

Maximum Likelihood Path Planning for Fast Aerial Maneuvers and Collision Avoidance

Ji Zhang, Chen Hu, Rushat Gupta Chadha, and Sanjiv Singh

Abstract—We propose a planning method to enable fast autonomous flight in cluttered environments. Typically, autonomous navigation through a complex environment requires a continuous search on a graph generated by a k-connected grid or a probabilistic scheme. As the vehicle travels, updating the graph with data from onboard sensors is expensive as is the search on the graph especially if the paths must be kinodynamically feasible. We propose to avoid the online search to reduce the computational complexity. Our method models the environment differently in two separate regions. Obstacles are considered to be deterministically known within the sensor range and probabilistically known beyond the sensor range. Instead of searching for the path with the lowest cost (typically the shortest path), the method maximizes the likelihood to reach the goal in determining the immediate next step for navigation. With such a problem formulation, the online method realized by a trajectory library can determine a path within 0.2-0.3ms using a single CPU thread on a modem embedded computer. In experiments, it enables a lightweight UAV to fly at 10m/s in a cluttered forest environment (see Fig. 1 as an example).

I. INTRODUCTION

The paper aims to solve a path planning problem to enable fast autonomous flight in complex environments. The problem remains challenging because planning paths to avoid obstacles discovered by onboard sensors requires creating and updating a representation of the environment that can be searched for kinodynamically feasible paths. The process is computationally expensive. Since computational resources available for lightweight aerial vehicles are limited, we need a method that can guide an aerial vehicle with low computational complexity. A typical way is to use a hierarchical approach that separates the planning problem into two subproblems. The first problem solves a global planning problem possibly assisted by a heuristic to ensure the path does not fall into local minima. A second problem solves a local planning problem that runs in parallel to track the global path as well as avoid obstacles. This method has been used successfully in autonomous navigation [1]–[3] but still requires considerable computation. In this paper, we propose a method that reduces the computational complexity considerably such that it can ensure safe flight using very lightweight computation onboard the aerial vehicle.

The key idea to make the low computational complexity possible is avoiding the online search. Instead of searching a graph that is continuously being updated by onboard sensors, we formulate the planning problem from a likelihood point of



Fig. 1. A photo from a flight experiment where our method enables a lightweight aerial vehicle to maneuver at 10m/s in a cluttered forest environment. More details regarding the experiment are in Section V-B.

view. The method does not seek the path with the lowest cost (typically the shortest path) but maximizes the probability to reach the goal in determining the immediate next step for execution of the navigation. This is through modeling of the configuration space differently in two separate regions. Obstacles are considered to be deterministically known within the sensor range as they are perceived by onboard sensors, and probabilistically known beyond the sensor range as they are from a prior map. If the prior map is unavailable, the method can also use a simple heuristic to guide the navigation. A trajectory library is used to bridge the probabilities across the sensor range, where the trajectories are separated into groups. During the navigation, the method evaluates each of the groups to determine the path.

Solving the planning problem with a probabilistic representation of the environment also associates new behaviors to the vehicle. Most of the existing methods find a single path with the lowest cost. The path can be the shortest in length but may guide the vehicle through narrow pathways leaving very few choices for the vehicle to avoid more obstacles, if more obstacles are discovered due to dynamic obstacles or environmental changes such that the prior map is outdated. The proposed method, seeking the highest probability of successful navigation, prefers opener spaces for traversal even though the resulting path can be longer, leaving more choices for obstacle avoidance during the navigation. Further, we have identified certain cases where existing methods based on deterministic representations of the environment encounter difficulty in finding feasible paths. The proposed method handles the cases (see Section V-A for details).

During the navigation, the method can find a path within 0.2-0.3ms using a single CPU thread on a modem embedded computer. Our experiment results are in a public video¹.

¹Experiment video: <https://youtu.be/VYtQt2NcY0Q>

II. RELATED WORK

The proposed method is most related to the literature in path planning and collision avoidance. The problem involves solving for a path for a vehicle to travel from start to goal given a representation of the environment. Graph search-based methods such as Dijkstra [4], A* [5], and D* [6] algorithms traverse different states on the graph to search for paths. On the other hand, sampling based methods cover the graph with random samples. Paths are generated by connecting selected samples. Contemporary sampling-based methods such as Rapidly-exploring Random Tree (RRT) [7] and its variants [8]–[11] have shown promising results to handle maps in large scales, generating paths in a relatively short amount of time. However, these methods require updating the graph and searching the graph continuously during the navigation. The computational complexity can be excessive if the environment is cluttered and complex. Finding a path is not guaranteed within a fixed amount of time.

Certain planning methods pre-process a map to extract traversable information as a means to facilitates the online search. For example, Probabilistic Roadmap (PRM) [12], [13] based methods randomly sample on the map to create a connectivity graph. Paths are then found by searching the graph. Other examples include Voronoi graph [14] and vector field [15]. In essence, these methods share the insight of moving part of the processing offline before the navigation starts to accelerate the online processing. However, the online processing still needs to traverse the graph to search for the path, and hence can be computationally expensive.

The proposed method employs a probabilistic representation of the environment. The concept of modeling uncertainty has been introduced to the path planning literature [16]. For example, the method of Berg et al. [17] uses a linear-quadratic controller with Gaussian models to take into account the uncertainties of the robot motion and state. Melchior and Simmons [18] extend RRT with particles on each node to handle the uncertainty of terrain friction. These methods involve probabilities to model the motion or state of the vehicle. Chung et al. [19] model the edge costs on a graph with uncertainties for graph search. Heiden et al. [20] use probabilities to model the traversability of map voxels. The proposed method, however, models obstacles within the sensor range deterministically, and beyond the sensor range probabilistically as they are known from a prior map.

