

SEER: Safe Efficient Exploration for Aerial Robots using Learning to Predict Information Gain

Yuezhan Tao, Yuwei Wu, Beiming Li, Fernando Cladera, Alex Zhou, Dinesh Thakur and Vijay Kumar

Abstract—We address the problem of efficient 3-D exploration in indoor environments for micro aerial vehicles with limited sensing capabilities and payload/power constraints. We develop an indoor exploration framework that uses learning to predict the occupancy of unseen areas, extracts semantic features, samples viewpoints to predict information gains for different exploration goals, and plans informative trajectories to enable safe and smart exploration. Extensive experimentation in simulated and real-world environments shows the proposed approach outperforms the state-of-the-art exploration framework by 24% in terms of the total path length in a structured indoor environment and with a higher success rate during exploration.

I. INTRODUCTION

Autonomous exploration requires a robot to navigate through an unknown region and map it in real time. This problem has always attracted great attention in the robotics community due to its direct applications in tasks such as search and rescue, precision agriculture, and autonomous inspection. All of these tasks require the robot to navigate through cluttered and complex environments. Therefore, Micro Aerial Vehicles (MAV), especially quadrotors, have gained great popularity due to their agility.

Autonomous exploration has been widely studied and recent works have measured its performance in terms of map uncertainty [1]–[3], total travel distance [4], energy efficiency [5], and exploration time [4, 6]. Beyond optimizing for different objectives, different exploration strategies and planning methods have been proposed to tackle different environments, such as structured indoor environments [7]–[11] and subterranean environments [6, 12].

Nevertheless, most of the existing works follow heuristics to greedily expand the map or maximize the information to mimic human behavior; they do not consider inferring the importance of the environment through semantic objects. Although some methods do consider semantic information to guide the exploration, this information is not fully utilized during local planning. Furthermore, in many works information gain is optimistically estimated without considering the environment. Most of these works assume that all space inside the robot FOV at the goal point will be observed, leading to an overestimate of the information gain.

All authors are with GRASP Laboratory, University of Pennsylvania {yztao, yuweiwu, beimingli, fclad, alexzhou, tdinesh, kumar}@seas.upenn.edu.

We gratefully acknowledge the support of ARL DCIST CRA W911NF-17-2-0181, NSF Grants CCR-2112665, ONR grant N00014-20-1-2822, ONR grant N00014-20-S-B001, NVIDIA, Israel Minister of Defense, and C-BRIC, a Semiconductor Research Corporation Joint University Microelectronics Program program cosponsored by DARPA. We also thank Yifei (Simon) Shao for providing helpful insights and being our safety pilot in the real-world experiments.

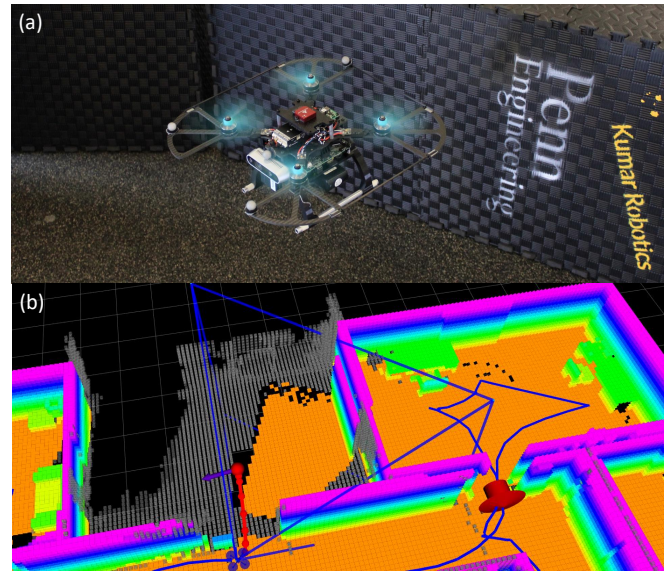


Fig. 1: (a) Falcon 250 Platform. The Falcon 250 platform is equipped with an Intel Realsense D435i camera, an Intel NUC 10 onboard computer with an i7-10710U CPU, and a Pixhawk 4 Mini flight controller. (b) A robot exploring a simulated environment with the proposed method. The predicted map is represented by gray cubes, executed trajectories are shown as blue curves, detected semantic objects are shown as red arrows, the next viewpoint is the purple arrow, and the planned trajectory is shown in red.

Motivated by these limitations, we propose an indoor MAV 3-D exploration framework that uses high-level behavioral policies and predicts information gain for intelligent and efficient exploration. Our work takes inspiration from the behavior of humans during exploration, i.e., inferring the importance of the environment by identifying and using semantic information. In addition, humans predict the structure of the unknown area by making assumptions based on observed regions. Therefore, in our proposed framework, we propose a detection and prediction module that detects frontier clusters, extracts semantic information, and predicts occupancy and information incrementally during exploration. This abstracted information is then used by a perception-aware planning module in which the behavior planning sets the global navigation goal following the behavior policies and the predicted informative planning generates local trajectories that allow safe and informative maneuvers. An illustration of our proposed system is shown in Fig. 1.

