



HAL
open science

FLYBO: A Unified Benchmark Environment for Autonomous Flying Robots

Anthony Brunel, Amine Bourki, Olivier Strauss, Cédric Démonceaux

► **To cite this version:**

Anthony Brunel, Amine Bourki, Olivier Strauss, Cédric Démonceaux. FLYBO: A Unified Benchmark Environment for Autonomous Flying Robots. 3DV 2021 - 9th International Conference on 3D Vision, Dec 2021, Virtual, United Kingdom. pp.1420-1431, 10.1109/3DV53792.2021.00149 . hal-03380109

HAL Id: hal-03380109

<https://hal.science/hal-03380109v1>

Submitted on 15 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

FLYBO: A Unified Benchmark Environment for Autonomous Flying Robots

Anthony Brunel^{1,2}, Amine Bourki², Olivier Strauss¹ and Cédric Demonceaux³

¹LIRMM, Univ. Montpellier, CNRS, 860 rue de St Priest, 34095 Montpellier, France.

²Gambi-M R&D, Champs-sur-Marne, France.

³ImViA, Université Bourgogne Franche-Comté (UBFC), Le Creusot, France.

{anthony.brunel, olivier.strauss}@lirmm.fr, amine.bourki@gmail.com,

cedric.demonceaux@u-bourgogne.fr

www.flybo.org




Figure 1: An MAV equipped with odometry- and active depth sensors autonomously explores a complex synthetic area from FLYBO (a) while gradually mapping the scene throughout different exploration stages and planning trajectories online (b–d). Simultaneously, the perceived surfaces are also reconstructed online (close-up views). FLYBO provides datasets, references and a framework to benchmark such systems w.r.t their volumetric exploration and online surface reconstruction capabilities.

Abstract

The use of Micro-Aerial Vehicles (MAVs) equipped with odometry- and depth sensors has become predominant for a wide variety of challenging industrial applications such as the autonomous exploration (i.e., digital mapping), and inspection (i.e., online surface reconstruction) of unknown facilities. However, despite the ongoing attention these topics receive, autonomous exploration systems still lack common evaluation grounds to assess their relative performance in terms of data and experimental tools. We address this deficit by introducing FLYBO, the first unified benchmark environment that focuses on the performance of such flying robots in terms of autonomous exploration and online surface reconstruction. It includes (i) 11 challenging realistic indoor- and outdoor datasets of increasing complexity and size, with ground-truth, (ii) a comprehensive benchmark of 7 of the top-performing autonomous exploration algorithms including methods without publicly available code. (iii) A unified experimental system factorizes the routines shared by autonomous planners in order to fairly and accurately assess their exploration performance in a controlled environment.

1. Introduction

Over the past decade, Micro-Aerial Vehicles (MAVs) have considerably gained in popularity and commercial maturity, giving birth to even consumer-grade products at an affordable price. In particular, MAVs equipped with depth- and odometry sensors have become a staple equipment to help map, digitize, locate and reconstruct relatively intricate facilities, even in absence of prior knowledge about their structure, or comprising objects [1]. As a result, the ability of such autonomous flying robots to *efficiently explore* unknown environments [8, 6, 65] while *accurately reconstructing surfaces* [63] that are perceived therein has become critical to help minimize risks and operational costs.

Autonomous exploration planning has been a long-standing problem for over two decades [78, 79] and has received an ongoing attention within the scientific communities of Computer Vision [33], Robotics [6, 14] and Computer Graphics [23, 77]. Despite this activity, state-of-the-art autonomous planning systems (commonly referred-to as “planners”) are evaluated on disparate data, in sometimes drastically diverse experimental conditions [65, 19], and overall, lack a unifying public benchmark.




Figure 2: Overview of the 11 areas we consider in FLYBO, through a few selected views and levels.

In terms of datasets, the evaluation of autonomous exploration and their online surface reconstruction capabilities requires 3D models that are (i) watertight (to avoid aberrant flight behaviors, especially indoors) and (ii) realistic in terms of their structural complexity, size and geometry. CAD-derived data in particular are a gold-standard w.r.t these criteria and are favored when available as a benchmarking reference in a controlled, simulated-environment for autonomous exploration [8, 19], and surface reconstruction [47, 43].

We address the aforementioned challenges by making the following main contributions:

- *System:* We present FLYBO, the first unified benchmark environment specifically tailored for the evaluation of autonomous flying MAVs for the tasks of volumetric exploration and online surface reconstruction.
- *Datasets:* The package includes 11 challenging CAD-derived datasets with ground-truth, that represent realistic indoor- and outdoor areas of varying nature, size and structural complexity.
- *Benchmark:* We evaluate 7 of the leading state-of-the-art autonomous exploration planners, including methods with no public code using our own re-implementations.
- *Tools:* In order to fairly and accurately assess the pure exploration planning part of existing systems, we provide an experimental framework that factorizes the common routines that are shared by all planners. Our system also allows to generate the quantitative and qualitative results seamlessly and on the fly.
- A project website with a leaderboard will be published upon acceptance along the datasets, our re-implementations, and unified evaluation system.

2. Related Work

Autonomous Exploration Planning. Autonomous exploration planning algorithms are commonly categorized into frontier-based, sampling-based or hybrid strategies. Frontier-based methods rely on the concept of frontier voxels that separate *empty* and *unknown* spaces in the explored

environment. The first method to introduce this paradigm was Classic [78] and has originally done so in 2D exploration planning. The method proceeds by visiting the nearest available frontier at each algorithm iteration, until all achievable frontiers are visited. Despite its apparent simplicity, the method also performs very well in 3D and is still considered relevant as a baseline system in terms of 3D exploration efficiency [8, 82]. The seminal work of Yamauchi et al. [78] has led to numerous extensions including multi-agent exploration [79], and more sophisticated strategies to extract frontiers using stochastic differential equations [67], information theory and a multi-objective utility function [21], or more generally, in a Next-Best View (NBV) fashion, i.e., by maximizing a given utility function greedily [24]. Another trend in frontier-based planning is to enhance its main original drawbacks of speed and efficiency. Rapid [14] addresses these shortcomings by alternating between the original behavior of Classic [78] and minimizing the deviation angle between the current trajectory of the MAV and its incidence with frontier candidates. This simple heuristic has led to impressive gains in terms of volumetric efficiency and is still considered a strong-performing planner w.r.t current state-of-the-art [82, 14]. The current main difficulty to assess the performance of Classic [78] and Rapid [14] outside the scope of their originally documented performances lies in the absence of their public implementations. We address this by providing common, efficient implementations of their systems.

