

# Bag of Views: An Appearance-based Approach to Next-Best-View Planning for 3D Reconstruction

Sara Hatami Gazani<sup>1</sup>, Matthew Tucsok<sup>2</sup>, Iraj Mantegh<sup>3</sup>, and Homayoun Najjaran<sup>1</sup>

**Abstract**—UAV-based intelligent data acquisition for 3D reconstruction and monitoring of infrastructure has experienced an increasing surge of interest due to recent advancements in image processing and deep learning-based techniques. View planning is an essential part of this task that dictates the information capture strategy and heavily impacts the quality of the 3D model generated from the captured data. Recent methods have used prior knowledge or partial reconstruction of the target to accomplish view planning for active reconstruction; the former approach poses a challenge for complex or newly identified targets while the latter is computationally expensive. In this work, we present Bag-of-Views (BoV), a fully appearance-based model used to assign utility to the captured views for both offline dataset refinement and online next-best-view (NBV) planning applications targeting the task of 3D reconstruction. With this contribution, we also developed the View Planning Toolbox (VPT), a lightweight package for training and testing machine learning-based view planning frameworks, custom view dataset generation of arbitrary 3D scenes, and 3D reconstruction. Through experiments which pair a BoV-based reinforcement learning model with VPT, we demonstrate the efficacy of our model in reducing the number of required views for high-quality reconstructions in dataset refinement and NBV planning.

**Index Terms**—Reactive and sensor-based planning, aerial systems: perception and autonomy, intelligent data acquisition.

## I. INTRODUCTION

**A**CTIVE vision is characterized by the ability of a robot to make decisions about placing or reconfiguring its sensors to complement its perception of the environment [1]. This ability leads to meaningful actions of the robot based on interpretations of its surrounding environment and grants the robot a planning strategy, namely view planning, for actively replacing its sensors to uncover most amount of information about the target. In the context of active 3D reconstruction, view planning is used to optimize the robot’s path until the task requirements are satisfied. For the application of 3D reconstruction of infrastructure using UAV-based imaging which is the concern of this work, the view planning problem dictates the data acquisition process and significantly impacts the reconstruction results. Previous work in this domain either relies on a given proxy of the target to build upon while

<sup>1</sup>S. Hatami and H. Najjaran are with the Department of Mechanical Engineering, University of Victoria, Victoria, BC, CA, V8P 5C2 (e-mail: sarahatami@uvic.ca; najjaran@uvic.ca).

<sup>2</sup>M. Tucsok is with the Okanagan School of Engineering, University of British Columbia, Kelowna, BC, CA, V1V 1V7 (e-mail: tucsok@student.ubc.ca).

<sup>3</sup>I. Mantegh is with the National Research Council (NRC) Canada, QC, CA (e-mail: iraj.mantegh@nrc-nrc.gc.ca).

The authors have provided supplementary material (code and video) available at <https://github.com/ACIS2021/ViewPlanningToolbox> and [https://youtu.be/RhK2\\_HhaJoo](https://youtu.be/RhK2_HhaJoo).

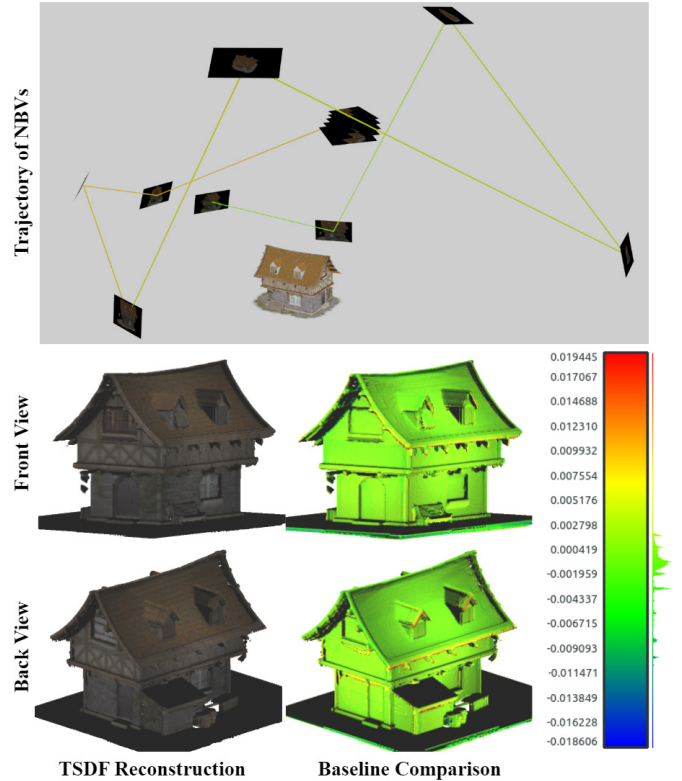


Fig. 1. Results of the proposed Bag-of-Views model applied to next-best-view planning trained and evaluated using our View Planning Toolbox. The output reconstruction shows a 94.6% surface coverage by only 13 captured views with 2.78 cm Hausdorff distance and 0.80 cm Chamfer discrepancy as compared to the results from a complete coverage baseline scan of the target with 288 views.