We benchmarked our proposed framework with classical and state-of-the-art methods in simulation experiments. We also conducted real-world experiments in which our proposed framework runs fully on board a Size Weight and Power (SWaP) constrained platform in a structured indoor environment. In summary, **our contributions are:**

- An incremental detection and prediction module that detects semantic objects, classifies frontiers, predicts information gain, and samples viewpoints around clustered frontiers.
- A perception-aware planning approach that guides the exploration through behavior planning and plans predicted informative trajectories to ensure intelligent, safe, and fast exploration in a structured indoor environment.
- We validate the proposed framework and algorithms extensively in simulation and real-world experiments. We will open-source our proposed system and algorithm at <https://github.com/tyuezhhan/SEER>.

II. RELATED WORK

A. Autonomous Exploration Strategy

Various exploration strategies have been proposed to autonomously map unknown regions. There are three major exploration strategies: frontier-based, information-based, and the hierarchical mixed approach.

Following the idea of purely expanding the map into unknown regions, frontier-based exploration [13] has been widely adopted. Real-world demonstrations in [8, 14] show its applicability in both 2-D and 3-D environments. To achieve high energy efficiency during exploration, [5] proposes to select the frontier within the agent’s FOV to maintain a high flight speed. Instead of searching frontiers, maximizing the information gain serves as another heuristic to guide exploration. Next Best View (NBV) is sampled from Rapidly Exploring Random Tree (RRT) [15, 16] or motion primitives [6] to maximize information gain. Different methods are also proposed to evaluate the information gain [1, 2, 17]. However, classical methods assume all voxels in the ray-casting process will be updated, resulting in an overestimate of total information gain. Hierarchical approaches have inherent advantages by benefiting from multiple strategies. In [3, 18], global paths are planned towards frontiers while local paths are sampled through motion primitives and optimized for information gain. [4, 19] plans a global coverage path through Travelling Salesman Problem (TSP) to efficiently travel through all viewpoints and sample viewpoints around frontiers to maximize coverage or information gain. To further improve the efficiency and consistency of exploration, semantic and topological information is used in [7, 11] to assist high-level robot behaviors. However, directly incorporating the semantic information into the utility function sometimes fails to guide the exploration, as the semantic information might be labelled incorrectly. In [10], NBV is selected following a decision model to guide the exploration behavior. However, it assumes that a semantic mapping model will provide perfect semantic information of the environment.

In this paper, we present a perception-aware planning framework in which the behavior planning module takes the detected semantic information into consideration to select a candidate goal, and the predicted informative planning module utilizes the predicted information gain to improve exploration efficiency.

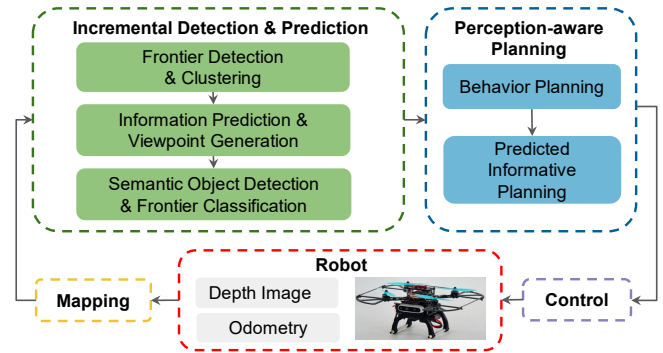


Fig. 2: Proposed system architecture. The proposed system contains two main components: the incremental detection & prediction module, and the perception-aware planning module.

B. Map Completion and Prediction

Predicting the occupancy of unknown environments can provide prior information to help with navigation and exploration tasks. Database-based approaches and deep learning-based approaches are commonly used. To infer unknown information, historical data is stored and used in [20, 21]. Without using a database, [22] extrapolates the local 2D structure to estimate the potential information gain. With the capability of learning from and adapting to a large amount of historical data, map prediction using deep learning has gained popularity in recent years. 2D occupancy maps are predicted in [9, 23] to estimate the total information or uncertainty for exploration. In [12], topological information is predicted in a large-scale underground environment. Egocentric 2D occupancy and semantic maps are predicted in [24, 25] to help the navigation tasks. To achieve safe and fast navigation, [26] proposed egocentric 3-D occupancy prediction of occluded areas. Most existing work focuses on using 2D structural information to predict the occupancy and estimate the total potential information gain within a bounding box. This approach provides a good estimate of the total information that can be obtained in a large region. However, it overestimates the information, as it cannot infer unseen environment. It also does not help with estimating the information gain at certain viewpoints in a 3-D environment.

In this paper, utilizing 3-D occupancy prediction, we present a method to predict information gain in unknown regions beyond frontiers to sample the viewpoint which maximizes the potential information gain accurately.

III. SYSTEM OVERVIEW

As illustrated in Fig. 2, the proposed system contains two main components. The incremental detection and prediction module (Sect. IV) first takes in a bounding box of the updated map, reevaluates existing frontiers, detects new frontiers, and clusters them. Then, semantic object detection, frontier classification, and occupancy prediction are executed around clusters. Viewpoints are sampled and sorted according to the predicted information gain. The perception-aware planning module (Sect. V) then uses this information to generate a two-stage exploration plan that enables safe and informative flight.

IV. INCREMENTAL DETECTION AND PREDICTION

During the exploration phase, the map of the environment is built incrementally, so it is natural to evaluate existing frontiers and detect new frontiers from the updated regions of the map. Similarly, we incrementally detect semantic objects, label frontier clusters, and sample informative viewpoints through accurate information prediction.

