ETH *zürich*

Informed Sampling Exploration Path Planner for 3D Reconstruction of Large Scenes

Journal Article

Author(s):
Kompis, Yves (ID); Bartolomei, Luca (ID); Mascaro, Ruben; Teixeira, Lucas (ID); Chli, Margarita (ID)

Informed Sampling Exploration Path Planner for 3D Reconstruction of Large Scenes

Yves Kompis, Luca Bartolomei, Ruben Mascaro, Lucas Teixeira, Margarita Chli

Abstract—As vision-based navigation of small aircraft has been demonstrated to reach relative maturity, research into effective path-planning algorithms to complete the loop of autonomous navigation has been booming. Although the literature has seen some impressive works in this area, efficient path-planning that can be used in tasks such as inspection and coverage is still an open problem. In this spirit, this letter presents an online path-planning algorithm for fast exploration and 3D reconstruction of a previously unknown area of interest. Micro Aerial Vehicles (MAVs) are an ideal candidate for this task due to their maneuverability, but their limited computational power and endurance require efficient planning strategies. Popular sampling-based methods randomly sample the MAV's configuration space and evaluate viewpoints according to their expected information gain. Most often, however, valuable resources are spent on information gain calculations of unpromising viewpoints. This letter proposes a novel informed sampling approach that leverages surface frontiers to sample viewpoints only where high information gain is expected, leading to faster exploration. We study the impact of informed sampling in a wide range of photo-realistic scenes, and we show that our approach outperforms state-of-the-art exploration path planners in terms of both speed and reconstruction quality.

Index Terms—Aerial Systems: Perception and Autonomy, Motion and Path Planning

I. INTRODUCTION

AS vision-based robot navigation has been demonstrating its ability to perform in increasingly challenging environments, research effort has been shifting towards higher-level tasks, such as path-planning techniques, promising to automate key tasks. For example, inspections of large infrastructure for faults/cracks, digitisation of archaeological or construction sites for monitoring, and 3D reconstruction of existing structures as an architectural aid, all promise great impact, highlighting the importance of efficient, autonomous robotic exploration and mapping of previously unknown structures of interest. The agility and power of Micro Aerial Vehicles (MAVs) have set them the *de facto* platform of choice for such applications as soon as it was evident that they could run vision-based perception onboard [1].

With online exploration promising efficient mapping of structures of interest, the limited resources onboard an MAV
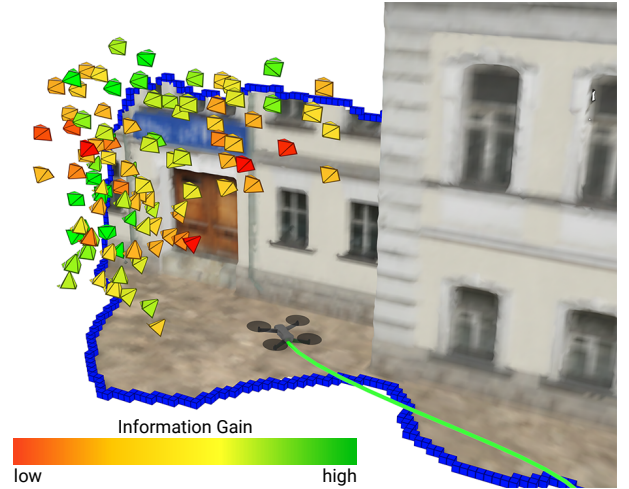
Authors are with the Vision for Robotics Lab, Department of Mechanical and Process Engineering, ETH Zurich, Switzerland. {ykompis, lbartolomei, rmascaro, pilucas, chlim}@ethz.ch

Fig. 1: This work leverages surface frontier voxels (blue) to sample relevant viewpoints, which are evaluated based on their expected information gain. The proposed method guides the MAV through the most promising viewpoints demonstrating faster and more complete coverage of the structure of interest than state of the art in a wide variety of scenes, such as the *City Building* shown here [1].

pose an additional challenge that needs to be considered. One key problem in exploratory path-planning is the need for the robot to explore areas where no information is available while operating inside the boundaries of the currently known map. Thus, a planning strategy must be able to ensure safe navigation around the boundaries of the ever-evolving map and to incorporate new information quickly. Moreover, while fast exploration is desirable, so is the high accuracy of the resulting 3D reconstruction. Notably, purely frontier-based methods ([2], [3]) lack the ability to include the reconstructed surface quality during planning, while early sampling-based methods often used volumetric information-gain formulations ([4], [5]), which do not explicitly consider surface quality. Addressing this, newer sampling-based methods include the quality of the reconstruction in the gain formulation ([6], [7], [8]), combining the exploration and reconstruction tasks into a single objective function. Random sampling approaches, however, waste valuable computational resources on unpromising information gain calculations, as a randomly sampled viewpoint is neither guaranteed to observe any unknown space nor to observe the area of interest. Nonetheless, such viewpoints get evaluated and discarded later due to their low information gain. As a result, uninformed sampling most often leads to fewer proposed viewpoints with high information gain, while considering additional degrees of freedom during sampling, such as yaw and pitch of the sensor,

only exacerbate this issue. For this reason, many sampling-based approaches only sample the position in $\mathbb{R}^3$ by either relying on an omnidirectional sensor [4] or by separately optimizing for the best yaw, such as [8], [9], [10].

Inspired by these shortcomings, we propose an *informed* sampling-based exploration path planner for 3D reconstruction for an MAV equipped with a gimbal-actuated stereo camera[2]. The method makes use of the current version of the reconstructed model to propose viewpoints orthogonal to surface frontiers. Using an Artificial Potential Field (APF) based ranking method, the most suitable proposed viewpoints are determined, which then get evaluated according to an information gain formulation that considers exploration and surface quality. In parallel, a path planner computes a path to the Next-Best-View (NBV), considering the current state of the reconstruction and optimizing the path and camera gimbal orientation, such that surface frontiers are visible during flight. We present evaluations in a variety of photo-realistic simulated scenes, revealing the power and efficacy of the proposed approach for informed viewpoint sampling and the benefit of the APF ranking.