planning the views [2], [3] or generates a partial reconstruction using the knowledge of the agent about the target so the camera is navigated to complete the model [4], [5]. In this setting, the agent iteratively calculates the next waypoint to attend where it can capture the next-best-view (NBV) with the highest predicted information gain. However, when the purpose is to capture views from newly recognized targets or in the case of targeting complex structures, a geometric proxy of the target might not be accessible and online 3D reconstruction to achieve guidance can be computationally expensive and time-inefficient. On the other hand, under the assumption that adequate computation resources exist onboard the drone, algorithms that use an external model for guidance purposes mostly focus on the coverage completeness of the area where less attention is paid to the relative visual information contained in consecutively captured views.

In this work, we propose a novel approach to a fully

appearance-based view planning for online NBV planning and offline dataset refinement applications. The key uniqueness of our work lies on its model-free nature that makes it independent of the true state of the environment, namely the actual 3D model, both during training and at inference time. This allows the method to be applied to a wide range of settings and customized to fit different applications. In order to introduce the concept, we draw parallels between computational representation of views and how humans perceive objects by recognizing and interpreting distinct visual cues [6]. We employ the Bag-of-Visual-Words [7] as a simplified vision model to represent an agent's knowledge of the target and introduce spatial information to the visual vocabularies to form a Bag-of-Views (BoV), a model to record and retrieve the visual features of the target from different viewpoints. In the context of the application explored in this study where UAV flight duration is restricted due to power consumption concerns, the agent's task is to identify viewpoints that avoid having the camera repeatedly traverse back and forth around the target. Instead, it should chart a continuous and consistent path around the target. In summary, our approach is distinguished by two distinctive attributes: i) it does not necessitate real-time reconstructions during training or inference, and ii) it ensures a smooth and uninterrupted planned trajectory. First, through experimental results, we demonstrate how the selection of the views using our appearance-based heuristics affects the 3D reconstruction process and use the observations to refine already acquired datasets. Next, we bring this approach to an active level and utilize a Soft Actor-Critic (SAC) method [8] to train an agent that seeks to capture the target from views that cause a more drastic change in how it remembers the environment so far. The core idea of this method is to use the local visual features of the scene and their positioning to guide the agent to predict the NBVs that would result in revealing the highest number of unseen visual features as the agent remembers them. In addition, we present the View Planning Toolbox (VPT) as a comprehensive solution for simulating the view capturing process as part of developing view planning algorithms. The introduction of this toolbox fills a significant gap in the area of simulating machine learning-based view planning frameworks and provides developers with an open-source, user-friendly solution to streamline their training and evaluation process. The main contributions of this work can be summarized as follows:

- Bag-of-Views, a novel fully appearance-based view selecting model for offline dataset refinement. This model needs no pre-training and is modular and customizable for different applications.
- A reinforcement learning approach to appearance-based NBV planning with no complete or partial reconstructions required during training or at inference time.
- View Planning Toolbox (VPT), which provides an environment for training and testing of machine learning-based view planning and 3D reconstruction models.