A. Frontier Detection and Clustering

We implemented the frontier detection and clustering algorithm of [4]. The incremental frontier detection and clustering module receives a bounding box that represents the updated region of the map. Existing frontiers inside the bounding box are reevaluated, and frontiers that have been changed are removed. New frontiers are searched and clustered inside the bounding box by a region-growing algorithm, and large clusters are split into smaller ones recursively by Principal Component Analysis (PCA) along the principal axis if the largest eigenvalue exceeds a predefined threshold. Each time a new cluster F_i is detected, a frontier cluster node is created and saved.

B. Information Prediction & Viewpoint Generation

As frontiers are the boundary between the known and unknown regions, it is natural to navigate towards them to expand the map into the unknown regions. To find a goal that provides the most information from the unknown environment, we present a novel, combined information prediction and viewpoint generation method that takes the occupancy prediction into account to predict future information gain accurately.

The occupancy grid map is represented as a vector \mathbf{M} of size $N = N_x \cdot N_y \cdot N_z$ where N_x, N_y, N_z are the sizes of the grids along each axis. Each element $m_{xyz} \in (\text{clamp}_{min}, \text{clamp}_{max})$ represents the log-odds of the probability of the occupancy for the voxel. For occupancy prediction, we cut-off the occupancy values m_{xyz} into trinary $o_{xyz} \in \{-1, 0, 1\}$ to generate \mathbf{O} , where $\{-1, 0, 1\}$ denotes $\{\text{unknown}, \text{free}, \text{occupied}\}$, respectively.

1) *3-D Occupancy Prediction*: We generate realistic training data pairs to simulate the partially observed scenes during exploration for the self-supervised 3-D occupancy prediction network. After clustering the frontiers, we extract the local region centered at the centroid of each frontier cluster F_i from the volumetric map and perform occupancy prediction for unknown regions beyond the frontiers. Taking into account the gap between simulation environments and real-world scenarios, we choose the MatterPort3D dataset [27], which contains depth images and 6-DoF camera poses collected from 90 real-world buildings. To generate the target occupancy data \mathbf{O}^{target} , we first reconstruct the 3-D occupancy grid map of the buildings from the raw depth images with 10cm resolution. Then we crop a $8m \times 8m \times 2.4m$ region at the center of each scene. Since some of the buildings contain multiple layers, we align the ground to the bottom of the cropped region to ensure consistency of the data. To generate the corresponding input data \mathbf{O}^{in} , we

first subsample 50% of the input depth images at each scene during map reconstruction. To make \mathbf{O}^{in} more similar to the exploration scenarios, we further crop it with straight lines with different slopes through the center of the box. With this operation, we also increased the number of training data pairs. As a result, we generate 93,768 $(\mathbf{O}^{in}, \mathbf{O}^{tar})$ pairs from 78 buildings in MatterPort3D.

To predict $\mathbf{O}_{known}^{tar} \setminus \mathbf{O}_{unknown}^{in}$, the known occupancy value in the target that is unknown in the input, we use OPNet [26] which contains a U-Net [28] with Atrous Spatial Pyramid Pooling (ASPP) [29] to improve the prediction on a multiscale level. As the goal of the occupancy prediction is to provide occupancy information beyond the frontiers for the estimation of information gains, we design a loss function \mathcal{L} that contains an occupancy loss \mathcal{L}_{occ} that emphasizes the prediction of the occupancy of individual voxels and a structural loss \mathcal{L}_{struct} that focuses on the inference of unknown structures. The loss is defined as follows:

$$\mathcal{L} = \omega_{occ} * \mathcal{L}_{occ} + \omega_{struct} * \mathcal{L}_{struct} \quad (1)$$

where ω_{occ} and ω_{struct} are hyperparameters. For occupancy loss \mathcal{L}_{occ} , we compute weighted binary cross-entropy (BCE) between the prediction \mathbf{O}^{pred} and the target \mathbf{O}^{tar} :

$$\begin{aligned} \mathcal{L}_{occ} = & -\frac{1}{N} \sum_{x=1}^{N_x} \sum_{y=1}^{N_y} \sum_{z=1}^{N_z} \lambda_{xyz} [o_{xyz}^{tar} \cdot \log o_{xyz}^{pred} \\ & + (1 - o_{xyz}^{tar}) \cdot \log(1 - o_{xyz}^{pred})] \\ \lambda_{xyz} = & \begin{cases} 0 & \text{if } o_{xyz}^{tar} = -1 \\ \alpha & \text{if } o_{xyz}^{in} = -1 \text{ and } o_{xyz}^{tar} = 1 \\ 1 & \text{otherwise} \end{cases} \end{aligned} \quad (2)$$

where α is a weight to emphasize the prediction of occupied voxels on the target occupancy map. The structural loss considers the weighted $L1$ norm between the number of occupied voxels in the target along the Z axis and the number of occupied voxels in the prediction along the Z axis to emphasize the prediction of unseen structures. The structural loss is defined as:

$$\begin{aligned} \mathcal{L}_{struct} = & \frac{1}{N_x \cdot N_y} \sum_{x=1}^{N_x} \sum_{y=1}^{N_y} \varphi_{xy} \left| \sum_{z=1}^{N_z} \mathbb{1}_{\{o_{xyz}^{pred}=1\}} \right. \\ & \left. - \sum_{z=1}^{N_z} \mathbb{1}_{\{o_{xyz}^{tar}=1\}} \right| \\ \varphi_{xy} = & \begin{cases} \beta & \text{if } \sum_{z=1}^{N_z} \mathbb{1}_{\{o_{xyz}^{tar}=1\}} > N_z/2 \\ 1 & \text{otherwise} \end{cases} \end{aligned} \quad (3)$$

where we set weight β on the structural information along the z axis. As the frontiers are clustered, the occupancy prediction is performed on each F_i incrementally.