Sampling-based solutions are the second dominant trend in autonomous exploration. They are based on sampling sparse viewpoint candidates to visit in order to limit the size of the search space, and ultimately, the volumetric efficiency of the process altogether. Such approaches often rely on tree data structures that connect the sampled positions and allow to produce feasible trajectories through queries [42, 49]. The top-performing sampling strategies are based on Rapidly-exploring Random Trees (RRT*) [49] and develop one [63] or multiple tree structures [6, 65] from randomly sampled seed positions. The motivating theoretic

	MAZE	OFFICES A	OFFICES B	OFFICES C	FACILITY A	FACILITY B	FACILITY C	WAREHOUSE	PLATFORM A	PLATFORM B	POWERPLANT
Area (m^2)	400	470	120	534	115	323	478	800	384	480	1000
Volume (m^3)	1000	1410	1443	1602	690	1450	2868	9600	3456	12000	26598
Bounding volume (m)	$20 \times 20 \times 2.5$	$22 \times 28 \times 3.4$	$8 \times 20 \times 12$	$30 \times 16 \times 7.4$	$12 \times 14 \times 6$	$19 \times 17 \times 4.6$	$38 \times 22 \times 6$	$40 \times 20 \times 12$	$16 \times 24 \times 9$	$24 \times 20 \times 25$	$33 \times 31 \times 26$
#Rooms	-	2	4	23	1	2	5	1	-	-	-
#Levels	1	1	4	2	1	1	2	2	3	3	1
#CAD mesh faces	130	6910	227k	857k	373k	1120k	1072k	793k	149k	470k	2000
Structural complexity	Low	Low	Intermediate	High	High	High	High	High	Intermediate	High	Low
Scale	Small	Intermediate	Intermediate	Large	Intermediate	Small	Large	Large	Large	Large	Large
Indoor / outdoor	Indoor	Indoor	Indoor	Indoor	Indoor	Indoor	Indoor	Indoor	Outdoor	Outdoor	Outdoor

Table 1: Statistics of the 11 Areas that comprise the dataset part of FLYBO.

cal purpose behind this paradigm lies in the scalability and ease of convergence although the early strong-performing methods of the kind such as NBVP [6] tend to get stuck in local exploration. In contrast, AEP [65] combines the NBVP [6] strategy with the storage of previously-explored node positions and re-assessing their information gain as well as considering frontiers, leading to a hybrid strategy that has a sparing effect regarding globally-exploring views. More recently, SplatPlanner [8] proposed a hybrid scheme that combines the flexibility of RRT* with frontiers that are sampled based on their spatial density using highly-efficient bilateral filtering operations, leading to impressive performance in terms of volumetric efficiency and coverage.

The task of autonomous exploration is inherently tied to the one of autonomous inspection, or online surface reconstruction in assessing the accuracy of the geometry that is produced by autonomous planners [37]. Several works jointly consider criteria for online surface reconstruction in a mono-agent scenario [37, 77, 63] or even as multi-agent solutions [23]. Despite the existence of a rich but much more specialized literature on reconstruction in general [5, 43] and online surface reconstruction in particular [68, 37], our proposed benchmark strictly focuses on the online surface reconstruction of general autonomous exploration systems in order to provide new insights regarding their performance beyond the traditional volumetric criteria, as a *trade-off* between their volumetric efficiency and the accuracy of their reconstructed surfaces in a controlled environment. This is made possible thanks to the unified nature of our FLYBO benchmarking framework.

Datasets and Benchmarks for Autonomous Flying Robots. Current state-of-the-art autonomous planners commonly evaluate their performance either on a handful of disparate, non-public CAD datasets [8, 82], or by resorting to borrowing datasets from other tasks [23, 69]. Such datasets are RGBD-derived [9, 17, 3, 69, 76], which can suffice to evaluate autonomous planners embedded on non-flying robots [77, 23], but are however impractical for simulating flying robots that can fly through unscanned regions and are also far from ideal to evaluate surface reconstruction in presence of sometimes severe structural artifacts and missing parts, that can also bias frontier-based explorations.

In contrast, FLYBO proposes watertight and highly-structured CAD-derived environments, comprising the first public dataset and benchmark focusing on MAV-based autonomous exploration systems.

3. Flying through FLYBO

3.1. Dataset and Acquisition

Our proposed benchmark is comprised of **11 areas**, as detailed in Table 1 and illustrated in Figure 2. They consist of CAD models that were manually designed by human experts and directly serve our experimental framework to evaluate state-of-the-art methods for the considered tasks.

- **MAZE.** This is the most structurally simple area of our dataset. It allows to isolate exploration behaviors on the simplest form of indoor scenes, i.e., consisting of one ground level and flat walls. The performance on this area also serves as a baseline to assess the robustness to increasing levels of structural complexity and scene size.
- **OFFICES A.** Incrementally more complex than MAZE, it depicts two interconnected furnished office areas of moderately-sized open-spaces with few objects.
- **OFFICES B.** Another set of interconnected office areas with more geometric complexity spanning multiple levels and that contain an important number of present objects.
- **OFFICES C.** Similarly, area C is much more challenging in terms of structural complexity, amount of objects and details as well as in terms of overall scale.
- **FACILITY A, B AND C.** These facilities depict complex industrial indoor environments, comprising piping, ducts, complex equipments and very challenging intricate spaces to move through. FACILITY A–C are of high structural complexity and of increasing sizes and scales. They span the plausible usecases that autonomous exploration algorithms would typically have to deal with in a real-life inspection scenario, ranging from moderate difficulty to highly-challenging overall.
- **WAREHOUSE** is a large-scale area that contains a large central empty space that is surrounded by many complex equipments that span two interconnected levels.

- **PLATFORM A AND B** represent large industrial outdoors with increasing amounts of cluttering objects and complex equipments. They present numerous intricate spaces that are difficult to maneuver through for an MAV.
- **POWERPLANT** was originally available as an open-source model and used by several state-of-the-art planners to evaluate performance [8, 20, 14]. The area represents a very large-scale powerplant for which the exact crop considered for evaluation was undisclosed and potentially different. We address this by providing one fixed crop to this model with its corresponding experimental settings.

3.2. Tasks and Metrics

The proposed benchmark is specifically tailored for autonomous exploration planning systems, a.k.a planners. However, given that the most common cause of failure in MAV-based inspection campaigns generally is poorly reconstructed geometry provided by the offline, *a posteriori* assessment of the mapped 3D data, we consider both the tasks of (i) pure autonomous exploration planning per se, and (ii) online surface reconstruction to evaluate such systems. In practice, the online surface reconstruction may often constitute the deciding factor of whether an MAV exploration campaign is deemed successful in a production-ready scenario, sometimes regardless of the efficiency and speed of the exploration process.