## II. RELATED WORK

### A. Solutions to the View Planning Problem

Based on the amount of available information to the system about the environment and visual complexity of the target, view planning approaches are often categorized as model-based or model-free (a.k.a. non-model-based) approaches [9]. In model-based view planning problems, a viewing plan is obtained using a previously built or given model of the target [2], [3]. In more general applications and in cases where the target is unknown or introduced to the system at runtime, the viewing strategy should be generated without prior information of the target [1]. In such cases, the goal is to manipulate the camera position and orientation in a manner that most unrevealed information about the target is exposed to the agent in each step. Most model-free methods follow the NBV approach. Typically, these systems build an interpretation of the environment using acquired information and base the planning of the next view(s) on it. That interpretation of the scene can be in different forms based on the specific task and application. Among methods that follow this approach are frontier-based and volumetric-based methods. In frontier-based view planning first introduced in [10], the core idea was to visit unexplored regions of an initial map, i.e. frontiers, to update the map based on the newly collected information in those regions. Methods belonging to this category usually represent the target zone of the environment using a 2D [11] or 3D [4] occupancy grid or directly use a point cloud to map the boundary between explored and unexplored regions [12]. As opposed to frontier-based methods that mostly focus on exploring an environment, volumetric-based approaches are usually concerned with modelling a single target and focus more on the completeness of the coverage. In this regard, to guarantee a successful registration, overlaps between the views are also considered while planning the views [13]. In addition to 3D modelling, such algorithms are also used for scene inspection [14], [15]. A subcategory of the model-free view planning algorithms, namely appearance-based planning methods, carries out the decision making process based solely on the visual inputs such as RGB or grayscale images [16], [5], [17]. These methods either rely on an *a priori* model of the target or a partial reconstruction of the scene, at least during the training stage of their models. For example, [16] proposed a method to compute a candidate sequence of viewpoints for a micro aerial vehicle to attend based on the information gain from different camera poses. More recently, with the advancements of deep reinforcement learning, appearance-based view planning has been visited more often. Accordingly, [5] utilized only the captured images to plan the views without tracking a partial reconstruction of the true 3D model. They used the surface coverage percentage to guide the agent. Similarly, [17] used the surface coverage as well as reconstruction error as part of the reward for guiding their reinforcement learning agent towards complete reconstruction. Unlike these methods, our proposed model omits the need for the true state of the environment or any partial reconstructions for guiding the agent towards capturing high-utility views for the task of reconstruction. We treat 3D reconstruction as a downstream task

as opposed to a parallel task in the view planning framework and pay more attention to the conditions that should be met by the views in order to have a satisfactory reconstruction.

### B. Simulation Environments for The View Planning Problem

View planning research has explored various simulation environments, each exhibiting unique advantages and constraints. For instance, the PredRecon framework [18] utilized AirSim [19] for its simulation environment, taking advantage of its high-fidelity simulation capabilities. However, their work also relied on Unreal Engine 4 (UE4) [20] for data generation of other 3D models and Blender [21] for partial pointcloud reconstructions. This use of multiple software environments can be resource-intensive and complex to set up. In a similar manner, [22] used Gazebo [23]. While Gazebo is a feature-rich simulator providing a versatile environment for robotic simulations, using it for view planning in photo-realistic environments introduces computational overhead and performance degradation.

Previous work addressed the need for flexibility and customization by employing environments inspired by OpenAI Gym [24] to train their reinforcement learning agent for NBV planning [17]. Their approach offered greater flexibility in defining the simulation environment and incorporating different perceptual elements. Additionally, the *gym-collision-avoidance* package [25] used in [26] provided a simulation environment for informative trajectory planning. With a greater focus on collision avoidance, this package allowed researchers to simulate and evaluate view planning algorithms within a controlled environment.

It is worth noting that simulation environments often require dealing with trade-offs between fidelity, computational resources, and the level of control provided over environmental factors. Striking the right balance between these aspects is dependent on the specific task being studied. For view planning, factors affecting high-level control of the simulation are of most importance. These factors include camera controls (pose, focal length, resolution), scene manipulation (lighting, object placement and scaling), and scripting tools for automating the aforementioned functionalities. To address these trade-offs, we have developed the View Planning Toolbox (VPT), a comprehensive solution outlined in Sec. V.

### III. BAG OF VIEWS: APPEARANCE-BASED APPROACH FOR SELECTING BEST VIEWS

Following the work in [27], we introduce a computational representation of the views in terms of the visual features of the scene captured from the respective viewpoints. As discussed in [27], two key conditions must be satisfied for a successful multi-view 3D reconstruction:

- The views in an input set must present features of the target that are distinct from the ones presented by other views in the same set,
- Each view by itself must be rich in the number of visual features it is revealing of the target.