2) *Information Prediction and Viewpoint Generation*: The 3-D viewpoint generation and information prediction steps run after the occupancy prediction stage. Inspired by [4], we design a two-stage approach to sample viewpoints $VP_i = \{(\mathbf{p}_{i,1}, \xi_{i,1}, g_{i,1}), (\mathbf{p}_{i,2}, \xi_{i,2}, g_{i,2}), \dots, (\mathbf{p}_{i,n}, \xi_{i,n}, g_{i,n})\}$ at each F_i , where $(\mathbf{p}_{i,j}, \xi_{i,j}, g_{i,j})$ represents the 3-D position,

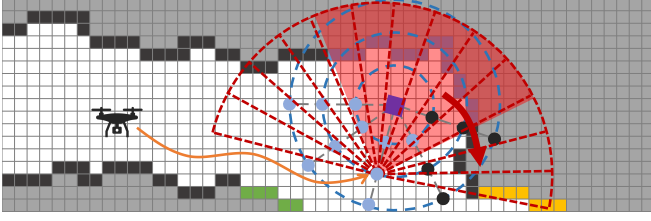


Fig. 3: Two-stage viewpoint generation. Feasible Viewpoints (blue dots) are sampled in cylindrical coordinate system from each cluster centroid (purple box). At each viewpoint, predicted information gain is computed in parallel at all sampled yaw angles following a sliding window.

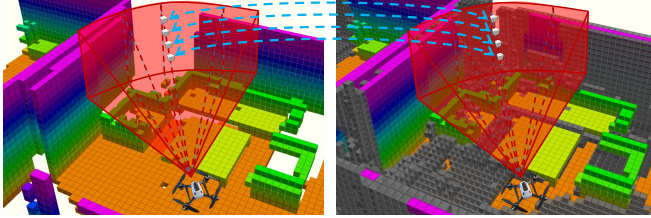


Fig. 4: Information prediction process. During ray-casting, ‘unknown’ voxels in the observed map (left panel) are looked up from the map with occupancy prediction (right panel) to predict information gain.

yaw angle, and predicted information gain at each sampled viewpoint, respectively. As shown in Fig. 3, in the first stage, the position $\mathbf{p}_{i,j}$ of each viewpoint is uniformly sampled from the cylindrical coordinates that originate from the center of the cluster. At the second stage, centered along the line connects $\mathbf{p}_{i,j}$ to the center of the cluster, multiple yaw angles $\{\xi_{i,j,1}, \xi_{i,j,2}, \dots, \xi_{i,j,m}\}$ are uniformly sampled. Then, we use the methods from [30] to avoid repeated ray-castings and compute in parallel the information gain prediction for all yaw angles sampled. The information prediction process is illustrated in Fig. 4. To reduce redundant computations, we calculate the information gain along each slice and aggregate the gains inside each FOV to get the predicted information at different sampled yaw angles. We designed a new information prediction approach by checking voxels from both the current map and the predicted map to better predict the information during ray-casting, as detailed in Algorithm 1. After predicting the information gain $g_{i,j,k}$ at each yaw angle $\xi_{i,j,k}$, we save the one with the highest information gain as $\xi_{i,j}$ together with $\mathbf{p}_{i,j}$ and $g_{i,j}$. At the end of the viewpoint generation process, we sort all viewpoints in descending order of predicted information and save the top N_{vp} viewpoints to maximize utility during the planning process.

C. Semantic Object Detection and Frontier Classification

Indoor environments normally contain rich semantic information which can be used to guide the exploration. In this paper, we study this by detecting doors as semantic objects and classifying frontiers to help the exploration behavioral planning. Unlike [7, 11], which detect doors from laser scans, we directly detect doors from occupancy maps around frontier clusters. This allows the detection to happen in parallel at all frontier clusters.

We design a hand-crafted two-stage detection algorithm to detect doors on the occupancy map efficiently. At each

Algorithm 1 Information Gain Prediction along a ray

```

1: ray_start ← ComputeRayStart();
2: ray_end ← ComputeRayEnd();
3: pred_gain ← 0;
4: while ! reach ray end do
5:   vox ← next voxel on the ray;
6:   if vox ! in map or vox == OCCUPIED then
7:     break;
8:   if vox == UNKNOWN then
9:     in_pred_map ← checkCellInPredMap();
10:    if ! in_pred_map then
11:      pred_gain ++;
12:      continue;
13:    pred_vox ← getVoxInPredMap();
14:    if pred_vox == UNKNOWN then
15:      pred_gain ++;
16:    else if pred_vox == OCCUPIED then
17:      pred_gain ++;
18:      break;
19:    else
20:      pred_gain ++;

```

frontier cluster F_i , we extract a small 2D occupancy grid, detect gaps between parallel lines from it using a canny edge detector and probabilistic hough line transform, and compute the position and direction of the gap. In the second stage, we extract the plane perpendicular to the direction of the gap centered at the gap position, which contains the contour of the door. The distance between vertical walls is measured to verify it meets a typical door size. We remove false positive detection by continuously rechecking the detected doors after map updates. All detected doors are initially stored as ‘to-be-confirmed’ until the robot navigates to Semantic Object of Interest (SoI) as in Fig. 5 to re-evaluate whether it’s a door.