Our previous work dedicated to enabling fast autonomous flight in cluttered environments [21], [22]. These methods use a prior map to pre-plan alternative paths offline. The online navigation chooses one of the pre-planned paths to execute. The contribution of this paper is proposing a method to enable the capability without the necessity of a prior map. Based on a probabilistic representation of the environment, the method maximizes the likelihood of successful navigation to the goal. If a prior map is unavailable, the method uses a simple heuristic to guide the navigation. To the best of our knowledge, the resulting capability of aerial maneuver without a prior map has not yet been demonstrated.

III. PROBLEM DEFINITION

Define $\mathcal{Q} \subset \mathbb{R}$ as the configuration space of a vehicle. Let $A \in \mathcal{Q}$ be the vehicle current position and $B \in \mathcal{Q}$ be the goal point. The vehicle is equipped with perception sensors. Define $\mathcal{S} \subset \mathcal{Q}$ as the space covered in the range of the perception sensors, namely sensor range. Obstacles are modeled to be deterministically known in \mathcal{S} and probabilistically known in $\mathcal{Q} \setminus \mathcal{S}$. Consider the vehicle has multiple directions to choose for the immediate first step as it starts to move from A . For convenience, let us name the state of the vehicle at this step the start state, denoted as x_s . Obviously, different choices of x_s can lead to different routes. As a convention of this paper, let us define $P_B(\cdot)$ to be the probability for the vehicle to successfully reach B from a given state. The probability associated with start state x_s is $P_B(x_s)$. Our planning problem can be defined as the following,

Problem 1: Given $A, B \in \mathcal{Q}$, $\mathcal{S} \subset \mathcal{Q}$, and obstacles in \mathcal{Q} , determine start state x_s^* to maximize the probability $P_B(x_s)$,

$$x_s^* = \arg \max_{x_s} P_B(x_s). \quad (1)$$

The above problem is solved at each step as the vehicle travels along the path, i.e. the vehicle maximizes the probability to reach B at every instant time during the navigation.

IV. METHOD

A. Probabilistic Model

The proposed method maximizes the likelihood for the vehicle to successfully travel from A to B . As stated in the problem definition, obstacles within sensor range \mathcal{S} are considered to be deterministically known as the information is acquired from the perception sensors. Obstacles beyond \mathcal{S} are considered to be probabilistically known if a prior map is available. Otherwise, however, the case is equivalent to no obstacle being present a priori. Fig. 2 illustrates \mathcal{S} as the gray area. Define $\mathcal{F} \subset \mathcal{S}$ as the sensor frontier indicated by the red solid curve. Given start state x_s , a path connects A and B as the black curve. For all possible paths, they must intersect with \mathcal{F} . Here, one can argue that the vehicle can move laterally and does not intersect with \mathcal{F} . In this case, due to the nature of the problem, \mathcal{S} has to be expanded based on the vehicle motion model so that the vehicle does not traverse an area uncovered by \mathcal{S} . Define x_f as the state of the vehicle while passing \mathcal{F} . The conditional distribution of x_f given x_s , $p(x_f|x_s)$, can be derived from the obstacle information provided by the perception sensors. Further, the probability

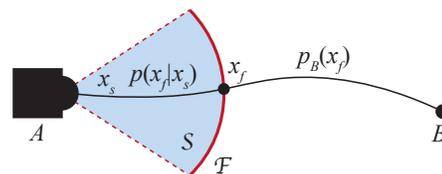


Fig. 2. Illustration of sensor range $\mathcal{S} \subset \mathcal{Q}$ as the gray area and sensor frontier $\mathcal{F} \subset \mathcal{S}$ as the red solid curve. The black curve is a path to navigate from A to B , which starts at x_s and intersects with \mathcal{F} at x_f .

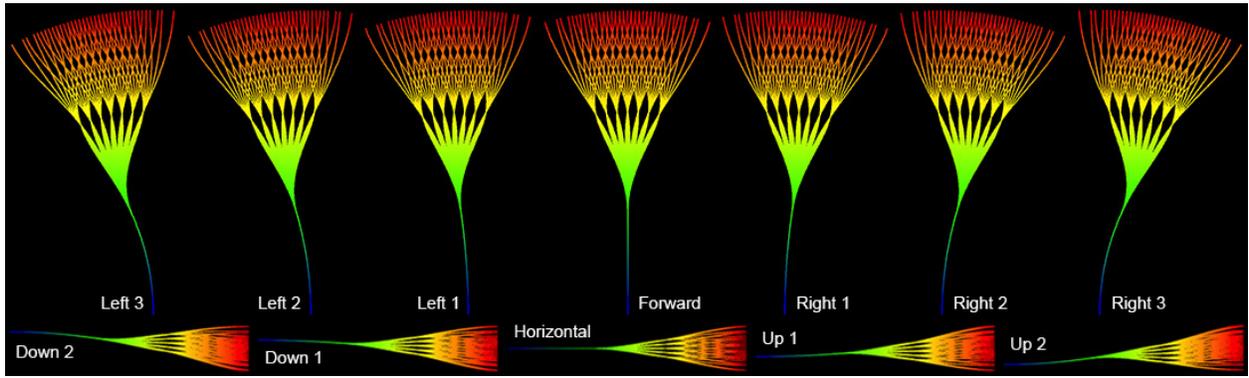


Fig. 3. Example path groups. On the top row, we show 7 path groups curving from left to right in top-down view. On the bottom row, we show 5 path groups curving from downward to upward in side view. Consider both horizontal and vertical directions, there are $7 \times 5 = 35$ path groups.