In summary, the contributions of this work are:

- an informed sampling-based exploration path-planner that proposes the most promising viewpoints for further evaluation,
- an effective ranking method of proposed viewpoints, which uses an artificial potential field to predict the value of proposed viewpoints saving the computation time of calculating their actual information gain, and
- a path-planner that optimizes the trajectory and the camera's gimbal orientation in order to keep the structure of interest in view.

II. RELATED WORK

The most popular approaches for exploration planners currently are frontier-based and sampling-based methods. In the former, exploration is driven by frontiers [11] defined as regions of space on the boundaries of known and unknown space. A frontier-driven path-planner can fully explore an environment by observing and eliminating all frontiers. This strategy can be performed in a high-speed fashion, as proposed in [3], by first eliminating frontiers that minimize the MAV's change in velocity. In case no more frontiers are visible, a less efficient global path search is used to drive the robot towards the remaining unexplored areas.

In sampling-based exploration, the robot is guided by several possible trails of sampled configurations, ranked by their expected information gain. Depending on the information gain formulation, both complete exploration (e.g. by including frontier visibility and visibility of unknown space) and accurate surface reconstruction (e.g. by including surface confidence) can be achieved.

Aiming to speed up sampling-based methods, informed sampling strategies can be employed. In [10], for example, viewpoints get sampled from clustered frontiers from an octree

map in an effort to reduce the map entropy, but do not explicitly use the reconstructed surface for sampling. If a 3D surface reconstruction is computed, the current reconstructed model can be utilized to make more educated choices on where to generate viewpoints. In this spirit, [4] propose to offset the reconstructed model surface with a distance transform and use the offset surface as input for uniform viewpoint sampling. Since they assume omnidirectional sensing, such viewpoints are guaranteed to have the object of interest in view. Interestingly, the information gain used by this planner only considers the visibility of surface frontiers, and in order to reach NBVs, an independent path-planner is used. Following a different strategy, [6] sample viewpoints from surface frontiers, which are on the boundaries of the reconstructed surface and unknown space, generating viewpoints at a fixed distance along the surface normal at every surface frontier, ensuring that every sampled viewpoint observes a frontier. Rather than using a computationally intensive ray-cast-based information gain, they only consider travel distances as a metric without proposing a path-planner to reach the target viewpoints with an MAV.

In order to address the efficiency of planning, several sampling-based methods use Rapidly-exploring Random Trees (RRTs) to randomly sample the MAV's configuration space [8], [9], [12], representing viewpoints as nodes in the RRT, and directly using edges as paths to reach every sampled viewpoint. Target nodes get evaluated by an information gain measure, and the MAV is guided along the tree towards the most promising viewpoints in a receding horizon fashion. Following this paradigm, [9] renew the RRT in every planning iteration using the most promising branch of the previous tree as a seed, while [8] rewire the RRT after every planning iteration to reduce unnecessary re-computations of information gain, allowing the tree to continuously grow with new nodes in every planning iteration. To reduce the computational load, [9] use a Gaussian Process to estimate the information gain wherever possible. Both methods assume a four-dimensional configuration space (i.e. the MAV's position and yaw), but sample only positions with the optimal yaw determined in a second step by optimizing the information gain from a $360°$ ray-casting operation. Moreover, [9] only consider unmapped volume in their information gain, while [8] propose an information gain that considers unmapped volume, frontier visibility, and surface quality.

In contrast to informed sampling methods, the aforementioned RRT approaches do not necessarily observe the structure of interest at every node while requiring omnidirectional ray-casts for information gain calculations. Thus, random sampling approaches spend more time on information gain calculations compared to informed sampling methods. In order to create an exploration planner that can efficiently make use of an actuated sensor, here we propose to use surface frontiers for informed viewpoint sampling. This reduces the computational cost of information gain calculations since omnidirectional ray-casts are not needed, while guaranteeing frontier observations from every sampled viewpoint. As in [8], we propose to maintain a global set of evaluated viewpoints for NBV planning to retain and reuse relevant information

---

[2]An RGBD camera could also be used as the method is independent of the sensor used as long as it captures depth and visual values.
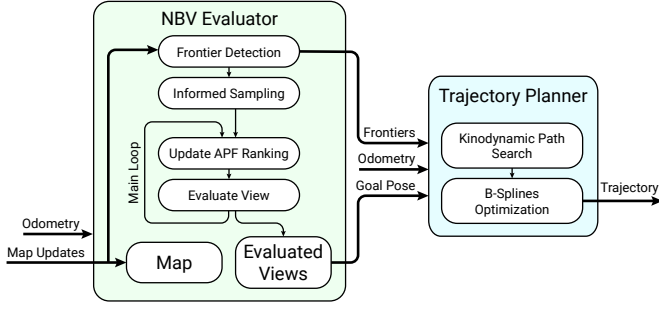
Fig. 2: System overview of the proposed exploration path-planner. The NBV evaluator samples viewpoints and maintains a global set of evaluated viewpoints. The trajectory planner finds feasible paths to these viewpoints.

efficiently, but we do not extrapolate information gains from nearby evaluations as in [9]. Using surface frontiers to guide exploration allows complete mapping of any environment with a connected surface, without the need to explore empty volumes of space. In contrast to RRT methods, our set of evaluated viewpoints is unconnected, meaning that we employ a separate path-planner to guide the MAV to the NBV. To this end, we adapt the perception-aware path-planner proposed in [13] to guide the MAV to its goal, while simultaneously directing the sensor at nearby frontiers.

III. METHOD

Our exploration path-planner consists of two processes running in parallel as shown in the system overview in Fig. 2. A mapping framework pre-processes raw sensor data captured by a stereo camera to build a map incrementally. The NBV evaluator processes this incoming map data into frontiers, which are used for informed viewpoint sampling. We evaluate a subset of sampled viewpoints one-by-one, selected and ordered by an APF ranking scheme, and add them to a global, persistent set of evaluated viewpoints. The trajectory planner successively guides the MAV to the NBV retrieved from the set of evaluated viewpoints.

A. Global Set of Evaluated Viewpoints

An evaluated viewpoint V consists of position $\mathbf{x}(V) \in \mathbb{R}^3$, yaw $\gamma(V)$, camera pitch $\theta(V)$, and information gain $g(V)$.