After semantic object detection, frontier clusters are classified according to the robot’s exploration state and detected semantic objects. In this work, we distinguish rooms and corridors and assign the corresponding label to each F_i . F_i is labeled as ‘corridor’ when the robot is exploring a corridor, navigating towards a door, or exiting a room. While the robot is entering or exploring a room, F_i is labeled as ‘room’. Frontier clusters classified as ‘room’ will be checked to confirm that they are inside the same room as the robot before being considered as a candidate goal during behavioral planning.

V. PERCEPTION-AWARE PLANNING FOR EXPLORATION

A. Behavior Planning

In the structured indoor environment, semantic objects provide strong information about the importance of the region, which classical frontier-based or information-based exploration cannot fully utilize. Assuming that the area of interest should be visited with higher priority and explored thoroughly, we propose a high-level Behavior Guidance State Machine (BGSM), which selects the appropriate next goal to explore, as illustrated in Fig. 5. The navigation behavior is determined by the existence of SoIs, robot state, and utility of frontiers. For the frontier cluster F_i , we define the heuristic utility function of viewpoint $vp_{i,j} = (\mathbf{p}_{i,j}, \xi_{i,j}, g_{i,j})$ as follows:

$$\mathcal{U}(F_i) = \max_{VP_i} u(vp_{i,j}) = \max_{VP_i} \frac{v_{p,m}}{\text{len}[\mathbf{p}_e, \mathbf{p}_{i,j}]} g_{i,j} \quad (4)$$

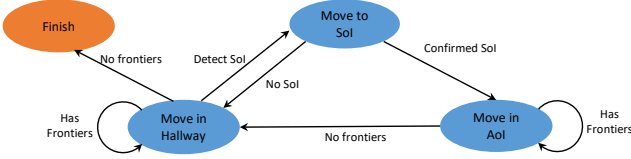


Fig. 5: BGSM. The BGSM takes the abstracted information from the detection and prediction module to make elaborate decisions on next navigation goal during exploration.

where $v_{p,m}$ is the maximum velocity, \mathbf{p}_e is the robot's position and $\text{len}[\cdot]$ is the length of a geometric path we search from the current position to the viewpoint. If a robot detects a new SoI in the corridor, it will navigate to it, confirm it, and enter the Area of Interest (AoI). Otherwise, it will continuously visit other frontiers according to their computed utilities. Once the robot enter the AoI, it will explore the area until no frontiers exist in AoI.

B. Predicted Informative Planning

1) *Polynomial Trajectory Generations*: Motion planning in 3-D for exploration requires generating flexible, dynamically feasible trajectories. Leveraging the differential flatness of quadrotor dynamics, we optimize the position and yaw trajectories of our agent separately. In particular, we parametrize the trajectory as the output of a 3-D m -order integrator, composed of M polynomial segments $\Phi(t)$. Therefore, our trajectory optimization problem becomes

$$\min_{\Phi(t)} J_c = \int_{t_0}^{t_M} \|\Phi^{(\kappa)}(t)\|^2 dt \quad (5a)$$

$$\text{s.t. } \mathcal{H}_c(\Phi(t), \dots, \Phi^{(\kappa-1)}(t)) = \mathbf{0}, \quad (5b)$$

$$\mathcal{G}_d(\Phi(t), \dots, \Phi^{(\kappa-1)}(t)) \preceq \mathbf{0}, \quad (5c)$$

$$\forall t \in [t_0, t_M]. \quad (5d)$$

Here \mathcal{H}_c encodes the smoothness, and \mathcal{G}_d represents the dynamic feasibility and obstacle avoidance constraints. We employ the method in [31] to represent the collision-free space as a safe flight corridor. We also leverage the framework in [32] to relax the constraints and solve an unconstrained spatial-temporal optimization problem, which is further extended for yaw trajectory optimization.

2) *Predicted Informative Yaw Planning*: To generate perception-aware trajectories, we optimize yaw angles with respect to both the predicted information gain and the relative view towards SoI. Specifically, we first search yaw primitives along the position trajectory with time intervals $\mathbf{T} = [T_1, \dots, T_M]^T$, $T_i = t_i - t_{i-1}$ to get a reference yaw $\psi(t)$ with high-value information. The total cost function to evaluate a primitive segment is defined by:

$$f_{\psi,i} = \beta_{\psi} \int_{t_{i-1}}^{t_i} \ddot{\psi}(t)^2 dt - \frac{T_i}{N_i} \sum_{\lambda=1}^{N_i} (\beta_s f_s(\tilde{t}) + \beta_u g_e(\tilde{t})), \quad (6)$$

where N_i is the sampling number, $\tilde{t} = t_{i-1} + \lambda T_i / N_i$, and $\beta_{\psi}, \beta_s, \beta_u$ are non-negative penalties associated with yaw control efforts, perception of semantics, and predicted information gain. The information gain g_e is evaluated using