density for the vehicle to reach B from x_f , $p_B(x_f)$, can be obtained from the obstacles on the prior map. We have

$$p_B(x_s) = \int p_B(x_f)p(x_f|x_s)dx_f. \quad (2)$$

Here, notation $P_B(x_s)$ in (1) is rewritten as $p_B(x_s)$ to denote the probability density. Consider $n \in \mathbb{Z}^+$ samples ξ_i , $i = 1, 2, \dots, n$, drawn from $p(x_f|x_s)$. According to the Monte Carlo theory of sampling [23], we can establish,

$$\int p_B(x_f)p(x_f|x_s)dx_f \approx_{n \uparrow \infty} \frac{1}{n} \sum_{i=1}^n p_B(\xi_i). \quad (3)$$

Combine (2)-(3) and consider n as a constant,

$$p_B(x_s) \approx \frac{1}{n} \sum_{i=1}^n p_B(\xi_i). \quad (4)$$

Eq. (4) indicates that the probability density to navigate to B from x_s , $p_B(x_s)$, can be approximated by $n \gg 1$ samples drawn from the conditional distribution $p(x_f|x_s)$.

B. Local Probabilities

Given start state x_s , the vehicle can follow different paths to reach sensor frontier \mathcal{F} . Here, let us name a path group as the set of paths sharing the same x_s . Consider a discrete model of x_s . Fig. 3 gives an example of path groups. On the top row, 7 path groups are present in top-down view, where x_s is at the start of the paths curving left or right. The path group in the middle corresponds to straight forward motion. On the bottom row, 5 path groups are shown in side view, where the paths curve upward or downward. Consider both horizontal and vertical directions, there are totally $7 \times 5 = 35$ path groups in this example. All paths end on \mathcal{F} .

Each path is generated as a cubic spline curve. The paths in a group split in multiple directions horizontally and vertically. In Fig. 3, the path first splits in 35 directions (7 horizontal and 5 vertical) and each splits in another 35 directions. This results in $35^2 = 1225$ paths in a group. Consider the 35 path groups, there are $35 \times 1225 = 42875$ paths in total. Fig. 4 shows all path groups together where color codes the group index. Note that these example paths are generated based on the vehicle motion constraints. The

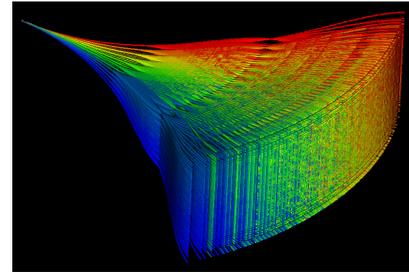


Fig. 4. All 35 path groups. The paths are color coded based on the group index. There are 1225 paths in each group and 42875 paths in total.

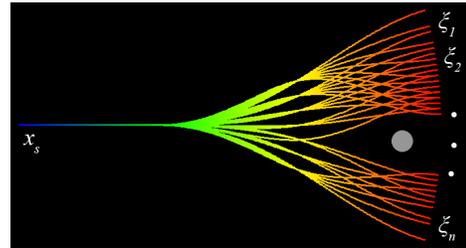


Fig. 5. Collision-free paths in a group with an obstacle as the gray dot. The paths start at x_s . The path ends are considered Monte Carlo samples ξ_i , $i = 1, 2, \dots, n$, whose distribution is drawn from $p(x_f|x_s)$.

method, however, is not limited to a specific motion model and can support various path group configurations.

The paths in a group can be considered as viable routes from x_s to \mathcal{F} . The states of the paths at the ends can be viewed as samples ξ_i , $i = 1, 2, \dots, n$, of x_f , where the distribution is drawn from $p(x_f|x_s)$. During the navigation, obstacles are detected by perception sensors occluding certain paths. Fig. 5 gives an example of an obstacle and the corresponding collision-free paths in a group. Define a Boolean function $c(\xi_i)$ to indicate the path clearance,

$$c(\xi_i) = \begin{cases} 1, & \xi_i \text{ is unoccluded,} \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

We can compute $P_B(x_s)$ based on (4),

$$P_B(x_s) \approx \frac{\sum_{i=1}^n c(\xi_i)p_B(\xi_i)}{\sum_{i=1}^n c(\xi_i)}, \quad (6)$$

Eq. (6) is applied to all path groups and x_s^* of the group with the highest $P_B(x_s)$ is chosen for the vehicle to execute.

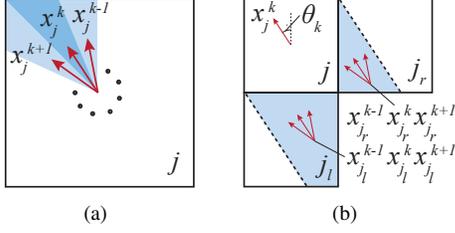


Fig. 6. (a) Voxel representation. Each voxel contains multiple directions at a constant angular interval. The state of a voxel is denoted as x_j^k , $j, k \in Z$, where j is the voxel index and k is the direction index in the voxel. (b) Probability transmission. The probabilities are transmitted to x_j^k from the adjacent voxels j_l and j_r in three directions from each voxel.

C. Global Probabilities

Our environment is represented with voxels. Different from the traditional voxel representation, our voxels contain both position and orientation information. As shown in Fig. 6(a), a voxel is separated into multiple directions based on a constant angular interval, denoted as δ . Define x_j^k , $j, k \in Z$ as the state of the voxel, where j is the voxel index and k is the direction index. The position associated with x_j^k is modeled to be uniformly distributed within the voxel and the orientation is modeled to be uniformly distributed within $[-\delta/2 \delta/2]$ around the direction of x_j^k . The probability density to reach B from x_j^k is denoted as $p_B(x_j^k)$.