We denote the  $i^{\text{th}}$  view in the set as  $\chi_i$ . Each view  $\chi_i$  is encoded to and is represented by its extracted features using a feature extracting algorithm such as Scale-Invariant Feature

Transform (SIFT) [28]. Thus,  $\chi_i$  will be a 2D matrix with dimensions  $m \times n$  where  $m$  is the number of detected keypoints and  $n$  is the number of values in the feature descriptors. In the case of SIFT, each row of this view matrix represents a 128-dimensional feature descriptor where each element represents a certain attribute of the local feature detected in the image patch. We can denote each view by iterating over its resulting feature descriptors as  $\chi_i(j, :) = \{f(j, k) : k \in \{1, 2, \dots, n\}\}$  where  $f(j, k)$  is the  $k^{\text{th}}$  value in the  $j^{\text{th}}$  descriptor of the image. These conditions for a set of views result in a greater distance between corresponding descriptors from two view representations denoted as  $\text{dist}(\chi_i(j, :), \chi_{i+1}(j', :))$  for consecutively selected random views  $\chi_i$  and  $\chi_{i+1}$ . A greater distance ensures a better reconstruction quality for a limited-length trajectory. In our case, the cosine distance metric is used to measure the dissimilarity between the descriptors and the visual words since its consistent range of outputs allows for easy interpretation and score comparison. Thus,

$$\text{dist}(\chi_i(j, :), \chi_{i+1}(j', :)) = 1 - 2 \times \cos(\chi_i(j, :), \chi_{i+1}(j', :)) \quad (1)$$

where  $j \in \{1, 2, \dots, m\}$  and  $j' \in \{1, 2, \dots, m'\}$  with  $m$  and  $m'$  being the number of representative descriptors of  $\chi_i$  and  $\chi_{i+1}$ :

$$\cos(\chi_i(j, :), \chi_{i+1}(j', :)) = \frac{|\chi_i(j, :)| \cdot |\chi_{i+1}(j', :)|}{|\chi_i(j, :)| \cdot |\chi_{i+1}(j', :)|} \quad (2)$$

Since the extracted features belonging to a target are prone to self-similarity, the feature descriptors of different regions of the target can be clustered based on their similarity and, instead of pair-wise comparisons of all feature descriptors belonging to all views in the set, they can be compared to the cluster cores. In addition, the necessity of there being a model-free view planning model with no pre-training requires learning these cluster cores iteratively from the incoming information. This clustering of the visual features is inspired by [29] where cluster centers of the quantified feature descriptors were used to form a Bag of Visual Words (BoVW).

In the context of view planning for a single target with one or more symmetry axes, learning a global visual vocabulary for the entire model can be prone to overconfidence in recognizing certain visual words. Therefore, we propose a method to track visual features of the target captured from different viewpoints via distinguished visual vocabularies for different regions included in what we call a Bag-of-Views (BoV). As well as mitigating the symmetry challenge, the computation cost will be significantly less with the viewpoint-based vocabularies; a smaller number of visual words is required to describe a portion of the target rather than all parts of it and the new view will only update its corresponding vocabulary among the BoV. Below are the steps involved:

**1) Feature Extraction:** Using a feature extraction algorithm, in our case SIFT, local features of the captured image at position  $T$  are extracted in the form of 128-dimensional vectors.  $T$  is the position of the camera in the spherical coordinate system with its origin located at the center of the scene.

**2) Utility Assignment:** Depending on the application at hand, assigning utilities to the views can be divided into two different cases:

**2.I)** In the case of dataset refinement where the views have already been captured, a decision should be made about including each view in the input set of the reconstruction algorithm based on its utility. The question simply is "does this new view help the reconstruction process?". To answer that, we look at the extent to which this view satisfies the two conditions previously mentioned. Each of the feature descriptors from the previous step are compared with their closest visual word through vector quantization [30]. Then, a distance metric is used to measure the dissimilarity between feature descriptors and the supposed visual word that would represent them in the corresponding vocabulary in the BoV, denoted as  $\nu_{id_T}$ . Here,  $id_T$  is the identification number for the vocabulary associated with the region that  $T$  belongs to. This process is repeated for all of the descriptors and the sum of the dissimilarity scores is used to decide whether to ignore or utilize the view in the reconstruction process. If the final score is a non-zero positive value, it is included in the set and proceeds to the third step, otherwise it is ignored.

**2.II)** In the second case, we use the BoV model to train a reinforcement learning agent to propose NBVs based on the appearance of consecutively captured views. Further elaboration on the learning process will be provided in Sec. IV. In this section, our focus lies primarily on the utilization of this model to shape the reward function within the specified context. While training the reinforcement learning agent, we use the change in the corresponding vocabulary of the BoV after capturing a new view at location  $T$  to shape the reward function. A greater difference between the new and previous BoV implies that the new view contains more unseen features and results in higher rewards.

**3) View Representation:** This step is a continuation of case 2.I. Depending on the position and orientation of the captured viewpoint, the extracted features update the specific vocabulary assigned to the range of views that the new view

belongs to. This updating includes clustering of the descriptors belonging to that region using a clustering algorithm such as K-means [31]. Thus, every group of the cluster centers in the BoV, namely every regional vocabulary, describes the appearance of the target from viewpoints that are close in position and orientation. Algorithm 1 showcases the pseudo-code for the creating a BoV model.