The global set of evaluated viewpoints $\mathcal{E}$ is maintained throughout the planning process. In each planning iteration, we add a new evaluated viewpoint and insert it into the set. However, since the underlying map is constantly changing and growing, the information gain $g(V)$ of a viewpoint V is subject to change. We discard evaluated viewpoints in a pruning step based on their age by removing them from the set $\mathcal{E}$. As $g(V)$ can only be affected by changes in the map and changes occur around the current MAV position, we do not discard viewpoints, which are sufficiently far from the current MAV position.

B. Map Representation and Information Gain

We use Voxblox [14] as the internal map representation of the scene for planning. Voxblox uses a voxel-based Truncated Signed Distance Field (TSDF) to incrementally build a Euclidean Signed Distance Field (ESDF), which can be used for online planning. Voxels are arranged in a uniformly spaced grid and each ESDF voxel m stores a distance value $d(m)$ and a weight $w(m)$. The distance value represents the distance to the closest obstacle and the voxel weight corresponds to the confidence in the distance value.

To calculate the expected information gain from a viewpoint V we use the formulation proposed by [8], which accounts for surface quality improvements. In a first step, an iterative ray-casting operation determines the set of all voxels $m \in \mathcal{V}is(V)$, which are visible from viewpoint V. Its expected information gain is then given by

$$g(V) = \sum_{m \in \mathcal{V}is(V)} \begin{cases} g_{frontier} & \text{if } m \in \mathcal{F} \\ \dfrac{w_{new}(m)}{w_{new}(m) + w(m)} & \text{if } m \in \mathcal{S} \\ g_{new} & \text{otherwise,} \end{cases} \quad (1)$$

where $\mathcal{F}$ is the set of all frontier voxels, $\mathcal{S}$ is the set of all surface voxels, and $g_{frontier}$ and g_{new} are constant gain values for frontier voxels and new voxels, respectively. For surface voxels, the gain is given by the voxel weight impact, which depends on the current voxel weight $w(m)$ and its expected weight $w_{new}(m)$ after integration of the data observed from the viewpoint V. The expected surface voxel weight from an observation is given by $w_{new}(m) = z^{-2}$, where z is the observation distance between the viewpoint V and the voxel m [14].

C. Surface Frontier Detection

Surface frontier detection is performed on the incoming map data provided by Voxblox, which integrates the raw sensor point cloud. We store the set of frontiers $\mathcal{F}$ in an unordered set which allows for $\mathcal{O}(1)$ frontier insertions, removals, and checks.

A surface frontier voxel $m \in \mathcal{F}$ fulfills three properties: (a) the voxel has been observed, (b) the voxel is part of the surface $m \in \mathcal{S}$, and (c) the voxel has at least one unobserved neighbor in its immediate surroundings. So formally, a voxel m is a surface frontier if

$$w(m) > w_{init}, \quad (2)$$
$$-d_v \leq d(m) \leq d_v, \text{ and} \quad (3)$$
$$\exists m_n \in \mathcal{N}_{26}(m), \text{ s.t. } w(m_n) = w_{init}, \quad (4)$$

where w_{init} is the initial weight of an unknown voxel, d_v is the voxel size, and $\mathcal{N}_{26}(m)$ is the 26-connected neighborhood of voxel m. In Fig. 1 we visualize the set of frontier voxels in an example scene.

For efficiency, we only check for frontiers in areas of the map that could be affected by the latest sensor measurements. Thus, out-of-sensor-range voxels are not checked. This incremental frontier detection is, therefore, independent of the size of the currently reconstructed scene and does not slow down as the map grows.

D. Informed Sampling

A proposed viewpoint V_p consists of position $\mathbf{x}(V_p) \in \mathbb{R}^3$, yaw $\gamma(V_p)$, camera pitch $\theta(V_p)$, and utility $u(V_p)$, which
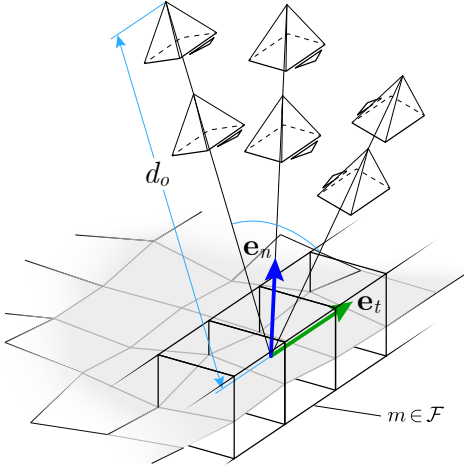
Fig. 3: Viewpoint proposals, drawn as camera frusta, from a single frontier voxel. The reconstructed surface is represented in grey, the frontier voxels in black, the surface normal in blue, and the frontier tangent in green.

is defined in Section III-E. We call $\mathcal{P}$ the set of proposed viewpoints.

For every surface frontier voxel $m \in \mathcal{F}$ we sample multiple viewpoints. The primary observation point of a frontier is placed at a fixed observation distance d_o along the surface normal $\mathbf{e}_n(m)$ from the frontier voxel. The yaw $\gamma(V_p)$ and pitch $\theta(V_p)$ are chosen, such that the frontier voxel m is at the center of the field of view.

Since the primary observation point might not be traversable, or the frontier in question might be part of an edge, it is not guaranteed that the frontier can be resolved from this viewpoint alone. Thus, we add additional viewpoint proposals at multiple distances along the surface normal $\mathbf{e}_n(m)$. Furthermore, to aid in corner situations we rotate the frontier normal $\mathbf{e}_n(m)$ around the frontier tangent $\mathbf{e}_t(m)$ to generate additional viewpoints as shown in Fig. 3.

This procedure is carried out for every surface frontier in $\mathcal{F}$ and all proposed viewpoints V_p are added to the set $\mathcal{P}$.

E. Artificial Potential Field (APF) Ranking

The information gain calculation described in Section III-B is computationally the most demanding task of the proposed exploration planner. It is infeasible for a mobile platform to calculate the information gain for every proposed viewpoint. Therefore, we devised an APF to rank every proposed viewpoint based on its location and surroundings. The APF signifies how useful it would be to calculate the information gain of a proposed viewpoint.