Autonomous Exploration Planning. Given $\mathcal{V} \subset \mathbb{R}^3$ a voxel-grid that represents the occupancy states of voxels initially labeled as ‘unknown’, the objective is to re-label such entities as ‘free’ or ‘occupied’ while computing collision-free trajectories online. The standard metrics [23, 8, 19, 65, 14] we consider are defined as follows:

- The volumetric efficiency of the explorations in m^3/s ,
- The volumetric coverage that we define as a fraction of the covered volume at a fixed exploration plateau relative to the top score, for simplification.

Online Surface Reconstruction. We adopt the standard metrics in the field for general-purpose surface reconstruction [23, 46, 43], namely:

- The surface coverage, expressed as the recall over time,
- The accuracy of the produced meshes as F-score and RMSE w.r.t ground-truth references [23, 46].

All of the aforementioned metrics are averaged over multiple runs (5 throughout our experiments) to account for the inherently stochastic nature of the task. This is also a common practice in the field [8, 19, 65, 14].

3.3. Ground-truth Generation

To evaluate autonomous planners according to the aforementioned tasks, we consider the following data:

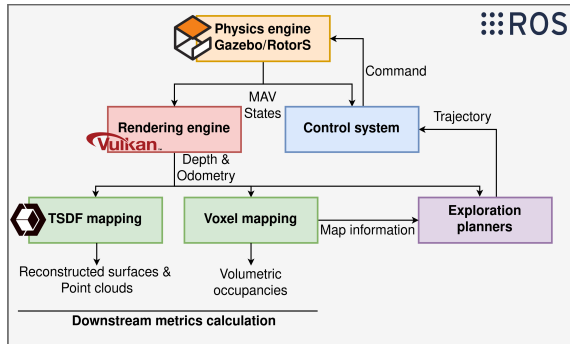


Figure 3: Schematic workflow and components of our proposed unified experimental framework.

Autonomous Exploration Planning. For all the considered metrics, the reference data are the input CAD models.

Online Surface Reconstruction. In this particular case, the direct use of CAD models is not adequate because the surface-based metrics such as recall and F-score account for the amount of geometric surfaces that can be reconstructed, hence perceived in a scene. To avoid introducing significant biases in our evaluation, we run an interactive exploration mapping mode without simulated noise that is manually controlled. This allows to recover only the exterior, perceivable parts of the CAD scenes and to store them as highly accurate and dense point clouds to compute the considered metrics w.r.t surface reconstruction.

3.4. Unified Simulation Framework

We propose the first unified experimental simulation framework tailored for MAV-based autonomous exploration planners (Figure 3). It factorizes the common simulation routines shared by all systems, and the generation of intermediate data outputs and the computation of downstream metrics in a seamless, unified way.

The proposed simulation framework builds upon the Robot Operating System (ROS Kinetic) library [61]. It simulates a Hummingbird quadrotors MAV model from Ascending Technologies and its physics are modeled by the popular Gazebo-based RotorS simulator [29]. The control is insured by a flying robots control framework [27]. Traditional collision-checking is ensured within a sphere centered at the position of the MAV. The computed trajectories are typically piecewise-linear and are smoothed out with a polynomial trajectory optimization [62] to send graceful paths to the control system. The rendering engine uses the Vulkan API¹, assuming perfect poses from RotorS. It simulates a 20Hz DepthVision camera synchronized with the odometry. We take into account depth and odometry controlled noise, similar to ESM [63] we use a gaussian depth error with quadratic scaling modeling recent depth

¹<https://www.khronos.org/vulkan/>

camera [44]. The system simulates noise in positions and depthmaps for realism. TSDF [38] and point clouds are continuously generated as well as the final meshes.

3.5. Featured Autonomous Exploration Systems

Our benchmark experiments consider 7 methods from state-of-the-art featuring the following systems: Classic [78], Rapid [14], NBVP [6], AEP [65], ESM_{IPP} [63], ESM_{Recon} [63] and SplatPlanner [8]. We have re-implemented Classic [78], Rapid [14] and SplatPlanner [8] in absence of public implementations for these methods. NBVP [6], ESM_{IPP} [63], ESM_{Recon} [63] and AEP [65] are used based on their public implementations. Finally, we integrate all of the considered systems in our evaluation framework. NBVP [6] and AEP [65] are inherently dependent on Octomap [35] and thus, they can not benefit from our continuous voxel mapping component (Figure 3).

4. Results

4.1. Autonomous Exploration

Figure 4 and Table 2 respectively summarize the benchmark results in terms of volumetric exploration efficiency averaged over 5 runs with min-max performance and volumetric coverage also averaged over the 5 same runs.

Volumetric efficiency. The first general trend that emerges from the reported data is that certain methods maintain a steady performance throughout the dataset areas. This is the case for SplatPlanner [8], AEP [65], NBVP [6], ESM_{IPP} [63] and ESM_{Recon} [63].

In contrast, Classic [78] and Rapid [14] seem to have sometimes severe performance drops on specific area types. Classic [78] seems to struggle on all the larger scenes (e.g., FACILITY C, WAREHOUSE, PLATFORM A–B, POWERPLANT). This can be explained by the fact that its exhaustive visit of nearest frontiers is linearly correlated to the increased input scale, hence the efficiency drops.

Rapid [14] on the other hand tends to suffer in presence of higher levels of structural complexity, such as indoor scenes with narrow, intricate spaces such as OFFICE C and FACILITY B. On such scenes, the planner tends to oscillate between the traditional Classic [78] heuristic, and its characteristic maximization of the MAV velocity. In particular, while geometric complexity increases, many frontiers are unfavorably oriented w.r.t to the current MAV trajectory.

Volumetric coverage. ESM_{IPP} [63] is a clear winner per the average relative scene coverage, followed by SplatPlanner [8] and ESM_{Recon} [63]. NBVP [6] and AEP [65] perform the worst among the considered planners and Classic [78] shows competitive performance rank-wise, despite its algorithmic simplicity, and performs especially well on smaller areas like MAZE, OFFICES A–C and FACILITY A–C. Rapid [14] shows competitive performance throughout.

4.2. Online Surface Reconstruction

Figure 5 and Table 3 respectively summarize the results regarding the surface coverages and accuracies of the final produced meshes. Both of these scores are also averaged over the same 5 runs and Figure 5 also shows the min-max scores per planner (as colored intervals).