#### IV. A REINFORCEMENT LEARNING APPROACH TO APPEARANCE-BASED NBV PLANNING

The problem of NBV planning is a sequential decision making process that can be defined as a Partially Observable Markov Decision Process (POMDP) and be solved through reinforcement learning algorithms. Our goal is to achieve this without any need for *a priori* knowledge of the target and without any full or partial reconstruction of the target during training or inference time. We begin by formulating different components of this process. The goal of the agent in our system is to iteratively propose next views for a limited number of steps to reach regions with a high number of features unfamiliar to the BoV. Seeking such views leads to drastic changes in the vocabularies of the BoV through each relocation of the camera.

The state space should provide the agent with enough information about the environment to enable meaningful actions towards the goal. We use the concatenation of down-sampled grayscale images captured through the last  $\tau$  consecutive frames as well as the concatenation of the normalized camera locations associated with each view. Thus, the state  $s_t$  at time  $t$  is defined as  $\{T_{t-\tau:t}, obs_{t-\tau:t}\}$ . The camera location  $T_t$  is presented using the spherical coordinate system in the form of  $\{R, \phi, \theta\}$  with three values for radial distance from the center, the azimuth angle, and the elevation angle. Also, given a deterministic state transition, the action  $a_t$  determines the next camera location  $T_{t+1}$ . Depending on the policy network's architecture,  $T_{t+1}$  may need to be adjusted to meet safety requirements. In our case, it is rescaled to conform to the specified ranges for its three components. Based on the introduction in case I of Sec. III, the reward received for this action represents the change in the part of the BoV that has been influenced by the new action, namely  $\nu_{T_{t+1}}$  which is the vocabulary associated with the region that the new location belongs to. This change is measured by comparing the same regional vocabulary before and after taking the action; the closest visual words in the two vocabularies are identified and their distance is measured through vector quantization with the cosine distance metric. The sum of these distances is used to represent the change in the vocabulary after taking the action. We also add a negative constant reward at each time step. Thus, we define the reward to be  $r_{t+1} = dist(\nu_{T_{t+1}}, \nu_{T_t}) - 1$ , with the distance between vocabularies being defined as:

$$dist(\nu_{T_{t+1}}, \nu_{T_t}) = \sum_i (1 - 2 \times \cos(\nu_{T_{t+1}}(i), \nu_{T_t}(k))) \quad (3)$$

where  $k = \arg \max_j \cos(\nu_{T_{t+1}}(i), \nu_{T_t}(j))$  and  $i$  iterates over each of the visual words in  $\nu_T$ . The number of vocabularies in the BoV and the size of each vocabulary are dependant on the resources available during training and runtime.

---

#### Algorithm 1 Bag-of-Views Model for Offline View Selection

---

**Require:** Number of view ranges  $N$ , Number of words  $W$ ,

Database  $\mathcal{D}$

```

1: Initialize BoV $\{\nu_{1:N}\}$ , Feature extractor  $\mathcal{F}$ 
2: while There is data to process in  $\mathcal{D}$  do
3:    $T, X \leftarrow$  Load data from  $\mathcal{D}$ 
4:    $id_{view} \leftarrow T.azimuth // \frac{2\pi}{N}$ 
5:    $\chi \leftarrow \mathcal{F}(X_i)$ 
6:    $dist \leftarrow 0$ 
7:   for  $j = 1$  to  $W$  do
8:      $C \leftarrow \arg \max_k (\cos(\chi(j), \nu_{id}(k)))$ 
9:      $dist \leftarrow dist + (1 - 2 \times \cos(\chi(j), \nu_{id}(C)))$ 
10:    if  $dist > 0$  then
11:      Update descriptors with  $\chi$  for  $\nu_{id_{view}}$ 
12:       $\nu_{id_{view}}.codebook \leftarrow$  Perform K-means on  $\nu_{id_{view}}$ 
13:    else
14:      Remove  $X$  from  $\mathcal{D}$ 