This utility field is influenced by the current position of the MAV $\mathbf{x}_{MAV}$, the set of evaluated viewpoints $\mathcal{E}$, and the set of visited viewpoints $\mathcal{V}$. Intuitively, we evaluate viewpoints close to $\mathbf{x}_{MAV}$, but farther away from already evaluated and already visited viewpoints. This ensures that the set of evaluated viewpoints $\mathcal{E}$ provides relevant information at the current MAV position, which are distinct and not repetitive.

Formulated as an APF, this yields three terms

$$u_M(V_p) = \|\mathbf{x}(V_p) - \mathbf{x}_{MAV}\|, \qquad (5)$$

$$u_{\mathcal{E}}(V_p) = \sum_{V_e \in \mathcal{E}} \mu(V_p, V_e) \frac{1}{\|\mathbf{x}(V_p) - \mathbf{x}(V_e)\|}, \qquad (6)$$

$$u_{\mathcal{V}}(V_p) = \sum_{V_v \in \mathcal{V}} \mu(V_p, V_v) \frac{1}{\|\mathbf{x}(V_p) - \mathbf{x}(V_v)\|}, \qquad (7)$$

where $u_M(V_p)$ ensures proximity to $\mathbf{x}_{MAV}$, $u_{\mathcal{E}}(V_p)$ ensures distinctiveness, while $u_{\mathcal{V}}(V_p)$ prevents repetition. To reduce influences from nearby viewpoints facing in another direction we use the angular discount factor

$$\mu(V_1, V_2) = \max(0, \cos(\Delta\gamma)) \cdot \max(0, \cos(\Delta\theta)) , \qquad (8)$$

where $\Delta\gamma = \angle(\gamma(V_1), \gamma(V_2))$ is the difference in yaw and $\Delta\theta = \angle(\theta(V_1), \theta(V_2))$ is the difference in pitch between two viewpoints. The resulting viewpoint utility of a proposed viewpoint V_p is given by

$$u(V_p) = \lambda_1 u_M(V_p) + \lambda_2 u_{\mathcal{E}}(V_p) + \lambda_3 u_{\mathcal{V}}(V_p) , \qquad (9)$$

where the coefficients λ_1, λ_2, and λ_3 are constant weights.

In every planning iteration we calculate the information gain $g(V)$ of a single proposed viewpoint given by

$$V = \underset{V_p \in \mathcal{P}}{\operatorname{argmin}}\ u(V_p). \qquad (10)$$

This viewpoint is inserted into the set of evaluated viewpoints $\mathcal{E}$, which is visualized in Fig. 1.

F. Perception-aware Path-Planner

When exploration planning starts or when a previous goal is reached, the perception-aware path-planner starts planning towards a new goal selected from the set of evaluated viewpoints $\mathcal{E}$. To prevent long back-and-forth travel moves, the goal is determined by the attenuated information gain

$$\underset{V \in \mathcal{E}}{\operatorname{argmax}} \left(\min \left(g(V), \frac{g(V)}{1 + \alpha(\|\mathbf{x}(V) - \mathbf{x}_{MAV}\| - r)} \right) \right), \qquad (11)$$

with distance threshold radius r, and decay α.

The perception-aware path-planner presented in [13] uses a two-stage planning process with an initial A^* path search based on motion primitives followed by a B-Spline trajectory optimization. Initially, A^* searches a collision free path only considering the MAV's position in $\mathbb{R}^3$. The orientation is considered in the following B-Spline optimization, which optimizes the path for collision avoidance, smoothness, and visibility of frontiers. We extend the state space from $[x, y, z, \gamma]$ to $[x, y, z, \gamma, \theta]$ while keeping the same cost formulations as in [13], thereby including yaw and pitch in the gradient based optimization of the trajectory. The frontier visibility cost favours orientations with more visible frontier voxels. To achieve this, the field of view is modelled using five half-spaces which are formulated as continuous, differentiable indicator functions suitable for numerical optimization.

Since the set of evaluated viewpoints $\mathcal{E}$ is not connected, it is not guaranteed that a traversable path to each viewpoint exists. In certain situations, we allow the path-planner to request a new goal, for example if the current goal is no longer traversable, if the only feasible path leads too far from the goal, or if too many re-planning attempts were necessary, suggesting that the goal is not reachable from the current configuration.
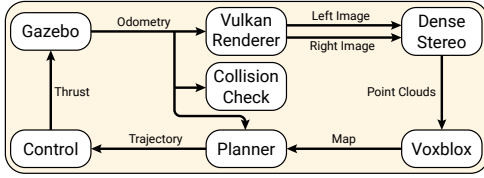
Fig. 4: Simulation overview for exploration experiments.

Max. velocity	$1\,\mathrm{m/s}$	Camera FoV	$80° \times 55°$
Max. acceleration	$1\,\mathrm{m/s^2}$	Camera resolution	752×480
Max. yaw rate	$\pi/2\,\mathrm{rad/s}$	Camera baseline	$0.11\,\mathrm{m}$
Max. sensor range	$5\,\mathrm{m}$	Fixed gimbal pitch	$0°$

TABLE I: Simulation parameters for all presented experiments.

G. Quality Improvement Phase

The exploration naturally finishes as soon as we can no longer generate new traversable viewpoints from any remaining surface frontiers. Instead of stopping the exploration entirely, we switch to a quality improvement phase. Here, the goal is to minimize the map uncertainty of the main surface that is reconstructed so far. Since the information gain formulation used already considers surface quality, we only need to adapt the viewpoint proposal step described in Section III-D.

Instead of proposing viewpoints from surface frontiers $\mathcal{F}$, we now generate viewpoints from surface voxels $m \in \mathcal{S}$ with a low weight $w(m) < w_{min}$. The minimum weight w_{min} is then incrementally increased as soon as all surface voxels below the current threshold are eliminated. To only consider relevant voxels, we generate viewpoints from voxels that are part of the largest reconstructed surface, suppressing unconnected voxels that were erroneously marked as surface voxels.

