to substantially lower energy cost even though the path length difference is modest ($< 5\%$).

6.5 Multi-Environment Results

Table 2 shows A* mean path lengths across all seven environment families (3 seeds, medium density). Canyon and Narrow Passage produce the longest paths (38.07 m and 38.27 m respectively) because the planners must route around full-height walls, adding significant detour cost. Random Clutter produces the shortest A* paths (33.97 m), as obstacles are scattered rather than forming barriers.

6.6 Success Rates

All three algorithms achieve 100% success (10/10 seeds) on the random clutter medium-density environment. Wilson score 95% CI: [96.9%, 100%] for each algorithm. At high density (22%) on narrow passage and maze environments, RRT* and Informed RRT* occasionally fail within the 2000-iteration budget; A* never fails within the search space because it is complete.

7 Discussion

The fundamental trade-off. No single planner dominates all objectives (Fig. 3). A* wins decisively on speed ($200\text{--}250\times$ faster) and completeness. Informed RRT* wins on path quality: shortest paths, smoothest trajectories, lowest curvature, lowest energy. RRT* sits in between: shorter paths than A*, faster than Informed RRT*.

Grid discretisation artefacts. A*'s high curvature ($11.5\times$ that of Informed RRT*) and energy cost ($2\times$ higher) arise entirely from lattice constraints, not from the algorithm being inherently worse. At infinite grid resolution, A*'s path would converge to the continuous-space optimum. For practical drone flight at 1 m voxel resolution, these artefacts are significant: the staircase pattern requires the drone to turn 45° or 90° at nearly every waypoint, which requires deceleration and wastes energy.

Informed RRT* curvature advantage. The $11.5\times$ curvature reduction of Informed RRT* over A* is explained by the geometry of ellipsoidal sampling. After finding an initial solution, the ellipsoid major axis aligns with the start-to-goal direction. Most samples lie near this axis, causing tree edges to be long and nearly colinear, which directly reduces Menger curvature at each waypoint.

Implementation correctness matters. The corrected Karaman–Frazzoli rewiring radius (Eq. 1) is critical: at $n = 100$ nodes, the correct radius is 7.50 m; at $n = 1000$ it is 5.07 m; at $n = 3000$ it is 3.69 m. A fixed radius of 4.5 m (a common implementation default) would be undersized at small n (missing rewiring opportunities) and oversized at large n (causing $O(n^2)$ nearest-neighbour queries), and does not satisfy the theoretical conditions for optimality.

Practical implications. For time-critical reactive planning (< 100 ms budget), A* is the clear choice. For offline planning where smoothness and energy matter, Informed RRT* is preferred. For medium-latency replanning (100 ms–1 s), RRT* offers a reasonable compromise.

8 Limitations

1. **Grid resolution.** All results use 1 m/voxel resolution. At finer resolutions, A*'s curvature artefacts decrease and its planning time increases. A scalability study at $50 \times 50 \times 20$ and $100 \times 100 \times 30$ is planned.
2. **Iteration budget.** RRT* and Informed RRT* use $N_{\max} = 2000$. Higher budgets would improve path quality further; we did not explore the full convergence curve.
3. **Hardware validation.** All results are from pure simulation. Physical drone experiments with tracking RMSE measurement are future work.
4. **No hardware-aware planning.** The energy proxy is a simplified kinematic model; true rotor energy depends on aerodynamics, motor curves, and payload.
5. **Python timing overhead.** All timing includes Python interpreter overhead. C++ implementations would be 10–50× faster in absolute terms, but relative comparisons remain valid.

9 Conclusion

We present a reproducible benchmark comparing A*, RRT*, and Informed RRT* for 3D drone navigation across 630 trials spanning seven environment families and three obstacle densities. Key findings: (1) A* is 200× faster but produces paths with 11.5× higher curvature and 2× higher energy cost than Informed RRT*; (2) all three planners achieve 100% success; (3) all paths satisfy the dynamic feasibility criterion $\kappa \leq 0.8 \text{ m}^{-1}$; (4) correct implementation of the Karaman–Frazzoli rewire radius and SVD ellipsoidal sampler are essential for reproducing the theoretical guarantees.

All source code, raw CSV results, and reproduction scripts are released at <https://github.com/MedisettiRenukeswar>.

Acknowledgements

The author thanks the open-source communities behind NumPy, SciPy, and Matplotlib, which made this work possible.

References

- [1] 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, 1968.
- [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, 2011.
- [3] 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)*, 2014, pp. 2997–3004.
- [4] I. Pohl, “Heuristic search viewed as path finding in a graph,” *Artificial Intelligence*, vol. 1, no. 3–4, pp. 193–204, 1970.
- [5] K. Daniel, A. Nash, S. Koenig, and A. Felner, “Theta*: Any-angle path planning on grids,” *Journal of Artificial Intelligence Research*, vol. 39, pp. 533–579, 2010.
- [6] S. M. LaValle, “Rapidly-exploring random trees: A new tool for path planning,” 1998, technical Report TR 98-11.
- [7] 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.
- [8] M. Elbanhawi and M. Simic, “Sampling-based robot motion planning: A review,” *IEEE Access*, vol. 2, pp. 56–77, 2014.
- [9] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006. [Online]. Available: <http://planning.cs.uiuc.edu/>







Figure 1: 3D path visualisation on a representative random-clutter environment (seed=42, density=14%). Green triangle: start (1, 1, 1). Purple star: goal (23, 23, 10). **Left:** A* path (33.97 m) showing characteristic staircase waypoints. **Centre:** RRT* path (33.08 m) with smoother continuous-space routing. **Right:** Informed RRT* path (32.64 m) with the smoothest trajectory. All figures generated from actual algorithm execution.




(a) Path length (m): mean $\pm$ SD over 10 seeds.



(b) Planning time (ms, log scale): mean $\pm$ SD.



(c) Mean Menger curvature $\bar{\kappa}$ (m^{-1}).



(d) Kinematic energy proxy (J).

Figure 2: Core benchmark metrics: random clutter environment, medium density (14%), 10 seeds. All values derived from actual code execution.




Figure 3: Pareto front: path length vs. planning time (log-scale x -axis). Mean $\pm$ SD over 10 seeds. No single planner dominates: A* is fastest but longest; Informed RRT* finds the shortest paths at higher computational cost.




Figure 4: A* mean path length across seven environment families (3 seeds, medium density). Canyon and Narrow Passage require the longest detours due to full-height wall obstacles.