Coverage. Regarding this metric, Classic [78] performs unexpectedly well throughout the areas, except for larger ones similarly to the reported performance drop in volumetric efficiency in such environments. NBVP [6] performs the worse overall, and Rapid [14] performs reasonably well throughout. The top-performing systems are ESM_{IPP} [63] and SplatPlanner [8].

Accuracy of the produced meshes. The reported scores for this metric are the most surprising ones. ESM_{IPP} [63] claims the first place in terms of RMSE while ESM_{Recon} [63] is the better one w.r.t F-scores, closely followed by Rapid [14] and SplatPlanner [8]. It is also worth mentioning that AEP [65] performs worse than the original NBVP [6] method in terms of RMSE and is almost as bad w.r.t F-score. This can be explained by the fact that AEP [65] induces a trade-off between local (NBVP [6] - like) and global (frontier-like) exploration behaviors, which seems to disserve it, accuracy-wise.

4.3. Volumetric Exploration vs. Surface Reconstruction

The Pareto-optimal planners on FLYBO are ESM_{IPP} [63], ESM_{Recon} [63] and SplatPlanner [8] that are the respective leaders w.r.t volumetric exploration (coverage) and surface reconstruction (F-score), while SplatPlanner [8] emerges as the best compromise between the two task-representative metrics (Figure 6).

4.4. Additional Statistics

Table 4 provides indicative timings and statistics.

5. Discussion and Perspectives

The impact of the present work can potentially lead to several avenues of extensions and future work, a few of which are specified in the following discussion.

Given the structured, unified nature of FLYBO, the proposed data and framework can be used to support the development of Reinforcement-Learning (RL) based approaches to autonomous exploration. Methods that build on RL typically require a moderate amount of training data [13] and an efficient means of scaling the number of experiments which could naturally benefit from the flexibility of our system.

An important component of autonomous planning lies in the generation of efficient, continuous, graceful trajectories [22, 55, 80]. Such methods as well as replanning strategies [55, 81, 82] could also directly benefit from our frame-

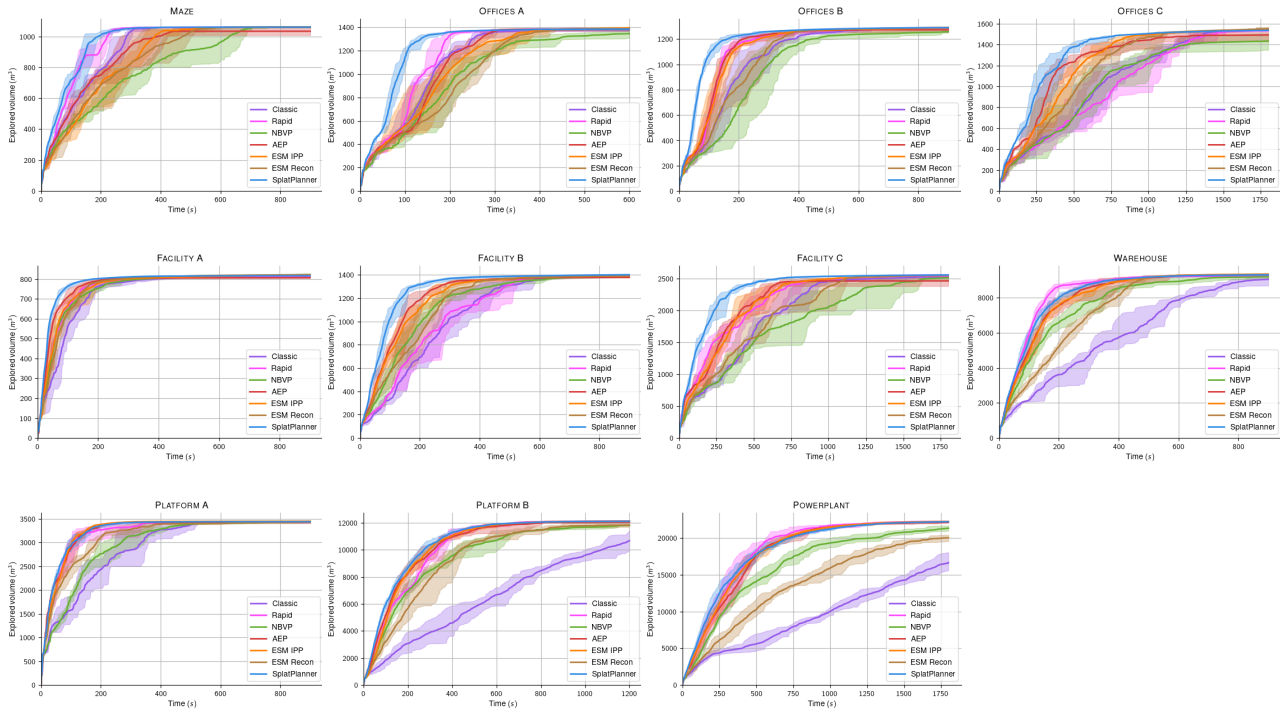


Figure 4: **Autonomous Exploration: Volumetric Efficiency** – Explored volumes over time, averaged over 5 runs.