The probabilities are transmittable between adjacent voxels. As illustrated in Fig. 6(b), consider the case that the probabilities are transmitted to x_j^k from the adjacent voxels, denoted as j_l on the left side and j_r on the right side. Let θ_k be the direction associated with x_j^k . As the position is modeled to be uniformly distributed within a voxel, the probabilities to be transmitted to x_j^k are from the gray regions in voxels j_l and j_r , with areas $1 - \tan \theta_k/2$ and $\tan \theta_k/2$ of a voxel, respectively. From each of the gray regions, the probabilities are transmitted from three adjacent directions. The probability density transmission is defined as,

$$p_B(x_j^k) = \left(\left(1 - \frac{\tan \theta_k}{2}\right)a + \frac{\tan \theta_k}{2}b \right) r_j, \quad (7)$$

where

$$\begin{aligned} a &= w_y p_B(x_{j_l}^{k-1}) + w_f p_B(x_{j_l}^k) + w_y p_B(x_{j_l}^{k+1}), \\ b &= w_y p_B(x_{j_r}^{k-1}) + w_f p_B(x_{j_r}^k) + w_y p_B(x_{j_r}^{k+1}). \end{aligned}$$

In (7), r_j represents the traversability of voxel j due to obstacles, where $r_j = 1$ means complete clearance and $r_j = 0$ means complete occlusion. w_f and w_y determine the probability distribution corresponding to forward motion and tuning in yaw, respectively. We require that,

$$w_f + 2w_y = 1. \quad (8)$$

The 3D case is a direct extension from the 2D case. Each voxel is separated into multiple directions not only in yaw but also in pitch. When transmitting probabilities, we take into account the probabilities transmitted between voxels at the same level as well as adjacent levels. Here, we omit the equations of the probability density transmission

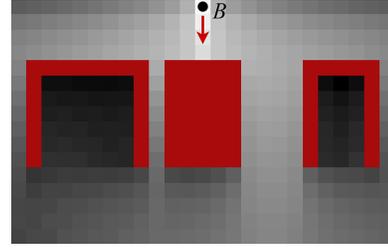


Fig. 7. Propagated probability densities in a 2D environment. The arrow represents the direction in the voxels that the probability densities are associated. Red areas are obstacles with the traversability defined in (7) set at $r_j = 0.01$. Brighter voxels have higher probability densities to navigate from itself to B and darker voxels have lower probability densities.

in 3D case due to their redundancy. During initialization, the probability densities are evenly distributed among all directions in the voxel containing B . Propagation of the probability is through an iteration process. Fig. 7 gives an example of the propagated probability densities in a 2D environment. Brighter indicates higher probability density.

D. Method Implementation

The path groups described in Section IV-B are generated offline. For collision check, we use a voxel grid overlaid with sensor range \mathcal{S} . The correspondences between the voxels and paths are pre-established and stored in an adjacency list. In the adjacency list, each row is associated with a voxel and consists of indexes of the paths that are occluded by the voxel. Here, the vehicle radius is taken into account for calculating the occlusions. Upon system starts, the paths and adjacency list are loaded into the vehicle computer memory. The online collision check processes all perception sensor data points and labels the corresponding paths to be occluded according to the adjacency list. Then, the algorithm traverses all paths in each group to compute $P_B(x_s)$ based on (6) and chooses the path group with the highest $P_B(x_s)$.

Theorem 1: The online processing has a computational complexity of $O(mn)$, where m is the number of perception sensor data points and n is the number of paths.

Theorem 1 analyzes the computational complexity in the worst case where every perception sensor data point blocks all paths in each group. In practice, a data point can block a few paths so that the computation is much lighter.

The probability propagation in the global scale uses a second voxel grid covering the environment, run only once before the navigation. This uses an implementation similar to the A* algorithm [24], where only the probability densities in the voxels adjacent to those in the open set are updated. This process terminates if the changes to the probability densities in the voxel containing A are smaller than a threshold. In the case that a prior map is unavailable, we can alternatively use a heuristic function to compute $p_B(\xi_i)$ in (6),

$$p_B(\xi_i) = -|\Delta p_i \Delta y_i|, \quad (9)$$

where Δp_i and Δy_i are the relative angles between ξ_i and the goal direction in pitch and yaw, respectively.

V. EXPERIMENTS

The underlying experiments use the path groups described in Section IV-B. Sensor range S is set at 30m in front of the vehicle. The collision check uses a voxel grid overlaid with S at 0.1m resolution. The probability propagation uses another voxel grid covering the environment at 1m resolution. A 3.1GHz i7 computer records the CPU processing time.

A. Simulation

We first validate the method in simulation. Fig. 8 shows the result of a test containing a narrow pathway and a wide pathway. Existing planning methods mostly find the shortest path connecting A to B regardless of the width of the pathway. In Fig. 8, the orange curve is generated by RRT* [10]. Our method, however, seeks the highest probability to reach the goal and therefore prefers the wider pathway. Wide pathways are preferable in aerial navigation keeping the vehicle safe as well as leaving more choices for obstacle avoidance. Note that the same behavior can be produced in the existing methods by considering the distance from a node to the closest obstacle in the cost function. The resulting methods need tuning between different environments.