Algorithm 1. The component of the cost encouraging the SoI to lie within the horizontal FOV θ is given by:

$$f_s(t) = \sigma \left(\frac{\theta}{2} - |\psi(t) - \text{atan2}(e_2^T \mathbf{p}_r(t), e_1^T \mathbf{p}_r(t))| \right), \quad (7)$$

where $\sigma(\cdot) = \max(\cdot, 0)$, $e_1 = [1, 0, 0]$, $e_2 = [0, 1, 0]$, and $\mathbf{p}_r(t) = \Phi(t) - \mathbf{p}_s$ is the relative position to the center point \mathbf{p}_s of the semantic object that the robot navigates to. After identifying an optimal reference yaw sequence $\gamma = [\gamma_1, \dots, \gamma_{M-1}]$ using a search-based procedure, we directly solve for the optimal yaw trajectory $\Psi(t)$ that interpolates the yaw angles $\psi = [\psi_1, \dots, \psi_{M-1}]$ at specified times. The cost function can be expressed as:

$$J_{\Sigma}(\psi) = \sum_{i=1}^M \int_{t_{i-1}}^{t_i} \Psi^{(3)}(t)^2 dt + \rho_p \sum_{i=1}^{M-1} |\psi_i - \gamma_i|^2 + \quad (8)$$

$$\rho_v \sum_{i=1}^{M-1} \sigma \left(\left| \frac{\psi_{i+1} - \psi_{i-1}}{T_{i+1} + T_i} \right| - v_{\psi,m} \right)^3 +$$

$$\rho_a \sum_{i=1}^{M-1} \sigma \left(\left| \frac{(\psi_{i+1} - \psi_i)/T_{i+1} - (\psi_i - \psi_{i-1})/T_i}{(T_{i+1} + T_i)/2} \right| - a_{\psi,m} \right)^3,$$

where $v_{\psi,m}, a_{\psi,m}$ are the maximum yaw velocity and acceleration. ρ_p, ρ_v, ρ_a are constant weights to prevent from undesired and infeasible yaw.

VI. RESULTS AND ANALYSIS

A. Implementation details

To achieve an efficient map update, we used a lightweight volumetric mapping module adapted from [33]. We implemented OPNet in PyTorch and trained it on the generated dataset. We used an Adam optimizer with an initial learning rate of 10^{-3} , decays by 0.1 every 8 epochs, and a batch size of 10. We set $\omega_{occ} = 2$, $\omega_{struct} = 1$, $\alpha = 5$ and $\beta = 2$. It takes 2.87 hours to train 20 epochs on a machine with an NVIDIA RTX A5000 GPU.

B. Simulation Experiments

1) *Information Gain Prediction*: To validate the effectiveness of our proposed information gain prediction module, we first compute the ground truth information gain $g_{i,j}^{gt}$ by comparing the occupancy between the observed map and the ground truth map during exploration at every VP_i sampled from every F_i . Then, we compute the information gain with both the proposed method $g_{i,j}$ and the classic method $g_{i,j}^{cls}$ at the same viewpoint $vp_{i,j}$. We compare the mean percentage error between $g_{i,j}^{gt}$ and the information gain estimated by different methods from the **22,276** data groups. The mean percentage error between $g_{i,j}^{gt}$ and $g_{i,j}^{cls}$ is **206.637%** while between $g_{i,j}^{gt}$ and $g_{i,j}$ **124.641%**. The results show that our method outperforms the classical method by $\approx 40\%$.

2) *Benchmark Results*: We have extensively tested our proposed system in Gazebo simulations. Controlled experiments demonstrated how information prediction and perception-aware planning benefit exploration. We benchmarked our proposed method with the state-of-the-art exploration algorithm in both small ($175m^3$) and large ($392m^3$)

TABLE I: EXPLORATION STATISTICS FOR TWO TESTING ENVIRONMENTS

| Method | Small Office | | | | | | | | Succ(%) | Big Office | | | | | | | | |
|------------------------------|----------------------|--------------|--------------|-------------|-----------------|--------------|--------------|-------------|------------|----------------------|--------------|--------------|-------------|-----------------|--------------|--------------|-------------|------------|
| | Exploration Time (s) | | | | Path Length (m) | | | | | Exploration Time (s) | | | | Path Length (m) | | | | Succ(%) |
| | Min | Max | Avg | Std | Min | Max | Avg | Std | | Min | Max | Avg | Std | Min | Max | Avg | Std | |
| Frontier [13] | 69.79 | 102.2 | 79.88 | 10.61 | 36.52 | 51.69 | 42.90 | 5.47 | 30% | 165.4 | 209.1 | 181.1 | 17.13 | 96.31 | 118.9 | 106.4 | 8.09 | 20% |
| FUEL [4] | 54.68 | 97.65 | 66.18 | 10.88 | 40.21 | 68.28 | 46.76 | 6.51 | 85% | 150.5 | 165.4 | 160.8 | 4.46 | 111.2 | 134.1 | 119.7 | 6.29 | 40% |
| Frontier + Util ^a | 51.76 | 95.95 | 72.47 | 12.86 | 34.57 | 65.35 | 47.51 | 8.26 | 50% | 134.0 | 196.5 | 156.7 | 21.07 | 93.02 | 145.2 | 108.1 | 18.99 | 25% |
| Frontier + Pred ^b | 53.06 | 78.91 | 62.92 | 7.96 | 37.34 | 52.72 | 42.70 | 4.42 | 65% | 128.1 | 178.7 | 148.8 | 16.03 | 86.31 | 129.6 | 106.4 | 14.60 | 35% |
| SEER | 58.90 | 96.52 | 75.33 | 11.17 | 25.96 | 46.07 | 38.61 | 6.06 | 80% | 129.8 | 202.1 | 163.2 | 19.21 | 76.6 | 113.4 | 90.99 | 11.73 | 75% |

^a Frontier with highest utility, trajectory generation method same as FUEL.

^b Frontier with highest utility computed with information prediction, trajectory generation method same as FUEL.