```

---

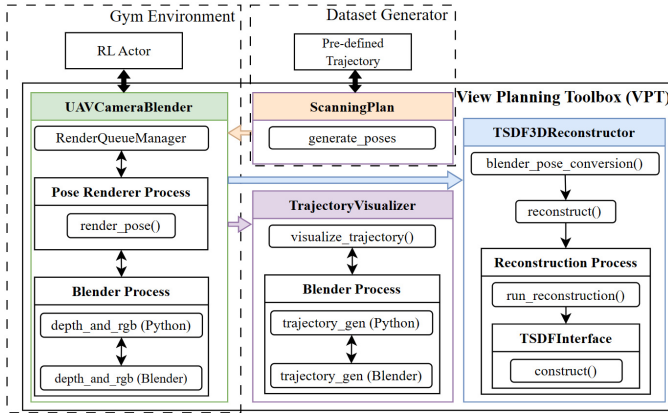


Fig. 2. Simplified visualization of the interactions between different components of the View Planning Toolbox (VPT) for reinforcement learning and dataset generation.

## V. SIMULATION ENVIRONMENT

To address the need for a lightweight, flexible, and easy-to-integrate simulation environment, we introduce the View Planning Toolbox (VPT). VPT contains tools for simulating the components required to train, visualize, and test view planning algorithms entirely within Python [32] and Blender [21]. At the core of VPT is the *UAV Camera*, which simplifies the visual data acquisition process carried out by an Uncrewed Aerial Vehicle (UAV) down to a camera floating in 3D space, where the specifications and pose of the camera can be determined either manually, programmatically, or through the use of a view planning algorithm to generate arbitrary 3D trajectories.

### A. UAV Camera

The *UAV Camera* simulates the base functionality of a drone traversing a 3D scene to gather visual information including both RGB and depth images from the environment. It is a simple pinhole model where the developer specifies the resolution of the captured image in pixels, the focal length in millimeters, and the depth range in meters. The *UAV Camera* implementation encapsulates BlendTorch [33], a framework designed to integrate PyTorch [34] deep learning models with Blender [21]. The BlendTorch framework provides a convenient, multiprocessing-safe means of communication between a Blender environments and native Python applications. In addition, the *UAV Camera* can be instantiated in an OpenAI gym [24] environment for use with reinforcement learning models.

### B. Data Generation

VPT serve a dual purpose as shown in Fig. 2:

- Generate state-action pair experiences for training reinforcement learning models. This requires accessing the *UAV Camera* directly which can be instantiated in a Gym environment.
- Generate offline datasets for training supervised learning models. This functionality is contained within *Scanning Plans* and adheres to traditional photogrammetric principles [35], ensuring proper determination of end and side overlap for aerial scanning.

As a way of evaluating our appearance-based view planning model in Sec. III, we compare the resultant reconstructions from the captured views with those generated by views that fully cover the targets. These views are generated using the *hemispherical scan* tool in VPT, which maintains a constant end and side overlap between images at a distance  $R$  from the center of the scene. The spacing of views around the hemisphere are calculated by making a series of flat plane and small angle approximations due to the high overlap between views. Each scene contains a photo-realistic structure which has been centered about the z-axis and lies on the XY-plane. The structure is also normalized by its largest dimension. This ensures structures of all sizes can be fully visible from all viewpoints. A resulting trajectory of views generated with this method can be seen in Fig. 3.

### C. TSDF-Fusion 3D Reconstruction

To evaluate the performance of our view planning algorithms, VPT integrates a Truncated Signed-Distance Function (TSDF) Fusion 3D reconstruction module from [36]. A wrapper has been implemented to isolate the CUDA environment of the TSDF-Fusion module to prevent conflicts with PyTorch. The TSDF-Fusion module creates either a 3D point cloud or a 3D mesh using the marching cubes method [37]. TSDF-Fusion reconstruction requires RGB views, the corresponding depth maps, and the computed  $4 \times 4$  rigid transformation matrices based on the poses for the corresponding views. This module acts as a lightweight photogrammetric stand-in to quickly evaluate the reconstruction quality from hundreds of views.

## VI. IMPLEMENTATION DETAILS AND EXPERIMENTS

### A. Network Architecture and Training

The policy network is composed of two different sub-networks for processing the two components of the state; the observation network and the location network. Details of these two sub-networks are shown in Fig. 4. We used the Soft Actor-Critic (SAC) algorithm by [8] to train the policy network. SAC is powered by the advantages of actor-critic methods and maximum entropy reinforcement learning. It is a model-free, off-policy algorithm that optimizes both the policy and the value function simultaneously through a combination of policy gradient and Q-value updates. It maximizes the expected return while incorporating a soft entropy regularization term [38].