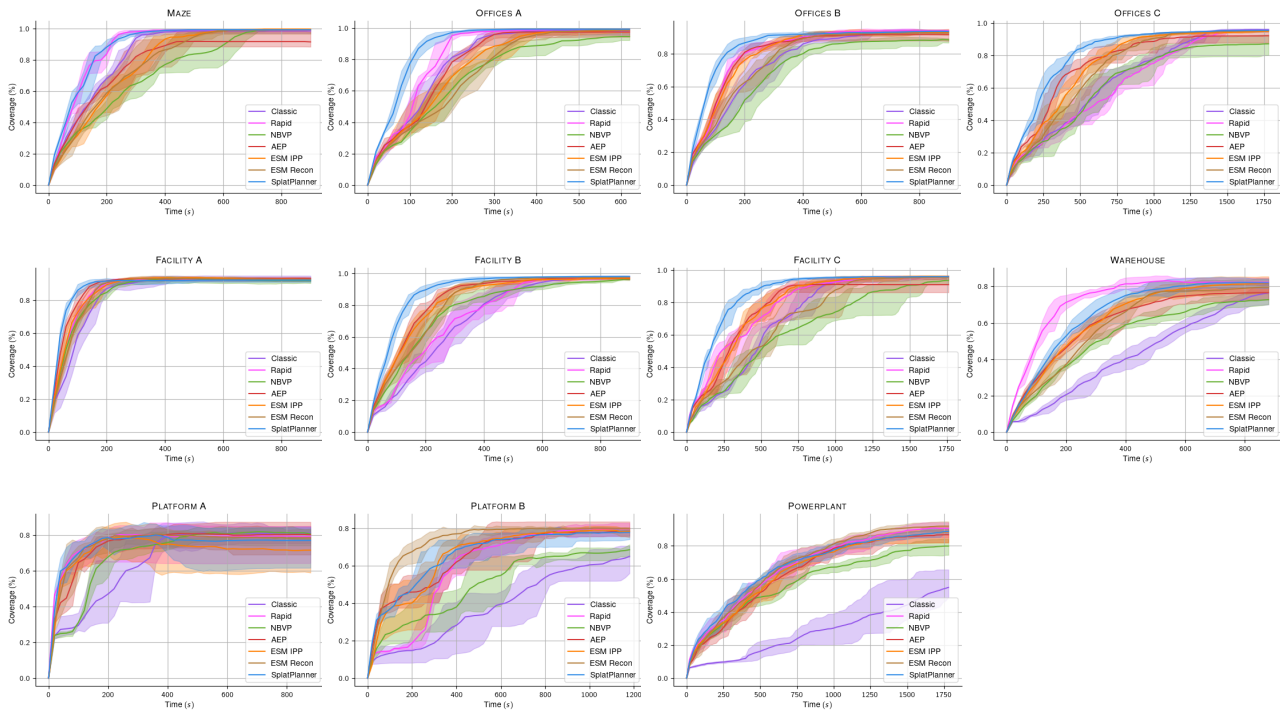


Figure 5: **Online Surface Reconstruction: Coverage** – Scene surface recall over exploration time, averaged over 5 runs.

	MAZE	OFFICES A	OFFICES B	OFFICES C	FACILITY A	FACILITY B	FACILITY C	WAREHOUSE	PLATFORM A	PLATFORM B	POWERPLANT	Mean	Rank
Classic [78]	0.9992	0.9945	0.9957	0.9926	0.9868	0.9911	0.9883	0.9681	0.9933	0.8795	0.7484	0.9579	5.3
Rapid [14]	0.9988	0.994	0.9987	0.9912	0.99	0.9932	0.995	0.9861	0.9934	0.9995	0.9999	0.9945	3.7
NBVP [6]	0.9925	0.9648	0.9685	0.9258	0.9973	0.9866	0.9774	0.9834	0.9887	0.9748	0.9587	0.9744	5.8
AEP [65]	0.975	0.9911	0.9882	0.9656	0.9836	0.9863	0.9642	0.9962	0.9939	0.993	0.9945	0.9847	5.5
ESM _{IPP} [63]	0.9997	1.0	0.9962	0.9976	0.9961	0.9925	0.9971	1.0	1.0	1.0	1.0	0.9981	2.2
ESM _{Recon} [63]	1.0	0.9953	0.9988	1.0	1.0	0.9927	0.9974	0.9974	0.9863	0.9758	0.9001	0.9858	2.9
SplatPlanner [8]	0.9998	0.9923	1.0	0.9935	0.9959	1.0	1.0	0.9952	0.9976	0.999	0.9984	0.9974	2.6
Plateau Time (s)	700	600	750	1650	700	750	1650	900	550	1200	1800		

Table 2: **Autonomous Exploration: Volumetric Coverage** – Scores at plateau, relative to the best value (5 run average).

	MAZE	OFFICES A	OFFICES B	OFFICES C	FACILITY A	FACILITY B	FACILITY C	WAREHOUSE	PLATFORM A	PLATFORM B	POWERPLANT	Mean	Rank
Classic [78]	0.9807	0.9792	0.9371	0.9661	0.9266	0.9459	0.9629	0.7781	0.8111	0.6738	0.6694	0.8755	3.9
Rapid [14]	0.9831	0.9905	0.9435	0.962	0.9257	0.9477	0.9574	0.8097	0.7912	0.7597	0.8673	0.9034	2.8
NBVP [6]	0.9888	0.9609	0.9086	0.916	0.9257	0.9351	0.9416	0.7289	0.778	0.7009	0.8251	0.8736	5.8
AEP [65]	0.9362	0.9687	0.9182	0.9385	0.9132	0.9353	0.9214	0.7581	0.7647	0.7616	0.8484	0.8786	5.8
ESM _{IPP} [63]	0.9944	0.982	0.9344	0.9575	0.9338	0.9429	0.9613	0.8077	0.7028	0.768	0.8513	0.8942	3.6
ESM _{Recon} [63]	0.9862	0.9831	0.9355	0.9603	0.9355	0.9468	0.9538	0.7927	0.7965	0.7986	0.8809	0.9064	2.7
SplatPlanner [8]	0.9848	0.9891	0.9391	0.9597	0.9201	0.9527	0.9602	0.8118	0.7576	0.7575	0.8674	0.9	3.3
Classic [78]	0.018 ± 0.014	0.019 ± 0.025	0.021 ± 0.02	0.015 ± 0.015	0.021 ± 0.021	0.02 ± 0.019	0.02 ± 0.019	0.03 ± 0.034	0.03 ± 0.032	0.046 ± 0.046	0.032 ± 0.028	-	2.6
Rapid [14]	0.021 ± 0.014	0.018 ± 0.014	0.021 ± 0.019	0.016 ± 0.015	0.022 ± 0.021	0.02 ± 0.021	0.02 ± 0.018	0.029 ± 0.033	0.032 ± 0.034	0.043 ± 0.046	0.034 ± 0.03	-	2.7
NBVP [6]	0.018 ± 0.013	0.019 ± 0.016	0.023 ± 0.022	0.017 ± 0.016	0.021 ± 0.02	0.021 ± 0.019	0.022 ± 0.02	0.031 ± 0.032	0.037 ± 0.035	0.044 ± 0.045	0.034 ± 0.027	-	5.
AEP [65]	0.023 ± 0.015	0.021 ± 0.016	0.024 ± 0.022	0.017 ± 0.017	0.026 ± 0.026	0.021 ± 0.02	0.025 ± 0.022	0.032 ± 0.033	0.037 ± 0.037	0.043 ± 0.046	0.034 ± 0.028	-	6.1
ESM _{IPP} [63]	0.016 ± 0.012	0.018 ± 0.015	0.021 ± 0.019	0.016 ± 0.016	0.019 ± 0.017	0.02 ± 0.019	0.02 ± 0.018	0.029 ± 0.032	0.036 ± 0.03	0.042 ± 0.045	0.036 ± 0.03	-	2.3
ESM _{Recon} [63]	0.018 ± 0.013	0.019 ± 0.015	0.022 ± 0.019	0.016 ± 0.015	0.02 ± 0.018	0.02 ± 0.019	0.021 ± 0.019	0.029 ± 0.031	0.031 ± 0.029	0.038 ± 0.044	0.031 ± 0.027	-	2.5
SplatPlanner [8]	0.02 ± 0.014	0.018 ± 0.015	0.021 ± 0.02	0.016 ± 0.016	0.022 ± 0.022	0.018 ± 0.017	0.021 ± 0.018	0.028 ± 0.032	0.033 ± 0.03	0.043 ± 0.045	0.033 ± 0.028	-	2.7

Table 3: **Online Surface Reconstruction: Mesh Accuracy** – Comparative evaluation of the produced meshes by means of F-score and RMSE, averaged over 5 runs.