IV. EXPERIMENTS

All experiments are conducted in the Gazebo RotorS simulation framework [15] using ground truth odometry. The input point cloud used for mapping and planning is retrieved from a stereo image pair through fully dense stereo depth estimation [16]. The stereo camera is mounted on an actuated gimbal, which is able to pitch the camera. The full simulation pipeline is presented in Fig. 4. Our path-planner is compared against two state-of-the-art algorithms, AEP [9] and RRT IPP [8], in a selection of outdoor and indoor scenes with varying size and difficulty as shown in Fig. 5 and Fig. 6.

While our planner makes use of the actuated gimbal, the two state-of-the-art planners we compare against, AEP and RRT IPP, do not. To get comparable results, we also test against an adapted version of our planner with a fixed gimbal, where the viewpoint proposal step is adapted to only generate viewpoints with a fixed sensor pitch. Simulation parameters are listed in Table I.

A. Efficacy of Informed Sampling and APF Ranking

To test the efficacy of each part of our planner, we performed a number of experiments with different modifications on our planner. Namely, although we use the same information gain formulation proposed in [8], we can get different results

Method	Modification	Exploration 15 min [%]	Exploration 30 min [%]
Ours	Volumetric Frontiers	86.5 ± 5.7	$\mathbf{99.3 \pm 0.1}$
Ours	No APF	87.6 ± 5.6	97.9 ± 2.3
Ours	-	$\mathbf{91.7 \pm 5.1}$	$\mathbf{99.3 \pm 0.05}$

TABLE II: The percentage of the observable surface coverage after a fixed time in the *Bridge* scene for the proposed approach against two modified versions. The mean and standard deviation over six experiments is reported with best performance shown in bold.

Method	Sampling	Exploration 10 min [%]	Exploration 20 min [%]
Ours Fixed	Random $[x, y, z, \gamma]$	73.5 ± 7.0	95.2 ± 1.6
Ours Fixed	Random $[x, y, z]$	73.1 ± 7.6	94.9 ± 2.2
Ours Fixed	Informed	$\mathbf{85.5 \pm 3.3}$	$\mathbf{99.7 \pm 0.01}$
Ours	Random $[x, y, z, \gamma, \theta]$	71.3 ± 6.1	95.5 ± 2.9
Ours	Random $[x, y, z]$	80.7 ± 7.3	99.3 ± 0.2
Ours	Informed	$\mathbf{94.7 \pm 2.4}$	$\mathbf{99.8 \pm 0.03}$

TABLE III: The percentage of the observable surface coverage after a fixed time in the *City Building* scene for the proposed method with and without a fixed gimbal using either informed or random sampling. The mean and standard deviation over six experiments is reported with the best performance in bold.

if we use our definition of a surface frontier instead of the original volumetric frontier implementation. Thus, in the first modification we check that we do not lose any performance by substituting the volumetric frontiers in the original information gain formulation with our surface frontiers. The results in Table II show that there is no significant impact on exploration time in the *Bridge* scene.

To test the value of our APF ranking approach for view evaluation, in the second modification we replace the ranking with a randomized view selection. The potential benefit of a random view selection is the much faster computation, while the potential risk is evaluating too many irrelevant viewpoints, while not evaluating necessary ones. In Table II, we can see that this modification without the APF ranking cannot map the scene to the same extent as the original, unmodified planner despite the fact that the unmodified planner spends, on average, less computation time on information gain calculations, freeing up valuable resources.

Finally, we test our informed sampling approach against random sampling. These tests were performed in the *City Building* scene with ('Ours') and without the actuated gimbal ('Ours Fixed'). In these experiments, we replaced the informed viewpoint sampling described in Section III-D with random sampling. The tested random sampling approaches sample viewpoints globally with a bias to sample in a small region around the MAV's current position. Since the high dimensionality of the sampling space might adversely affect random sampling, we also test random sampling with yaw and pitch optimization. This reduces the sampling space to a position in $\mathbb{R}^3$ as in [8] and [9]. We find the optimal yaw and pitch from a segmented 360° ray-cast for each sampled position while the APF ranking and view evaluation steps are not altered.