Fig. 9 shows the result of a test involving an environment change between the prior map and actual world. Fig. 9(a) shows the prior map where an opening is available to the right. Our method uses the prior map for probability propagation before the navigation starts. Obstacles on the prior map have the traversability set at $r_j = 0.01$. Fig. 9(b) shows the actual world where the opening on the prior map is closed. Another opening is now available to the front. After the navigation starts, our method generates a path curving to the right as an effect of the opening on the prior map. As the vehicle approaches, the method realizes the environment change from the perception sensors and then guides the vehicle toward the opening to the front. Thanks to the minor probabilities propagated through the obstacles on the prior map. On the other hand, traditional methods based on deterministic representations of the environment encounter difficulty. In Fig. 9(c) RRT* generates a path using

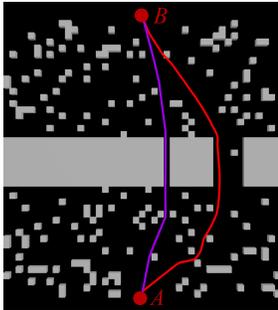


Fig. 8. Narrow pathway test result. The environment contains a narrow pathway on the left and a wide pathway on the right. The widths of the narrow and wide pathways are 5m and 20m, respectively. Existing planning methods such as RRT* (magenta curve) mostly seek the shortest path regardless of the width of the pathway. Our method (red curve) maximizes the probability to reach the goal and therefore prefers the wider pathway. Wider pathways help keep the vehicle safe during the flight and leave more choices for obstacle avoidance in the presence of dynamic obstacles.

the prior map when the navigation starts. As the vehicle travels, in Fig. 9(d), the environment seen by the perception sensors is updated indicating that the opening to the right is unavailable. However, the opening to the front has not been seen. As a result, RRT* finds no path from A to B .

The proposed method is further tested in 2D random world environments. The result is in Fig. 10, where our method is compared to RRT [7], RRT-Connect [8], RRT*, and BIT* [11] planners. Fig. 10(a) shows an example environment and representative paths generated by all methods. The obstacle ratio is set at 20%. Fig. 10(b) compares the runtime and corresponding success rate. The runtime of our method is from the probability propagation. Each method is tested in 50 randomly generated environments and executed 10 times in each of the environments. Note that RRT and RRT-Connect terminate after finding the first path. Given the same success rate, their runtime is faster than ours but the resulting paths are not practically useful. RRT* and BIT* are configured to terminate after finding the near-optimal path. Their runtime is much slower than ours given the same success rate.

We also test the method in 2D maze environments. As shown in Fig. 11, the size of the maze is set at 45×45 . Fig. 11(a) shows an example environment. All methods

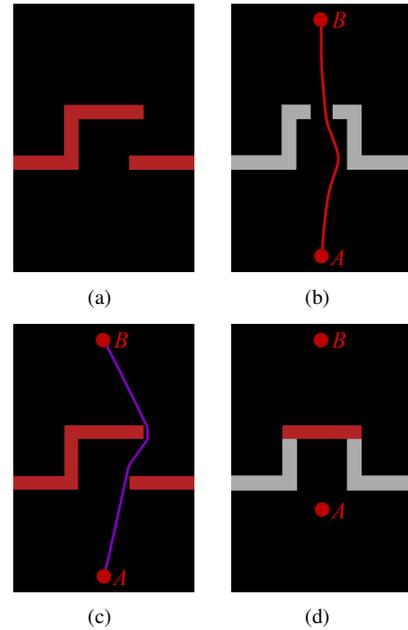
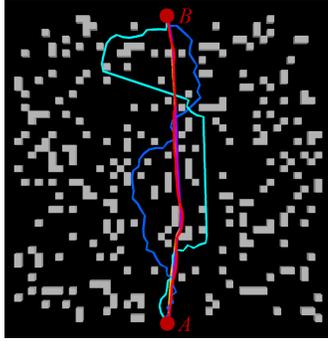
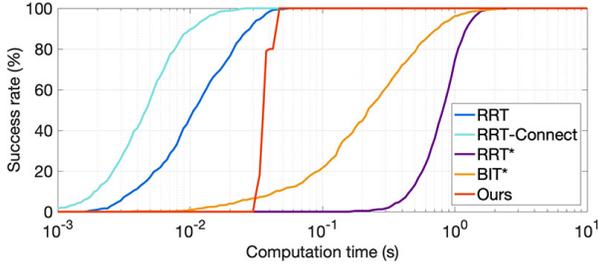


Fig. 9. Outdated prior test result. The test involves an environment change between the prior map and actual world. (a) shows the prior map where an opening is available to the right. Our method uses the prior map for probability propagation. Obstacles on the prior map have the traversability set at $r_j = 0.01$. (b) shows the actual world where the opening is to the front due to the environment change. Upon the navigation starts, the vehicle curves to the right as an effect of the opening on the prior map. As the vehicle approaches, perception sensor data indicates the environment change and the vehicle is then guided toward the opening to the front because of minor probabilities propagated through the obstacles on the prior map. On the other hand, existing methods based on deterministic representations of the environment encounter difficulty. In (c), at the start of the navigation, RRT* uses the prior map to plan a path. During the navigation, the environment within S is updated continuously by the perception sensors. In (d), the vehicle realizes the opening to the right is unavailable. However, the opening to the front has not yet been seen. RRT* finds no path.

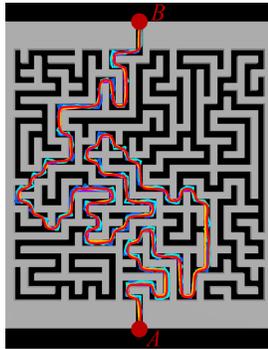


(a)

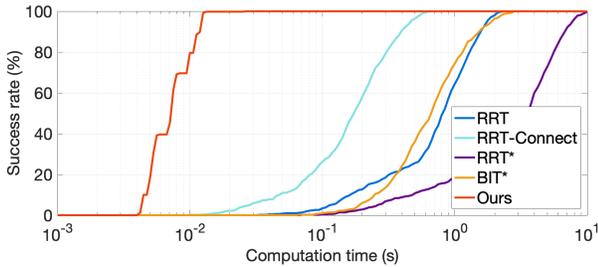


(b)

Fig. 10. 2D random world test result. The proposed method is compared to RRT, RRT-Connect, RRT*, and BIT* in random world environments. The obstacle ratio is set at 20%. (a) shows an example environment and the path from each method. (b) shows a comparison of the runtime by testing in 50 environments. Note that RRT and RRT-Connect terminate after finding the first path. Their runtime is faster than our method but the paths are not practically useful. RRT* and BIT* terminate after finding the near-optimal path. Their paths are similar to our method but the runtime is slower.