Figure 6: Summary of our benchmark results showing the *state-of-the-art*, Pareto-optimal planners (circled in black) w.r.t their joint Autonomous Exploration Coverage vs. Online Reconstruction Accuracy ranks on FLYBO.

work, planning implementations and data in order to improve existing planners with alternatives to their commonly supported piecewise-linear trajectories.

Future work includes the addition of state-of-the-art planners that are not yet published or available in our benchmark. Other exciting lines of work include the addition of data modalities such as LiDAR scans to the proposed indoor scenes within the realm of autonomous flying robots, to serve downstream applications such as online scene completion [20, 18], or by considering online shape recognition to speed up explorations by using the volumetric occupancy

	#Iterations	Per-iter. avg time (ms)	Tot. Path length (m)
Classic [78]	147 ± 7	35 ± 19	292 ± 21
Rapid [14]	52 ± 6	85 ± 110	442 ± 15
NBVP [6]	196 ± 2	858 ± 83	554 ± 8
AEP [65]	173 ± 5	37 ± 9	500 ± 15
ESM _{IPP} [63]	251 ± 1	N/A	673 ± 1
ESM _{Recon} [63]	246.9 ± 1	N/A	679 ± 2
SplatPlanner [8]	61 ± 2	236 ± 69	641 ± 31

Table 4: Comparative breakdown of exploration planning timings and statistics on the WAREHOUSE (at $t = 900s$).

of objects in a TSDF representation [36, 4].

6. Conclusion

We have presented FLYBO, the first unified benchmark environment that is specifically tailored for autonomous exploration using MAVs. It provides common material in terms of datasets, implementation of state-of-the-art methods, and a unified, exploration-centric, evaluation framework. In particular, our system allows to evaluate a planner according to all of the presented criteria and datasets in merely 17h30 of computation, which is a fraction of the time required to produce and package our framework and experiments. We hope that this significant speed-up will stimulate new contributions to the field in the near future.




Figure 7: **A few selected qualitative results** – The top four rows show results on the WAREHOUSE area, while the bottom three rows show the two-story OFFICES C at different exploration stages, running different methods.

References

- [1] <https://www.skydio.com/3d-scan>. 1
- [2] Benjamin Adler, Junhao Xiao, and Jianwei Zhang. Autonomous exploration of urban environments using unmanned aerial vehicles. *Journal of Field Robotics*, 31(6):912–939, 2014.
- [3] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. 3
- [4] Reem Ashour, Tarek Taha, Jorge Manuel Miranda Dias, Lakmal Seneviratne, and Nawaf Almoosa. Exploration for object mapping guided by environmental semantics using uavs. *Remote Sensing*, 12(5):891, 2020. 7
- [5] Matt Berger, Josh Levine, Luis Gustavo Nonato, Gabriel Taubin, and C Silva. An end-to-end framework for evaluating surface reconstruction. *Sci Comput Imag Inst*, 2011. 3
- [6] Andreas Bircher, Mina Kamel, Kostas Alexis, Helen Oleynikova, and Roland Siegwart. Receding horizon” next-best-view” planner for 3d exploration. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 1462–1468. IEEE, 2016. 1, 2, 3, 5, 7
- [7] Fredrik Bissmarck, Martin Svensson, and Gustav Tolt. Efficient algorithms for next best view evaluation. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5876–5883. IEEE, 2015.
- [8] Anthony Brunel, Amine Bourki, Cédric Demonceaux, and Olivier Strauss. Splatplanner: Efficient autonomous exploration via permutohedral frontier filtering. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2021. 1, 2, 3, 4, 5, 7
- [9] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Nießner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. In *2017 International Conference on 3D Vision (3DV)*, pages 667–676. IEEE Computer Society, 2017. 3
- [10] Devendra Singh Chaptol, Dhiraj Prakashchand Gandhi, Abhinav Gupta, and Russ R Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. *Advances in Neural Information Processing Systems*, 33, 2020.
- [11] Benjamin Charrow, Gregory Kahn, Sachin Patil, Sikang Liu, Ken Goldberg, Pieter Abbeel, Nathan Michael, and Vijay Kumar. Information-theoretic planning with trajectory optimization for dense 3d mapping. In *Robotics: Science and Systems*, volume 11, 2015.
- [12] Shengyong Chen, Youfu Li, and Ngai Ming Kwok. Active vision in robotic systems: A survey of recent developments. *The International Journal of Robotics Research*, 30(11):1343–1377, 2011.
- [13] Sanjiban Choudhury, Mohak Bhardwaj, Sankalp Arora, Ashish Kapoor, Gireeja Ranade, Sebastian Scherer, and Debadepta Dey. Data-driven planning via imitation learning. *The International Journal of Robotics Research*, 37(13-14):1632–1672, 2018. 5
- [14] Titus Cieslewski, Elia Kaufmann, and Davide 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)*, pages 2135–2142. IEEE, 2017. 1, 2, 4, 5, 7
- [15] Cl Connolly. The determination of next best views. In *Proceedings. 1985 IEEE international conference on robotics and automation*, volume 2, pages 432–435. IEEE, 1985.
- [16] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 303–312, 1996.
- [17] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 3
- [18] Angela Dai and Matthias Nießner. Scan2mesh: From unstructured range scans to 3d meshes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5574–5583, 2019. 7
- [19] Anna Dai, Sotiris Papatheodorou, Nils Funk, Dimos Tzoumanikas, and Stefan Leutenegger. Fast frontier-based information-driven autonomous exploration with an mav. In *2020 IEEE international conference on robotics and automation (ICRA)*, page 50. IEEE, 2020. 1, 2, 4
- [20] Angela Dai, Daniel Ritchie, Martin Bokeloh, Scott Reed, Jürgen Sturm, and Matthias Nießner. Scancomplete: Large-scale scene completion and semantic segmentation for 3d scans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2018. 4, 7
- [21] Di Deng, Runlin Duan, Jiahong Liu, Kuangjie Sheng, and Kenji Shimada. Robotic exploration of unknown 2d environment using a frontier-based automatic-differentiable information gain measure. In *2020 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, pages 1497–1503. IEEE, 2020. 2
- [22] Wenchao Ding, Wenliang Gao, Kaixuan Wang, and Shaojie Shen. An efficient b-spline-based kinodynamic replanning framework for quadrotors. *IEEE Transactions on Robotics*, 35(6):1287–1306, 2019. 5
- [23] Siyan Dong, Kai Xu, Qiang Zhou, Andrea Tagliasacchi, Shiqing Xin, Matthias Nießner, and Baoquan Chen. Multi-robot collaborative dense scene reconstruction. *ACM Transactions on Graphics (TOG)*, 38(4):1–16, 2019. 1, 3, 4
- [24] Christian Dornhege and Alexander Kleiner. A frontier-void-based approach for autonomous exploration in 3d. In *2011 IEEE International Symposium on Safety, Security, and Rescue Robotics*, 2011. 2
- [25] Christian Dornhege and Alexander Kleiner. A frontier-void-based approach for autonomous exploration in 3d. *Advanced Robotics*, 27(6):459–468, 2013.
- [26] Tom Duckett and Ulrich Nehmzow. Exploration of unknown environments using a compass, topological map and neural network. In *Proceedings 1999 IEEE International Symposium on Computational Intelligence in Robotics and Automation. CIRA’99*, pages 312–317. IEEE, 1999.
- [27] Matthias Faessler, Antonio Franchi, and Davide Scaramuzza. Differential flatness of quadrotor dynamics subject to rotor