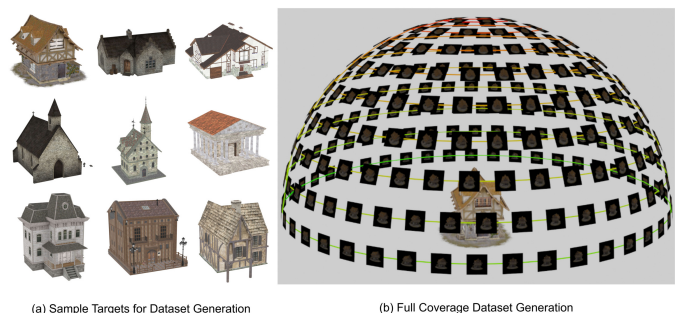


Fig. 3. VPT performs full coverage hemisphere scans in which the end and side overlap are specified on many targets.

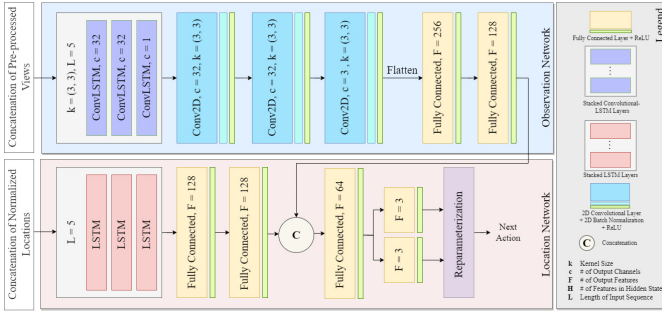


Fig. 4. Proposed architecture for the policy network. The policy network processes the two components of the state, namely the observations and locations throughout the last 5 frames, through two separate sub-networks and fuses the resultant feature vectors to produce the action distribution.

The critic network, which is responsible for mapping state-action pairs to their quality values, encodes the concatenation of the observation component of the state through a sequence of 3D convolutional layers followed by 3D batch normalization layers and ReLU activation functions. Then, the output is flattened and goes through two fully connected layers with an output size of 128. A sub-network also processes the concatenation of the location component of the state and the action through fully connected layers of output size 64 and 128, respectively. The resulting feature vector is then concatenated with the encoded observations and goes through two fully connected layers of output size 128 and 3 as the output quality value vector. Training is done using the hyperparameters listed in Table I and a simplified version of the proposed reinforcement learning cycle is depicted in Fig. 5.

TABLE I  
TABLE 1: HYPERPARAMETERS USED FOR TRAINING THE RL AGENT

Parameter	Value
Batch Size	8
Learning Rate	$5 \times 10^{-4}$
Temperature Parameter ( $\alpha$ )	0.2
Target Network Update Interval (steps)	20
Experience Replay Buffer Size	$10^6$
Soft Update Factor ( $\tau$ )	0.01
Discount Factor ( $\gamma$ )	0.99

### B. Evaluation of The View Planning Results

To evaluate the utility of the suggested views using the BoV model, we used the reconstruction method explained in Sec. V-C and compared the resulting 3D point clouds and meshes from the captured views with those from a complete coverage scan of the target discussed in Sec. V-B. We analyzed the reconstruction results of the offline dataset refinement and online NBV planning using the Chamfer discrepancy, Hausdorff distance, and the mesh-to-mesh comparison carried out in CloudCompare [39]. Chamfer discrepancy gives an overall measure of similarity or dissimilarity between two point sets, while Hausdorff distance focuses on the largest observed distance, highlighting extreme differences. When comparing scanning results using these metrics, a higher Chamfer discrepancy indicates poorer coverage of the target and lower overall accuracy while a higher Hausdorff distance suggests that there are missed areas in the scanning process, leading to a reconstruction with holes.

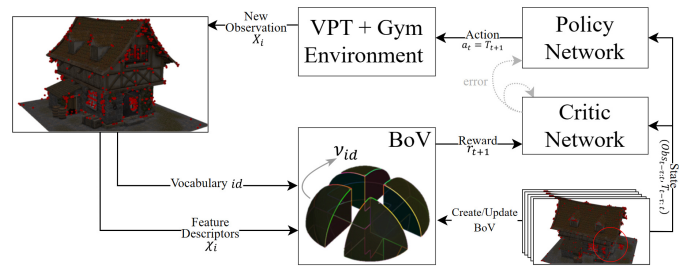


Fig. 5. Block diagram of the proposed reinforcement learning approach to NBV planning based on Bag-of-Views model.