(a)



(b)

Fig. 11. 2D maze test result. The proposed method is compared to RRT, RRT-Connect, RRT*, and BIT* in maze environments. The size of the maze is set at 45×45 . (a) shows an example environment and the paths where all methods generate similar paths. (b) shows a comparison of the runtime by testing in 50 environments. Our method is more than 10 times faster than RRT-Connect and about 100 times faster than RRT, RRT*, and BIT*.

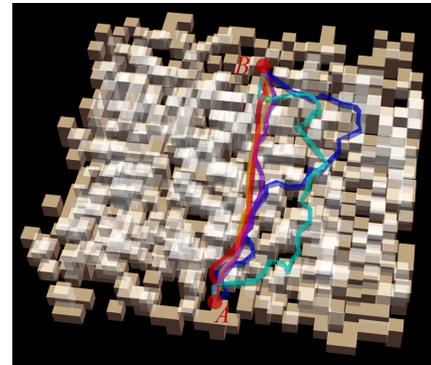
generate similar paths in the maze environments because of constrained space for navigation. Fig. 11(b) compares the runtime. Our method runs more than 10 times faster than RRT-Connect and about 100 times faster than RRT, RRT* and BIT* while producing the same success rate.

The proposed method is tested in 3D cases. Fig. 12 presents the result in 3D random world environments. Fig. 12(a) gives an example environment with a representative path from each method. The obstacle ratio is set at 20%. Fig. 12(b) shows the runtime and corresponding success rate, tested in 50 randomly generated environments and executed 10 times in each environment. Similar to the result in Fig. 10, RRT and RRT-Connect terminate after finding the first path. Even though their runtime is faster than ours, the resulting paths are practically useless. Compared to RRT* and BIT*, our method consumes much less runtime at the same success rate while the three methods produce similar paths.

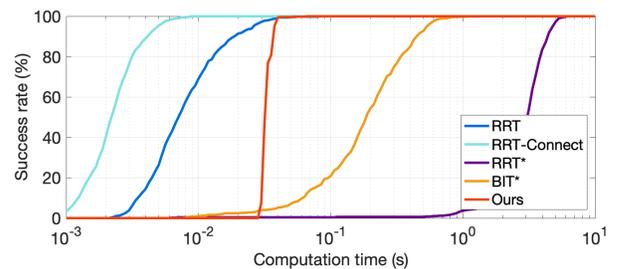
Finally, we evaluate the method in 3D maze environments. The result is in Fig. 13. The size of the maze is set at $25 \times 25 \times 25$. Fig. 13(a) shows an example environment. Similar to Fig. 11, all methods produce similar paths in the maze environments. Fig. 11(b) compares the runtime. Our method outperforms the other methods by a factor of 10 in terms of runtime while producing the same success rate.

B. UAV Experiment

Our experiment platform is shown in Fig. 14. This is a DJI Matrice 600 Pro aircraft carrying a DJI Ronin MX



(a)



(b)

Fig. 12. 3D random world test result. The obstacle ratio is set at 20%. (a) shows an example environment and the path from each method. (b) shows a comparison of the runtime by testing in 50 environments. Similar to Fig. 10, RRT and RRT-Connect terminate after finding the first path. Their runtime is faster than our method but the paths are practically useless. RRT* and BIT* are set to find the near-optimal path. Their paths are similar to our method but the runtime is slower given the same success rate.

gimbal. A sensor-computer pack is mounted to the gimbal and therefore is kept in the flight direction for obstacle detection. The sensor-computer pack consists of a Velodyne Puck laser scanner, a camera at 640×360 pixel resolution, and a MEMS-based IMU. A 3.1GHz i7 embedded computer carries out all onboard processing. The state estimation is based on our previous work [25], which integrates data from the three sensors to provide vehicle poses and registered laser scans. The software also builds a map during the flight.

The test site is in a forest as shown in Fig. 15. The flight test does not use a prior map but the heuristic function in (9) to guide the navigation. Fig. 15(a) shows an aerial overview of the test site. Fig. 15(b) presents an image logged by an onboard camera during the flight. Fig. 15(c) shows a render of the map built during the flight with the executed path overlaid on the map, from the same viewpoint as in Fig. 15(b). Fig. 15(d) presents the registered scans as the perception sensor data during the flight (colored points) and the determined collision-free paths (white curves). The yellow curve is the selected path for the vehicle to execute. The vehicle pose in Fig. 15(d) is the same as in Fig. 15(b). Fig. 15(e) shows the entire map and overall path of the flight. The vehicle position in Fig. 15(b) and Fig. 15(d) are labeled with number 1 in Fig. 15(e). The canopy of the forest is manually cropped to reveal the flight path. The traveling distance is approximately 300m and the maximum speed during the flight reaches 10m/s as indicated in Fig. 16.

Finally, let us inspect some metrics from the UAV test. Different from the simulation tests, the UAV test does not use

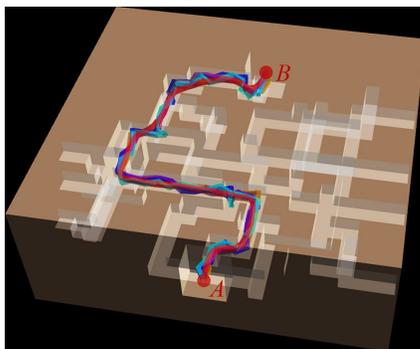
a prior map or propagate the probabilities. At initialization, the onboard navigation system reads the paths and adjacency list (described in Section IV-D) into the computer memory. The onboard processing time is listed in Table I. Collision check first processes all perception sensor data points to determine the collision-free paths, taking $213.7\mu\text{s}$ on average. Then, the processing traverses all paths to compute $P_B(x_s)$ for each group and select the path group with the highest $P_B(x_s)$, taking $38.4\mu\text{s}$ on average. The method runs at 5Hz. The resulting CPU load is $< 5\%$ of a single thread.