- drag for accurate tracking of high-speed trajectories. *IEEE Robot. Autom. Lett.*, 3(2):620–626, Apr. 2018. 4
- [28] Pedro F Felzenszwalb and Daniel P Huttenlocher. Distance transforms of sampled functions. *Theory of computing*, 8(1):415–428, 2012.
- [29] Fadri Furrer, Michael Burri, Markus Achtelik, and Roland Siegwart. *Robot Operating System (ROS): The Complete Reference (Volume 1)*, chapter RotorS—A Modular Gazebo MAV Simulator Framework, pages 595–625. Springer International Publishing, Cham, 2016. 4
- [30] Wei Gao and Russ Tedrake. Filterreg: Robust and efficient probabilistic point-set registration using gaussian filter and twist parameterization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11095–11104, 2019.
- [31] Héctor H González-Banos and Jean-Claude Latombe. Navigation strategies for exploring indoor environments. *The International Journal of Robotics Research*, 21(10-11):829–848, 2002.
- [32] Lionel Heng, Alkis Gotovos, Andreas Krause, and Marc Pollefeys. Efficient visual exploration and coverage with a micro aerial vehicle in unknown environments. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1071–1078. IEEE, 2015.
- [33] Benjamin Hepp, Debadeepta Dey, Sudipta N Sinha, Ashish Kapoor, Neel Joshi, and Otmar Hilliges. Learn-to-score: Efficient 3d scene exploration by predicting view utility. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 437–452, 2018. 1
- [34] Dirk Holz, Nicola Basilico, Francesco Amigoni, and Sven Behnke. Evaluating the efficiency of frontier-based exploration strategies. In *ISR 2010 (41st International Symposium on Robotics) and ROBOTIK 2010 (6th German Conference on Robotics)*, pages 1–8. VDE, 2010.
- [35] Armin Hornung, Kai M Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. Octomap: An efficient probabilistic 3d mapping framework based on octrees. *Autonomous robots*, 34(3):189–206, 2013. 5
- [36] Ji Hou, Angela Dai, and Matthias Nießner. Revealnet: Seeing behind objects in rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2098–2107, 2020. 7
- [37] Stefan Isler, Reza Sabzevari, Jeffrey Delmerico, and Davide Scaramuzza. An information gain formulation for active volumetric 3d reconstruction. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3477–3484. IEEE, 2016. 3
- [38] Mark W Jones, J Andreas Baerentzen, and Milos Sramek. 3d distance fields: A survey of techniques and applications. *IEEE Transactions on visualization and Computer Graphics*, 12(4):581–599, 2006. 5
- [39] Miguel Juliá, Arturo Gil, and Oscar Reinoso. A comparison of path planning strategies for autonomous exploration and mapping of unknown environments. *Autonomous Robots*, 33(4):427–444, 2012.
- [40] Mina Kamel, Thomas Stastny, Kostas Alexis, and Roland Siegwart. Model predictive control for trajectory tracking of unmanned aerial vehicles using robot operating system. In Anis Koubaa, editor, *Robot Operating System (ROS) The Complete Reference, Volume 2*. Springer, 2017.
- [41] Sertac Karaman and Emilio Frazzoli. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7):846–894, 2011.
- [42] Lydia E Kavraki, Petr Svestka, J-C Latombe, and Mark H Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE transactions on Robotics and Automation*, 12(4):566–580, 1996. 2
- [43] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)*, 32(3):1–13, 2013. 2, 3, 4
- [44] Leonid Keselman, John Iselin Woodfill, Anders Grunnet-Jepsen, and Achintya Bhowmik. Intel realsense stereoscopic depth cameras. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1–10, 2017. 5
- [45] Matthew Klingensmith, Ivan Dryanovski, Siddhartha S Srinivasa, and Jizhong Xiao. Chisel: Real time large scale 3d reconstruction onboard a mobile device using spatially hashed signed distance fields. In *Robotics: science and systems*, volume 4, page 1. Citeseer, 2015.
- [46] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017. 4
- [47] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. Abc: A big cad model dataset for geometric deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9601–9611, 2019. 2
- [48] Sebastian Koch, Albert Matveev, Zhongshi Jiang, Francis Williams, Alexey Artemov, Evgeny Burnaev, Marc Alexa, Denis Zorin, and Daniele Panozzo. Abc: A big cad model dataset for geometric deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9601–9611, 2019.
- [49] James J Kuffner and Steven M LaValle. Rrt-connect: An efficient approach to single-query path planning. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, volume 2, pages 995–1001. IEEE, 2000. 2
- [50] Steven M LaValle. Rapidly-exploring random trees: A new tool for path planning. 1998.
- [51] Charles Loop, Qin Cai, Sergio Orts-Escolano, and Philip A Chou. A closed-form bayesian fusion equation using occupancy probabilities. In *2016 Fourth International Conference on 3D Vision (3DV)*, pages 380–388. IEEE, 2016.
- [52] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4):163–169, 1987.
- [53] Daniel Mellinger and Vijay Kumar. Minimum snap trajectory generation and control for quadrotors. In *2011 IEEE international conference on robotics and automation*, pages 2520–2525. IEEE, 2011.