The results in Table III show that informed sampling outperforms random sampling with and without the actuated gimbal. With a fixed gimbal pitch, there is no apparent difference in
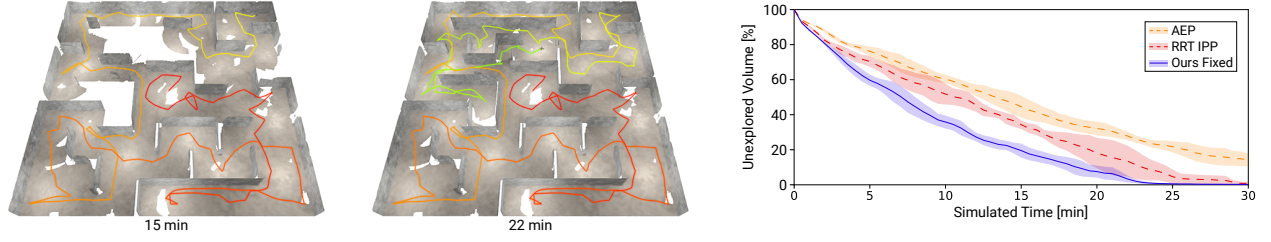
Fig. 5: We evaluate the *Maze* environment, introduced by [8], based on total unexplored volume over time (right). On the left we show the exploration progress and path of our planner after 15 and 22 minutes of exploration starting from the center of the maze. The path is color coded from start (red) to finish (green).
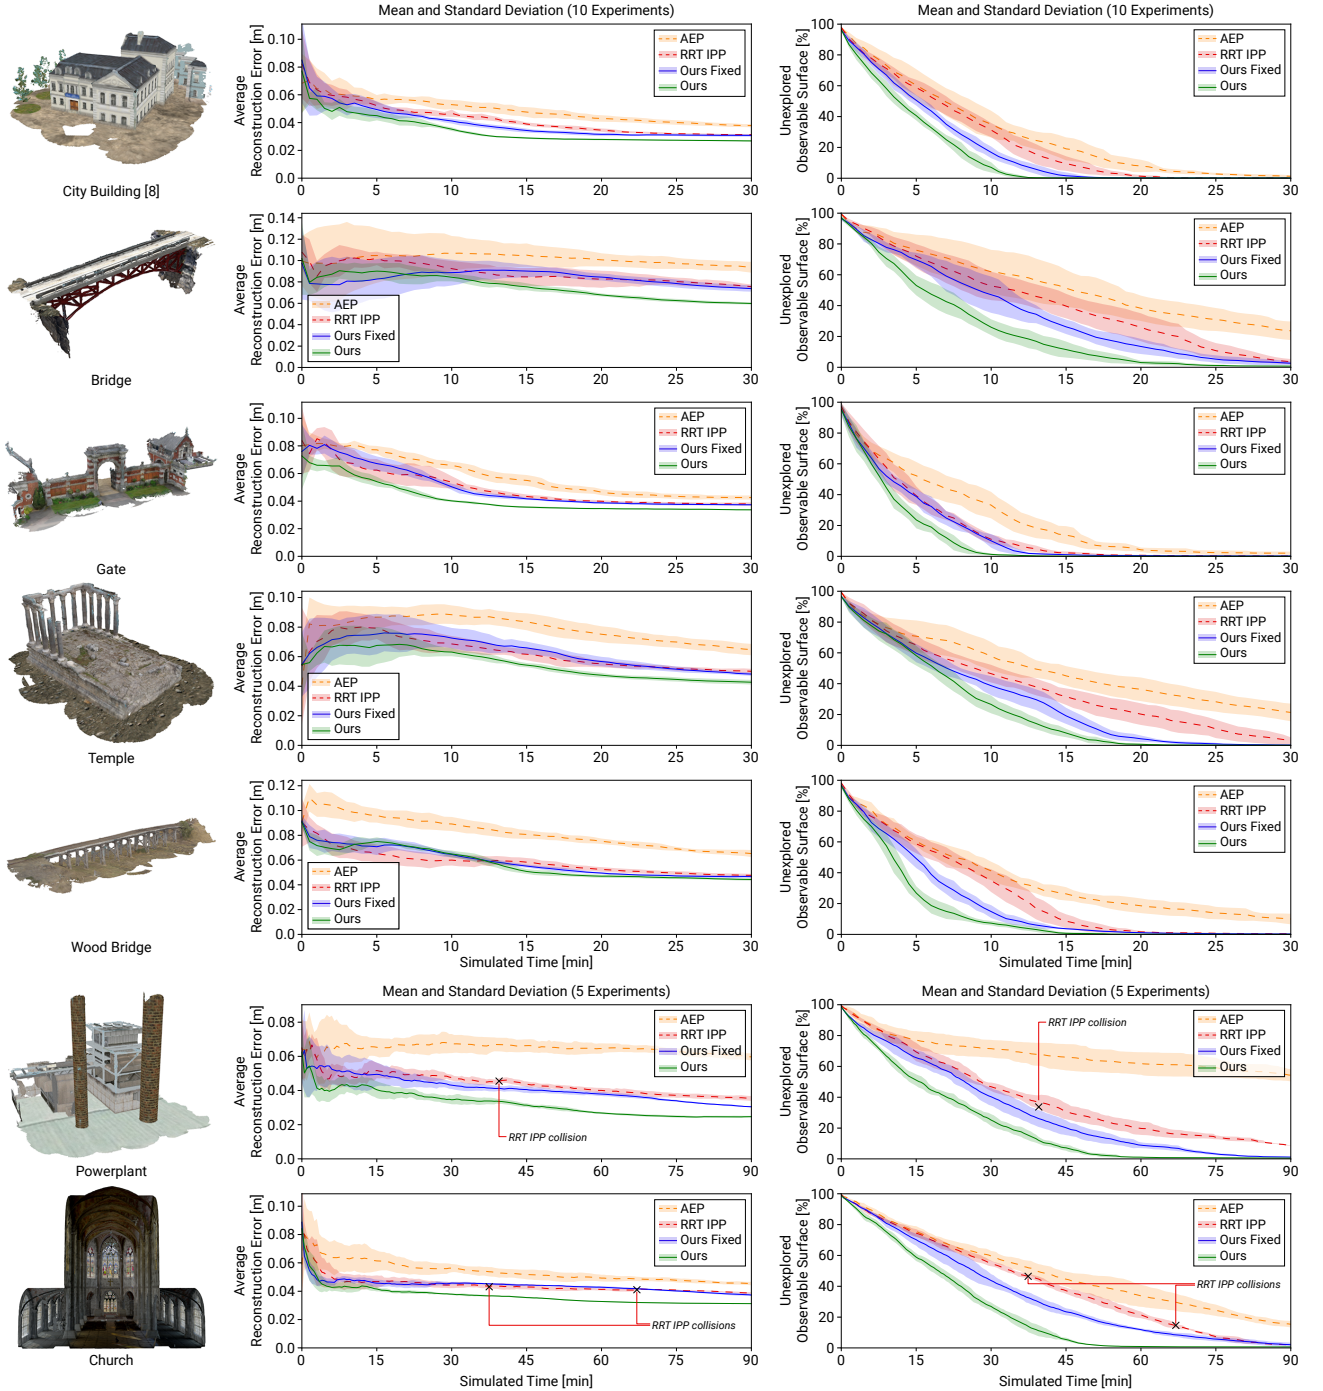