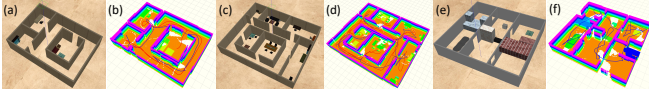


Fig. 6: Simulated 3-D Environment in Gazebo and 3-D maps after exploration. Small office (a) of volume $175m^3$. Big Office (c) of volume $392m^3$. Furniture models from [34]. Executed trajectories from the proposed method (blue) and FUEL (red) are visualized in the final maps (b, d). 3-D exploration results (f) of the proposed method is demonstrated in the indoor cluttered 3-D environment (e)

scale indoor structured environments as shown in Fig. 6. The box-bounded maximum velocity is set as $1 m/s$ and the maximum acceleration is $1 m/s^2$. The result of running 20 experiments for each method is shown in Tab. I. We record the statistics after a successful exploration, which is defined as having no collisions during the exploration and when the explored volume reaches 95% of the total space.

We observed that, by evaluating utilities at the frontiers, the selected candidate goal will lead to more immediate information in general. However, since the information is estimated blindly, it will lead to wrong decisions by selecting a goal that actually provides lower utilities. By utilizing information prediction, information gains at frontiers are estimated much more accurately, which directly improves the quality of selected goals, leading to faster exploration and shorter path lengths. The high-level behavior planning further improves this by reducing repeated visits to the explored region. Instead of maximizing immediate information gain, FUEL minimizes the length of the global tour that covers all frontiers. However, we observed that this coverage tour optimization sometimes leads to repeated visits to the explored regions as shown in Fig. 6. This results in a long path length in the structured indoor setting, which can be seen in Tab. I. However, our proposed method uses semantic information to avoid repeated visits to AoIs and to plan informative trajectories during exploration. As a result, our method achieves a **24%** shorter overall path length with a higher success rate and a comparable exploration time compared with the state-of-the-art method [4]. We also observed that our proposed method outperforms other methods in Tab. I more as the environment size increase.

C. Real World Experiments

The Falcon 250 UAV used for our experiments is a custom designed platform with a 402 mm tip-to-tip diameter with 6 in propellers, weighing 1.29 kg including a 331 g 3S 5200 mAh Li-Po battery. It carries an Intel Realsense camera and an onboard computer as shown in Fig.1. For real-

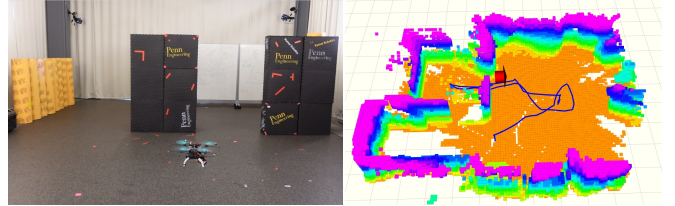


Fig. 7: Experiment in a structured 3-D indoor environment. The environment contains a room, a corridor and an open space. We bound the environment with a $12 m \times 8 m \times 2.5 m$ map. The right figure shows the occupancy map after exploration, executed trajectories (blue), and a detected door (red arrow).

world experiments, we use the odometry information from the Vicon Motion System. We utilized ONNX to run the trained network model onboard the quadrotor platform. The inference runs at 4 Hz with 4 threads on an i7-10710U CPU. Due to the space limit, we set the box-bounded maximum velocity at $0.4m/s$ and accelerate to $0.5m/s^2$. We conducted extensive real-world experiments in a $12m \times 8m \times 2.5m$ space to validate our proposed framework. As the robot takes off from the corridor and moves forward, it detects a door, navigates towards it, and confirms it. After the door is confirmed, the robot enters the room and completely explores it by navigating to the frontiers that maximize the utilities. The robot then exits the room back to the corridor, selecting the frontier with the highest utility. All utilities are computed considering predicted information gain during this process. The mapping result and the executed trajectories are shown in Figure. 7. Real-world experiments prove that our proposed framework is capable of handling realistic structured indoor environments. For more information, we kindly invite the reader to take a look at the video material at <https://youtu.be/5ZBkJmCKYwg>.

VII. CONCLUSION

In this paper, we develop a novel 3-D exploration framework for autonomous indoor exploration with MAVs. In particular, we propose an incremental detection and prediction module and a perception-aware planning module to enable safe and fast indoor exploration. Our proposed framework shows advantages in total trajectory length and success rate compared with classical and state-of-the-art methods. We validate our proposed framework in real-world experiments in which our proposed framework runs fully onboard a SWaP constrained platform.

REFERENCES

- [1] K. Saulnier, N. Atanasov, G. J. Pappas, and V. Kumar, "Information theoretic active exploration in signed distance fields," in *2020 IEEE*