- [54] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE international symposium on mixed and augmented reality*, pages 127–136. IEEE, 2011.
- [55] Helen Oleynikova, Michael Burri, Zachary Taylor, Juan Nieto, Roland Siegwart, and Enric Galceran. Continuous-time trajectory optimization for online uav replanning. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5332–5339. IEEE, 2016. [5](#)
- [56] Helen Oleynikova, Zachary Taylor, Marius Fehr, Roland Siegwart, and Juan Nieto. Voxblox: Incremental 3d euclidean signed distance fields for on-board mav planning. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (iros)*, pages 1366–1373. IEEE, 2017.
- [57] Emanuele Palazzolo and Cyrill Stachniss. Information-driven autonomous exploration for a vision-based mav. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 4:59, 2017.
- [58] Emanuele Palazzolo and Cyrill Stachniss. Effective exploration for mavs based on the expected information gain. *Drones*, 2(1):9, 2018.
- [59] Christos Papachristos, Shehryar Khattak, and Kostas Alexis. Uncertainty-aware receding horizon exploration and mapping using aerial robots. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 4568–4575. IEEE, 2017.
- [60] Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Real-time progressive 3d semantic segmentation for indoor scenes. In *2019 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1089–1098. IEEE, 2019.
- [61] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, page 5. Kobe, Japan, 2009. [4](#)
- [62] Charles Richter, Adam Bry, and Nicholas Roy. Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments. In *Robotics Research*, pages 649–666. Springer, 2016. [4](#)
- [63] Lukas Schmid, Michael Pantic, Raghav Khanna, Lionel Ott, Roland Siegwart, and Juan Nieto. An efficient sampling-based method for online informative path planning in unknown environments. *IEEE Robotics and Automation Letters*, 5(2):1500–1507, 2020. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#)
- [64] Alan C Schultz, William Adams, and Brian Yamauchi. Integrating exploration, localization, navigation and planning with a common representation. *Autonomous Robots*, 6(3):293–308, 1999.
- [65] Magnus Selin, Mattias Tiger, Daniel Duberg, Fredrik Heintz, and Patric Jensfelt. Efficient autonomous exploration planning of large-scale 3-d environments. *IEEE Robotics and Automation Letters*, 4(2):1699–1706, 2019. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#)
- [66] Shital Shah, Debadeepta Dey, Chris Lovett, and Ashish Kapoor. Airsim: High-fidelity visual and physical simulation for autonomous vehicles. In *Field and service robotics*, pages 621–635. Springer, 2018.
- [67] Shaojie Shen, Nathan Michael, and Vijay Kumar. Stochastic differential equation-based exploration algorithm for autonomous indoor 3d exploration with a micro-aerial vehicle. *The International Journal of Robotics Research*, 31(12):1431–1444, 2012. [2](#)
- [68] Soohwan Song and Sungho Jo. Surface-based exploration for autonomous 3d modeling. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–8. IEEE, 2018. [3](#)
- [69] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1746–1754, 2017. [3](#)
- [70] Yunlong Song, Selim Naji, Elia Kaufmann, Antonio Loquercio, and Davide Scaramuzza. Flightmare: A flexible quadrotor simulator. *arXiv preprint arXiv:2009.00563*, 2020.
- [71] Hang Su, Varun Jampani, Deqing Sun, Subhansu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. Splatnet: Sparse lattice networks for point cloud processing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2530–2539, 2018.
- [72] Ioan A Sucan, Mark Moll, and Lydia E Kavraki. The open motion planning library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, 2012.
- [73] Vladyslav Usenko, Lukas von Stumberg, Andrej Pangeric, and Daniel Cremers. Real-time trajectory replanning for mavs using uniform b-splines and a 3d circular buffer. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 215–222. IEEE, 2017.
- [74] Emanuele Vespa, Nikolay Nikolov, Marius Grimm, Luigi Nardi, Paul HJ Kelly, and Stefan Leutenegger. Efficient octree-based volumetric slam supporting signed-distance and occupancy mapping. *IEEE Robotics and Automation Letters*, 3(2):1144–1151, 2018.
- [75] Christian Witting, Marius Fehr, Rik Bähneemann, Helen Oleynikova, and Roland Siegwart. History-aware autonomous exploration in confined environments using mavs. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–9. IEEE, 2018.
- [76] Fei Xia, Amir R. Zamir, Zhi-Yang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: real-world perception for embodied agents. In *Computer Vision and Pattern Recognition (CVPR), 2018 IEEE Conference on*. IEEE, 2018. [3](#)
- [77] Kai Xu, Lintao Zheng, Zihao Yan, Guohang Yan, Eugene Zhang, Matthias Niessner, Oliver Deussen, Daniel Cohen-Or, and Hui Huang. Autonomous reconstruction of unknown indoor scenes guided by time-varying tensor fields. *ACM Transactions on Graphics (TOG)*, 36(6):1–15, 2017. [1](#), [3](#)
- [78] Brian 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*, pages 146–151. IEEE, 1997. [1](#), [2](#), [5](#), [7](#)

- [79] Brian Yamauchi. Frontier-based exploration using multiple robots. In *Proceedings of the second international conference on Autonomous agents*, pages 47–53, 1998. [1](#), [2](#)
- [80] Boyu Zhou, Fei Gao, Luqi Wang, Chuhao Liu, and Shaojie Shen. Robust and efficient quadrotor trajectory generation for fast autonomous flight. *IEEE Robotics and Automation Letters*, 4(4):3529–3536, 2019. [5](#)
- [81] Boyu Zhou, Jie Pan, Fei Gao, and Shaojie Shen. Raptor: Robust and perception-aware trajectory replanning for quadrotor fast flight. *IEEE Transactions on Robotics*, 2021. [5](#)
- [82] Boyu Zhou, Yichen Zhang, Xinyi Chen, and Shaojie Shen. Fuel: Fast uav exploration using incremental frontier structure and hierarchical planning. *IEEE Robotics and Automation Letters*, 6(2):779–786, 2021. [2](#), [3](#), [5](#)
- [83] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3D: A modern library for 3D data processing. *arXiv:1801.09847*, 2018.