Fig. 6: Each scene depicted on the left gets evaluated with all planners achieving the illustrated average reconstruction error over time in the middle, and the percentage of unexplored observable surface over time on the right. We plot the mean (line) and standard deviation (shaded region) averaged over ten experiments. Evidently, both the proposed planner ('Ours') and the fixed-gimbal version ('Ours Fixed') shown for fairness of comparisons, achieve the complete reconstruction of all scenes faster and with consistently better quality than the state of the art.
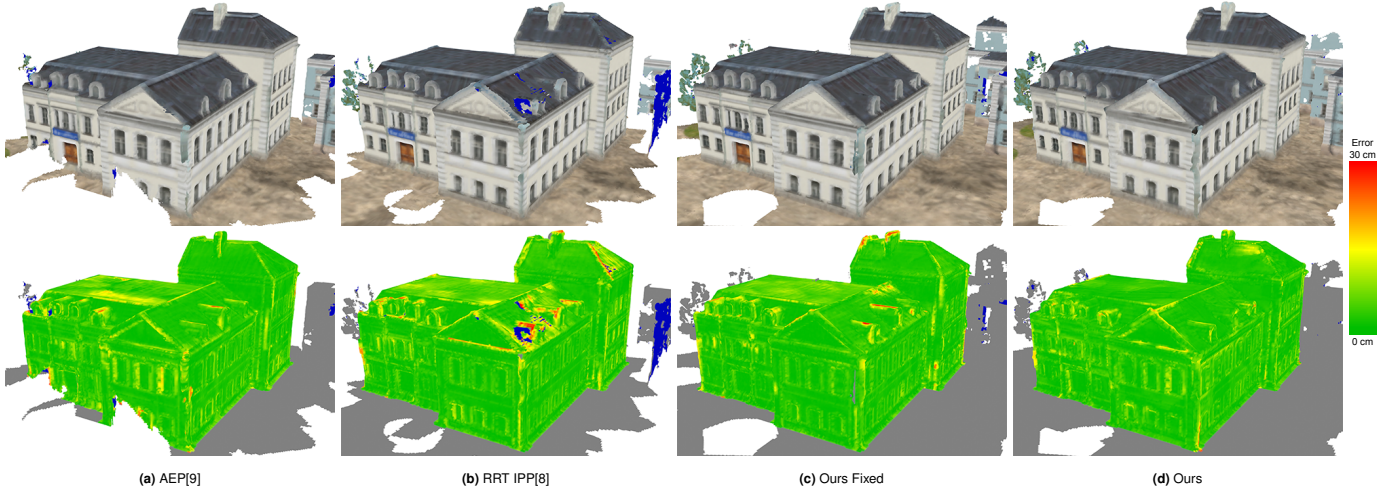
Fig. 7: Reconstruction results of the *City Building* scene for each planner after 20 minutes of exploration; the obtained meshes are shown in the top row, while the reconstruction error is visualised on the bottom. Out-of-bounds areas are marked grey, and back faces of meshes are marked blue.

Observation distances	$4.2\,\mathrm{m}, 2.5\,\mathrm{m}$	λ_1	2.0
Tangent rotation angles	$-\pi/4\,\mathrm{rad}, \pi/4\,\mathrm{rad}$	λ_2	1.0
Max. evaluated view age	$4\,\mathrm{s}$	λ_3	8.0
Threshold radius r	$3\,\mathrm{m}$	α	0.25

TABLE IV: Parameters for viewpoint sampling and path-planning.

performance between randomly sampled yaw and optimized yaw. If we include the additional degree of freedom from the gimbal pitch, we observe that the combined yaw and pitch optimization in the random sampling approach leads to higher coverage. Interestingly, the fully actuated, but randomly sampled method does not perform better than its counterpart with a fixed gimbal, despite the additional degree of freedom. Our gimbal actuated planner, on the other hand, is significantly faster than our fixed gimbal planner.

B. Volumetric Exploration Experiments

In a purely volumetric exploration task we are not interested in high quality surface reconstructions, but rather the complete exploration of a volume. Hence, we simulate AEP [9], RRT IPP [8] and our fixed planner in the indoor *Maze* environment. Since the focus is on exploration and not on surface quality, we use the unexplored volume as the main driver for all the planners. Therefore, we use volumetric frontiers for informed sampling by omitting the conditions imposed by (3), and we do not employ the actuated gimbal. Consequently, in Fig. 5 we report the evaluation based on the total explored volume over time. Our planner can quickly map the entirety of this hard-to-navigate scene.

C. 3D Reconstruction Experiments

We simulate AEP [9], RRT IPP [8], Ours Fixed, and Our planner in seven indoor and outdoor scenes. We used a common voxel size of $d_v = 0.1\,\mathrm{m}$ across all scenes. In Table IV we list the parameters used for our planners. The experiments were terminated after a fixed amount of time, depending on the size of the environment. The bounding boxes used to limit the exploration were slightly inflated for AEP, since we have discovered that a tight bounding box can limit

AEP's ability to explore certain regions fully. Every planner was simulated multiple times in every scene, starting at a chosen selection of different initial positions and orientations. In Fig. 6 we present a comprehensive overview of the results in each scene illustrating the percentage of unexplored observable surface achieved and the average reconstruction error recorded over time.

Our fixed gimbal planner consistently outperforms both state-of-the-art planners, with an even greater advantage if we make use of the actuated gimbal. Indicatively, both 'Ours' and 'Ours Fixed' manage to reconstruct 95% of the *Temple* scene before any state-of-the-art planner reaches a surface coverage of 80%. The results show that the average reconstruction error, measured at a specific extent of scene coverage, are relatively similar across all planners, with a slight advantage for our fully actuated planner, while our planners explore much faster than state-of-the-art planners. This means that the faster reconstruction times from our planners do not have a negative impact on the reconstruction quality. On the contrary, our fixed-gimbal planner is able to produce a reconstruction with a lower average error compared to state-of-the-art methods, and is only outperformed by our fully actuated planner. AEP's purely volumetric information gain might be the root cause of its worse performance, as it explores the largest overall volume (e.g. in the *Temple* scene our planner and RRT IPP explore a 14.6% and a 19.4% smaller volume than AEP, respectively).

In order to analyse the reconstructions of the *City Building* and the *Bridge* in more detail the results are illustrated in Fig. 7 and Fig. 8, respectively. Notably, both of our planners manage to reconstruct the *City Building* without any significant defects in the resulting mesh. They can establish the initial reconstruction relatively quickly, leaving enough time to enter the quality-improvement phase and allowing them to re-visit surface voxels with low confidence. In the larger *Bridge* scene, our planners reach a similar level of surface coverage, although with a larger average reconstruction error. This can be explained by the much finer details present in this scene, with some geometry being thinner than the chosen voxel size.