- International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4080–4085.
- [2] 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, no. 2, pp. 1500–1507, 2020.
 - [3] B. Charrow, G. Kahn, S. Patil, S. Liu, K. Goldberg, P. Abbeel, N. Michael, and V. Kumar, “Information-theoretic planning with trajectory optimization for dense 3d mapping,” in *Robotics: Science and Systems*, vol. 11, 2015, pp. 3–12.
 - [4] B. Zhou, Y. Zhang, X. Chen, and S. Shen, “Fuel: Fast uav exploration using incremental frontier structure and hierarchical planning,” *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 779–786, 2021.
 - [5] 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)*, 2017, pp. 2135–2142.
 - [6] M. Dharmadhikari, T. Dang, L. Solanka, J. Loje, H. Nguyen, N. Khedekar, and K. Alexis, “Motion primitives-based path planning for fast and agile exploration using aerial robots,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 179–185.
 - [7] C. Gomez, M. Fehr, A. Millane, A. C. Hernandez, J. Nieto, R. Barber, and R. Siegwart, “Hybrid topological and 3d dense mapping through autonomous exploration for large indoor environments,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 9673–9679.
 - [8] S. Shen, N. Michael, and V. Kumar, “Autonomous indoor 3d exploration with a micro-aerial vehicle,” in *2012 IEEE international conference on robotics and automation*. IEEE, 2012, pp. 9–15.
 - [9] R. Shrestha, F.-P. Tian, W. Feng, P. Tan, and R. Vaughan, “Learned map prediction for enhanced mobile robot exploration,” in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 1197–1204.
 - [10] C. Wang, D. Zhu, T. Li, M. Q.-H. Meng, and C. W. de Silva, “Efficient autonomous robotic exploration with semantic road map in indoor environments,” *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2989–2996, 2019.
 - [11] C. Gomez, A. C. Hernandez, and R. Barber, “Topological frontier-based exploration and map-building using semantic information,” *Sensors*, vol. 19, no. 20, 2019.
 - [12] M. Saroya, G. Best, and G. A. Hollinger, “Online exploration of tunnel networks leveraging topological cnn-based world predictions,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 6038–6045.
 - [13] 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, pp. 146–151.
 - [14] F. Fraundorfer, L. Heng, D. Honegger, G. H. Lee, L. Meier, P. Tanakanen, and M. Pollefeys, “Vision-based autonomous mapping and exploration using a quadrotor mav,” in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 4557–4564.
 - [15] 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, pp. 1462–1468.
 - [16] A. Bircher, M. S. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, “Receding horizon path planning for 3d exploration and surface inspection,” *Autonomous Robots*, vol. 42, 02 2018.
 - [17] B. Charrow, S. Liu, V. Kumar, and N. Michael, “Information-theoretic mapping using cauchy-schwarz quadratic mutual information,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4791–4798.
 - [18] 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, pp. 1699–1706, 2019.
 - [19] Z. Meng, H. Qin, Z. Chen, X. Chen, H. Sun, F. Lin, and M. H. Ang, “A two-stage optimized next-view planning framework for 3-d unknown environment exploration, and structural reconstruction,” *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1680–1687, 2017.
 - [20] D. P. Ström, F. Nenci, and C. Stachniss, “Predictive exploration considering previously mapped environments,” in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 2761–2766.
 - [21] A. J. Smith and G. A. Hollinger, “Distributed inference-based multi-robot exploration,” *Autonomous Robots*, vol. 42, no. 8, pp. 1651–1668, 2018.
 - [22] J. M. Pimentel, M. S. Alvim, M. F. Campos, and D. G. Macharet, “Information-driven rapidly-exploring random tree for efficient environment exploration,” *Journal of Intelligent & Robotic Systems*, vol. 91, no. 2, pp. 313–331, 2018.
 - [23] K. Katyal, K. Popek, C. Paxton, P. Burlina, and G. D. Hager, “Uncertainty-aware occupancy map prediction using generative networks for robot navigation,” in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 5453–5459.
 - [24] G. Georgakis, B. Bucher, A. Arapin, K. Schmeckpeper, N. Matni, and K. Daniilidis, “Uncertainty-driven planner for exploration and navigation,” in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 11 295–11 302.
 - [25] G. Georgakis, B. Bucher, K. Schmeckpeper, S. Singh, and K. Daniilidis, “Learning to map for active semantic goal navigation,” 2021. [Online]. Available: <https://arxiv.org/abs/2106.15648>
 - [26] L. Wang, H. Ye, Q. Wang, Y. Gao, C. Xu, and F. Gao, “Learning-based 3d occupancy prediction for autonomous navigation in occluded environments,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 4509–4516.
 - [27] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, “Matterport3d: Learning from rgb-d data in indoor environments,” *International Conference on 3D Vision (3DV)*, 2017.
 - [28] O. Ronneberger, P. Fischer, and T. Brox, “U-net: Convolutional networks for biomedical image segmentation,” in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
 - [29] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “DeepLab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 4, pp. 834–848, 2018.
 - [30] B. Zhou, J. Pan, F. Gao, and S. Shen, “Raptor: Robust and perception-aware trajectory replanning for quadrotor fast flight,” *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1992–2009, 2021.
 - [31] S. Liu, M. Watterson, K. Mohta, K. Sun, S. Bhattacharya, C. J. Taylor, and V. Kumar, “Planning dynamically feasible trajectories for quadrotors using safe flight corridors in 3-d complex environments,” *IEEE Robotics and Automation Letters (RA-L)*, pp. 1688–1695, 2017.
 - [32] Z. Wang, H. Ye, C. Xu, and F. Gao, “Generating large-scale trajectories efficiently using double descriptions of polynomials,” 2021.
 - [33] L. Han, F. Gao, B. Zhou, and S. Shen, “Fiesta: Fast incremental euclidean distance fields for online motion planning of aerial robots,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 4423–4430.
 - [34] A. Rasouli and J. K. Tsotsos, “The effect of color space selection on detectability and discriminability of colored objects,” 2017. [Online]. Available: <https://arxiv.org/abs/1702.05421>