VI. CONCLUSION AND FUTURE WORK

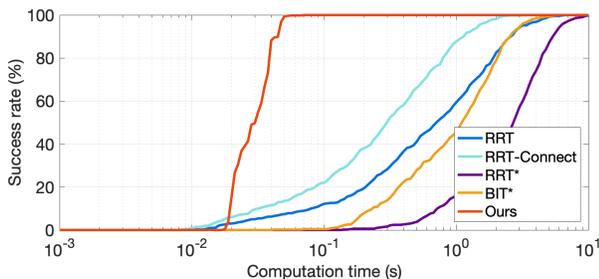
The paper proposes a planning method to enable fast autonomous flight in complex environments. The environment is modeled to be deterministically known within the sensor range where obstacle information is from the perception sensors, and probabilistically known beyond the sensor range. Instead of searching for the path with the lowest cost, the method maximizes the likelihood to successfully reach the goal in determining the immediate next step for execution of the navigation. If a prior map is available, probabilities are propagated offline through the environment. The online method realized by a trajectory library determines a path within 0.2-0.3ms using a single CPU thread on a modest embedded computer. In experiments, it enables a lightweight UAV to fly at 10m/s in a cluttered forest environment.

TABLE I
ONLINE PROCESSING TIME IN UAV TEST

| Collision check | | Path selection | | Overall | |
|--------------------|--------------------|-------------------|-------------------|--------------------|--------------------|
| Mean | Worst | Mean | Worst | Mean | Worst |
| $213.7\mu\text{s}$ | $286.2\mu\text{s}$ | $38.4\mu\text{s}$ | $41.3\mu\text{s}$ | $252.1\mu\text{s}$ | $327.5\mu\text{s}$ |



(a)



(b)

Fig. 13. 3D maze test result. The size of the maze is set at $25 \times 25 \times 25$. (a) shows an example environment and the paths where all methods generate similar paths. (b) shows a comparison of the runtime by testing in 50 environments. Given the same success rate, our method outperforms RRT, RRT-Connect, RRT*, and BIT* by a factor of 10 in terms of runtime.



Fig. 14. UAV experiment platform. A DJI Matrice 600 Pro aircraft carries our sensor-computer pack on a DJI Ronin MX gimbal. The gimbal keeps the sensors in the flight direction for obstacle detection. The sensor-computer pack consists of a Velodyne Puck laser scanner, a camera at 640×360 pixel resolution, and a MEMS-based IMU. An i7 embedded computer carries out all onboard processing. Note that GPS data is unused in the test.

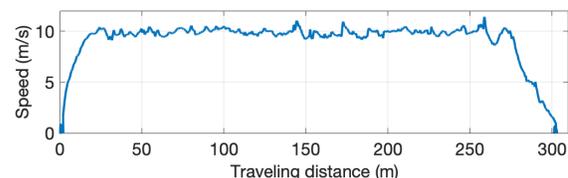


Fig. 16. Speed in UAV test. The maximum speed of the UAV reaches 10m/s while flying in a forest environment as presented in Fig. 15.