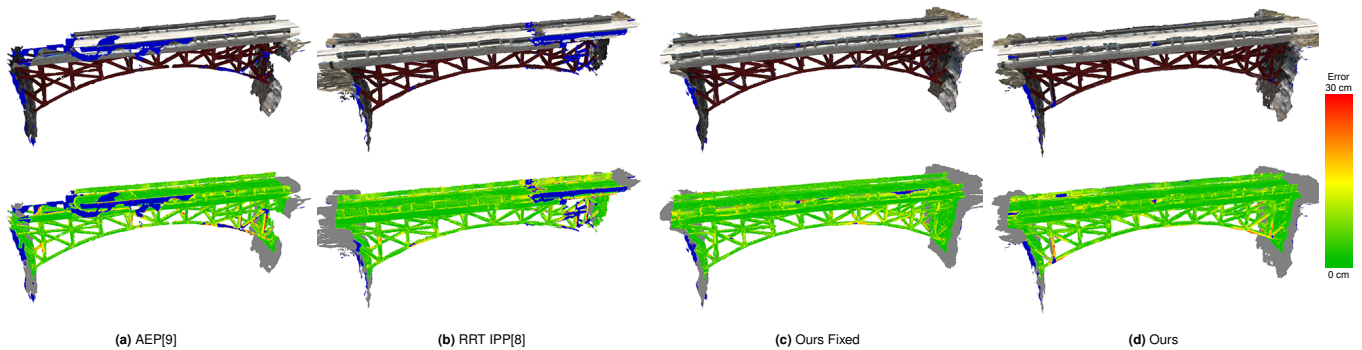The main difference in the final result between our planners

Fig. 8: Reconstruction results of the *Bridge* scene for each planner after 30 minutes of exploration; the obtained meshes are shown in the top row, while the reconstruction error is visualised on the bottom. Out-of-bounds areas are marked grey, and back faces of meshes are marked blue.

lies in the reconstruction quality with our gimbal-actuated planner having overall fewer problematic regions of high reconstruction error. Additionally, our gimbal-actuated planner is the only planner that manages to reconstruct the underside of the *Bridge* completely (not visible in Fig. 8, but visible in the accompanying video). This can be explained by the fixed pitch angle, which makes some surfaces below the bridge unobservable for the other planners.

AEP is not able to fully reconstruct the *City Building* nor the *Bridge* within the time limit. However, the *City Building* is almost fully mapped by RRT IPP with only a few areas of the roof missing and a few areas with a relatively large reconstruction error remaining. In the *Bridge* scene, RRT IPP leaves only a small area unexplored.

In summary, our planner outperforms the state-of-the-art planners both in reconstruction time as well as reconstruction quality. The additional degree of freedom from the actuated gimbal used by our planner seems to further improve the lead in both metrics over state-of-the-art methods and our fixed gimbal implementation.

V. CONCLUSION AND FUTURE WORK

Inspired by the need for effective exploration strategies for 3D reconstruction using small aircraft, this letter presents an online exploration path-planner for MAVs with a stereo-based depth sensor on an actuated gimbal. In particular, surface frontiers in the current reconstruction are used for informed viewpoint sampling, resulting in increased quality of the sampled and evaluated viewpoints that the MAV can travel to in order to complete the reconstruction of a structure of interest. Employing a perception-aware path-planner to guide the MAV to the next best viewpoint as evaluated at each step is shown to lead to shorter exploration times. In fact, comparisons on a testbed of 8 challenging photo-realistic scenes reveal consistently faster, more complete and more accurate reconstructions than the state of the art.

Future work will explore integrating a mapping framework that can deal with drifting odometry to improve the fidelity of navigation processes.

REFERENCES

[1] S. Weiss, M. W. Achtelik, S. Lynen, M. C. Achtelik, L. Kneip, M. Chli, and R. Siegwart, "Monocular Vision for Long-term MAV Navigation: A Compendium," *Journal of Field Robotics (JFR)*, vol. 30, 2013.

[2] L. Heng, D. Honegger, G. H. Lee, L. Meier, P. Tanskanen, F. Fraundorfer, and M. Pollefeys, "Autonomous visual mapping and exploration with a micro aerial vehicle," *J. Field Robot.*, vol. 31, no. 4, 2014.

[3] T. Cieslewski, E. Kaufmann, and D. Scaramuzza, "Rapid exploration with multi-rotors: A frontier selection method for high speed flight," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2017.

[4] L. Yoder and S. Scherer, *Autonomous Exploration for Infrastructure Modeling with a Micro Aerial Vehicle*. Springer International Publishing, 2016, pp. 427–440.

[5] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon "next-best-view" planner for 3d exploration," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016.

[6] R. Border, J. D. Gammell, and P. Newman, "Surface edge explorer (see): Planning next best views directly from 3d observations," *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018.

[7] S. Song and S. Jo, "Surface-based exploration for autonomous 3d modeling," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, May 2018.

[8] L. Schmid, M. Pantic, R. Khanna, L. Ott, R. Siegwart, and J. Nieto, "An efficient sampling-based method for online informative path planning in unknown environments," *IEEE Robotics and Automation Letters*, vol. 5, 2020.

[9] M. Selin, M. Tiger, D. Duberg, F. Heintz, and P. Jensfelt, "Efficient autonomous exploration planning of large-scale 3-d environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, 2019.

[10] A. Dai, S. Papatheodorou, N. Funk, D. Tzoumanikas, and S. Leutenegger, "Fast frontier-based information-driven autonomous exploration with an mav," *2020 IEEE International Conference on Robotics and Automation (ICRA)*, May 2020.

[11] B. Yamauchi, "A frontier-based approach for autonomous exploration," in *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. 'Towards New Computational Principles for Robotics and Automation'*, 1997.

[12] S. Song and S. Jo, "Online inspection path planning for autonomous 3d modeling using a micro-aerial vehicle," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017.

[13] L. Bartolomei, L. Pinto Teixeira, and M. Chli, "Perception-aware path planning for UAVs using semantic segmentation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2020)*, 2020.

[14] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.

[15] F. Furrer, M. Burri, M. Achtelik, and R. Siegwart, *RotorS – A Modular Gazebo MAV Simulator Framework*. Springer, 2016, vol. 625, pp. 595–625.

[16] H. Hirschmüller, "Stereo processing by semiglobal matching and mutual information," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 30, no. 2, 2008.