

# Supplementary Material for Comparison with Classical Sampling- and Search-Based Planners

We compare our method against classical planners such as RRT\* [Karaman and Frazzoli, 2011] and A\* [Hart et al., 1968], which are widely used in robotics. In this supplementary document we report such results for completeness and to contextualize our main experiments on `maze-large-v1`.

**Environment setting mismatch.** Classical search- and sampling-based planners assume access to an explicit environment model (e.g., an occupancy grid or cost map) and typically a known dynamics model. In contrast, SafeFlow-Matcher is designed for the setting where neither the dynamics nor the occupancy map is available at test time: the model is trained on expert state paths and plans directly in state space. Therefore RRT\* and A\* operate under a strictly *stronger* environment assumption than our generative planner and are not included as main baselines in the paper. Instead, we present them here as additional reference results.

## 1 Experimental Setup

We evaluate RRT\* and A\* in the same Maze2D layout as used in our main experiments. Both planners are given the **exact occupancy map** and the **start/goal states**. To make the spatial resolution comparable to our generative planner, we fix the edge step length to

$$\Delta s = 0.04,$$

which matches the average waypoint spacing of SafeFlowMatcher. Specifically, we generate 100 paths with SafeFlowMatcher, compute the Euclidean distance between consecutive waypoints along each path, and obtain an empirical average step length of approximately 0.0404; this value is rounded to 0.04 and used for both classical planners.

For RRT\* we use a rewiring radius of 1.0. Both RRT\* and A\* perform collision checking directly on the ground-truth occupancy map. We run each planner on 100 independent planning problems and report the same metrics as in the main Maze2D experiments: BS1, BS2, Score, T-Time, Trap rate, curvature,

and acceleration smoothness. Metric definitions are identical to those in the main paper and Appendix E.2.

## 2 Results

Table 1 summarizes the quantitative results. All three methods reach the goal without local traps and satisfy the safety constraints ( $BS1, BS2 \geq \delta = 0.01$ ). Among them, SafeFlowMatcher achieves the highest task score ( $1.632 \pm 0.003$  vs.  $0.874 \pm 0.005$  for RRT\* and  $0.854 \pm 0.000$  for A\*).

RRT\* and A\* exhibit much smaller curvature and acceleration because both planners generate predominantly piecewise-linear, near-straight-line paths, resulting in inherently smoother paths. RRT\* yields slightly better score and smoothness than A\*, but at a substantially higher computational cost (T-TIME 11.674 s vs. 0.020 s). The grid-based A\* search is orders of magnitude faster, but produces less smooth paths, consistent with its larger  $\kappa$  and  $a$  values.

Importantly, classical planners cannot leverage the intrinsic structure encoded in the expert path distribution; they rely solely on an explicit occupancy map and handcrafted search heuristics. By contrast, SafeFlowMatcher plans without any map and exploits statistical regularities learned from demonstrations, capturing high-level planning strategies unavailable to map-based methods. Thus, the classical planners serve only as reference baselines illustrating the behavior of explicit-map methods in the same maze.

Method	BS1 ( $\uparrow$ ) ( $\geq 0$ )	BS2 ( $\uparrow$ ) ( $\geq 0$ )	Score ( $\uparrow$ )	T-TIME (s)	TRAP RATE	$\kappa$ ( $\downarrow$ )	$a$ ( $\downarrow$ )
RRT* [Karaman and Frazzoli, 2011]	0.010	0.010	$0.874 \pm 0.005$	11.674	0%	$0.40 \pm 0.10$	$1.61 \pm 0.41$
A* [Hart et al., 1968]	0.010	0.010	$0.854 \pm 0.000$	0.020	0%	$5.76 \pm 0.00$	$9.57 \pm 0.00$
SafeFlowMatcher (ours)	0.010	0.010	$1.632 \pm 0.003$	9.957	0%	$69.19 \pm 1.02$	$91.90 \pm 0.77$

Table 1: **Comparison between classical planners and SafeFlowMatcher on maze-large-v1.** RRT\* and A\* are evaluated with access to the full occupancy map, whereas SafeFlowMatcher operates without any explicit map and uses only demonstration data and CBF-based corrections. All methods are evaluated over 100 trials using the same evaluation protocol and safety metrics as the Maze2D experiments reported in Table 1 of the main paper.

## 3 Qualitative Examples

Figure 1a and Figure 1b show representative paths generated by RRT\* and A\*, respectively. In both figures, the background displays the maze (gray: occupied space, white: free space), and the colored line shows the planned path from start (blue) to goal (red). The red ellipses highlight the two safety-constraint regions used in the main paper.




Figure 1: Example paths of classical planners in **maze-large-v1**. Both RRT\* and A\* use the exact occupancy map with step length  $\Delta s = 0.04$ . The colored path shows the solution from start (blue) to goal (red), and red ellipses indicate the safety-constraint regions used in the main paper. While both planners use the same step length, the grid-based A\* search produces a less smooth path—consistent with its larger  $\kappa$  and  $a$  values in Table 1—but is significantly faster than RRT\*.

## References

- Peter E. Hart, Nils J. Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968. doi: 10.1109/TSSC.1968.300136.
- Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The International Journal of Robotics Research*, 30(7):846–894, 2011. doi: 10.1177/0278364911406761.