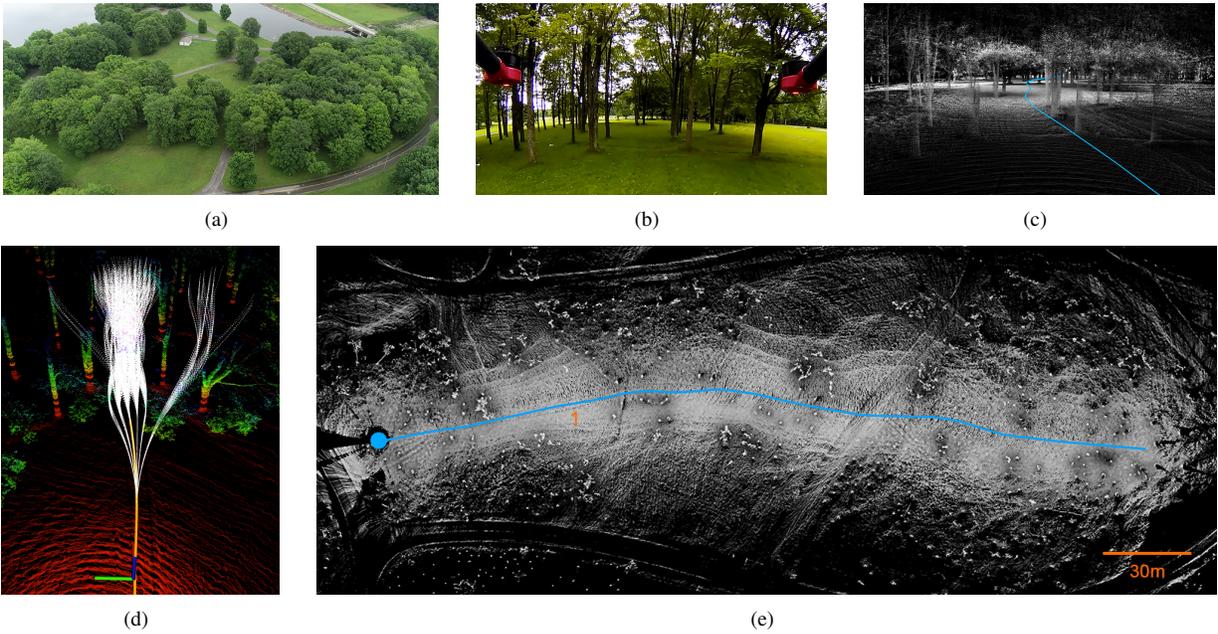


Fig. 15. UAV test result. The flight test is conducted in a forest environment. No prior map is used in the test. The method uses the heuristic function in (9) to guide the navigation. (a) shows an aerial overview of the test site. (b) is an image from an onboard camera captured during the flight. (c) shows a render of the map built during the flight and the executed path overlaid on the map. (d) presents the perception sensor data during the flight as the colored points and the corresponding collision-free paths as the white curves. The yellow curve is the selected path for the vehicle to execute. The vehicle pose is the same as in (b). (e) shows the entire map and overall path of the flight. The vehicle position in (b) and (d) is labeled with number 1 in (e). The canopy of the forest is removed to reveal the path. The flight has 300m of travel and the vehicle speed is at 10m/s through the course of the flight.

REFERENCES

- [1] D. Gonzalez, J. Prez, V. Milans, and F. Nashashibi, "A review of motion planning techniques for automated vehicles," *IEEE Trans. on Intelligent Transportation Sys.*, vol. 17, no. 4, pp. 1135–1145, 2016.
- [2] D. Droschel, M. Nieuwenhuisen, M. Beul, D. Holz, J. Stuckler, and S. Behnke, "Multi-layered mapping and navigation for autonomous micro aerial vehicles," *Journal of Field Robotics*, vol. 33, no. 4, pp. 451–475, 2016.
- [3] S. Scherer, S. Singh, and L. Chamberlain, "Flying fast and low among obstacles: Methodology and experiments," *The International Journal of Robotics Research*, vol. 27, no. 5, pp. 549–574, 2008.
- [4] R. Kala and K. Warwick, "Multi-level planning for semi-autonomous vehicles in traffic scenarios based on separation maximization," *J. of Intelligent and Robotic Systems*, vol. 72, no. 3/4, pp. 559–590, 2013.
- [5] B. MacAllister, J. Butzke, A. Kushleyev, H. Pandey, and M. Likhachev, "Path planning for non-circular micro aerial vehicles in constrained environments," in *IEEE International Conference on Robotics and Automation (ICRA)*, Karlsruhe, Germany, May 2013.
- [6] M. Ruffi and R. Y. Siegwart, "On the application of the D search algorithm to time-based planning on lattice graphs," in *The European Conf. on Mobile Robots (ECMR)*, Dubrovnik, Croatia, Sept. 2009.
- [7] S. M. LaValle, *Planning Algorithms*. New York, NY, USA: Cambridge University Press, 2006.
- [8] J. J. Kuffner and S. M. LaValle, "RRT-connect: An efficient approach to single-query path planning," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, San Francisco, CA, April 2019.
- [9] 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 *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Chicago, IL, Sept. 2014.
- [10] 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.
- [11] J. Gammell, S. Srinivasa, and T. Barfoot, "Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, Seattle, WA, May 2015.
- [12] D. Hsu, J.-C. Latombe, and H. Kurniawati, "On the probabilistic foundations of probabilistic roadmap planning," *The International Journal of Robotics Research*, vol. 25, no. 7, pp. 627–643, 2006.
- [13] L. E. Kavraki, M. N. Kolountzakis, and J.-C. Latombe, "Analysis of probabilistic roadmaps for path planning," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 1, pp. 166–171, 1998.
- [14] P. Beeson, N. K. Jong, and B. Kuipers, "Towards autonomous topological place detection using the extended Voronoi graph," in *IEEE International Conference on Robotics and Automation (ICRA)*, Barcelona, Spain, April 2005.
- [15] G. A. S. Pereira, S. Choudhury, and S. Scherer, "A framework for optimal repairing of vector field-based motion plans," in *Intl. Conf. on Unmanned Aircraft Systems (ICUAS)*, Arlington, VA, June 2016.
- [16] T. Fraichard and R. Mermond, "Path planning with uncertainty for car-like robots," in *IEEE International Conference on Robotics and Automation (ICRA)*, Leuven, Belgium, May 1998.
- [17] J. Van Den Berg, P. Abbeel, and K. Goldberg, "LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 895–913, 2011.
- [18] N. A. Melchior and R. Simmons, "Particle RRT for path planning with uncertainty," in *IEEE International Conference on Robotics and Automation (ICRA)*, Roma, Italy, April 2007.
- [19] J. Chung, A. Smith, R. Skeele, and G. Hollinger, "Risk-aware graph search with dynamic edge cost discovery," *The International Journal of Robotics Research*, vol. 38, no. 2-3, pp. 182–195, 2019.
- [20] E. Heiden, K. Hausman, G. Sukhatme, and A. Agha-mohammadi, "Planning high-speed safe trajectories in confidence-rich maps," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Vancouver, Canada, Sept. 2017.
- [21] J. Zhang, R. G. Chadha, V. Velivela, and S. Singh, "P-CAP: Pre-computed alternative paths to enable aggressive aerial maneuvers in cluttered environments," in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, Madrid, Spain, Oct. 2018.
- [22] —, "P-CAL: Pre-computed alternative lanes for aggressive aerial collision avoidance," in *The 12th International Conference on Field and Service Robotics (FSR)*, Tokyo, Japan, Aug. 2019.
- [23] C. Robert, *Monte carlo methods*. John Wiley & Sons, Ltd, 2004.
- [24] W. Zeng and R. L. Church, "Finding shortest paths on real road networks: the case for A*," *International Journal of Geographical Information Science*, vol. 23, no. 4, pp. 531–543, 2009.
- [25] J. Zhang and S. Singh, "Laser-visual-inertial odometry and mapping with high robustness and low drift," *Journal of Field Robotics*, 2018.