### C. Bag-of-Views for Dataset Refinement

To study the effect of the size of BoV, we tested the model performance in filtering two of the generated datasets consisting of 288 views uniformly covering two of the buildings in our dataset as explained in Sec. V-B. We used the method explained in Sec. III and Alg. 1 to reduce the dataset size and evaluate reconstruction performance resulting from the remaining views. Two effective variables in this study were the number of view ranges defining the size of BoV and the number of visual words in each vocabulary of the BoV. The reconstruction results were compared based on their Hausdorff distance and Chamfer discrepancy with the 3D point cloud reconstructed using the original dataset of 288 views. The results are listed in Table II. In addition to this quantitative comparison, we visualized the distance between the mesh reconstructions to those produced using the original dataset in Fig. 6.

This analysis explored the question of achieving optimal performance by balancing model efficiency (number of selected views) and reconstruction quality. Shown in Fig. 6, a larger BoV, achieved by increasing the number of vocabularies, resulted in a reconstruction that exhibits reduced spatial sparsity in the error surrounding the model. This interprets as a more uniform scan of the target when seeking to update visual vocabularies that are defined for a smaller view range. This means that for a fixed number of words in each vocabulary, each view has less opponents to be compared with and the resulting visual words are more local to that region. While each view has a higher chance to represent a certain view range in the vocabulary, there is a lower chance for its similar views to be accepted into the set, as the dominant local features have already been identified. This effect was reinforced when the number of visual words per each vocabulary in the BoV was increased. The higher the number of visual words attributed to each view range, the higher the chance of familiarity of the newly captured view and details exposed to it.

### D. Appearance-based Next-Best-View Planning: Baseline Comparison and Generalization

Based on the problem formulation in Sec. IV and the training algorithm in Sec VI-A, we trained an agent in a reinforcement learning cycle depicted in Fig. 5 to scan multiple structures from the dataset we generated using VPT and publicly available 3D models.

We set the termination condition as the completion of one pass around the target to demonstrate the ability of the

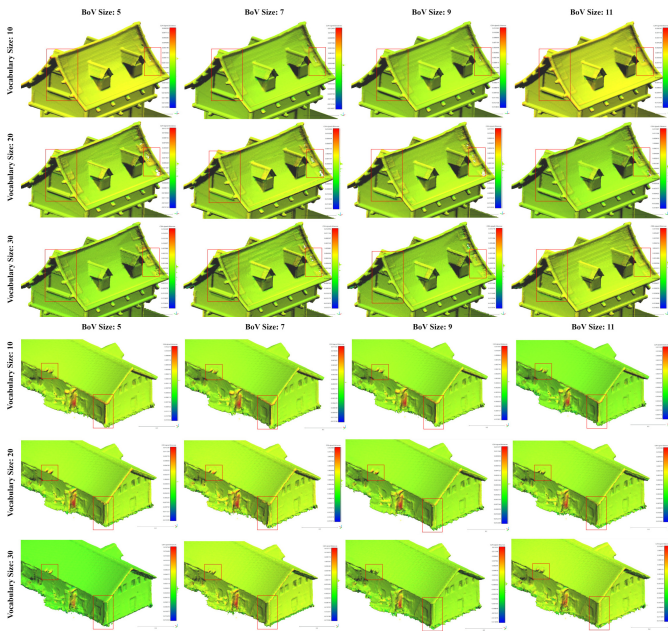


Fig. 6. Studying the effect of the size of BoV and the vocabularies within it. It can be seen that increasing the size of BoV lowers the error sparsity while increasing the vocabulary size reduces the error values.

model to iteratively find the optimal views, and compared the resulting reconstructions with the results of a full coverage scan of the target as the baseline. The resultant trajectory for a sample from the dataset is shown in Fig. 1. Furthermore, we applied the policy trained on a single structure in the initial phase to evaluate its performance on two unseen buildings. This allowed us to assess the policy’s ability to generalize across different structures and settings. Shown in Fig. 7, the results demonstrate the efficacy of the model in finding the optimal views which result in high-quality reconstructions. An average runtime of 0.893 seconds was measured while using an NVIDIA RTX 3060 with up to 16 GB of memory for the evaluation and testing of the model. This planning execution time is only a small fraction of the time required for scanning of an object in real-world conditions (i.e., on the order of several minutes). The implementation is thus practical with either an onboard GPU or a ground-based GPU calculating and communicating the mission with a UAV.

## VII. CONCLUSIONS AND FUTURE WORK

Our study tackled the challenge of model-free view planning by introducing a novel appearance-based computational representation of reconstruction targets that can be of utmost utility for UAV-based aerial photogrammetry. Our model enables

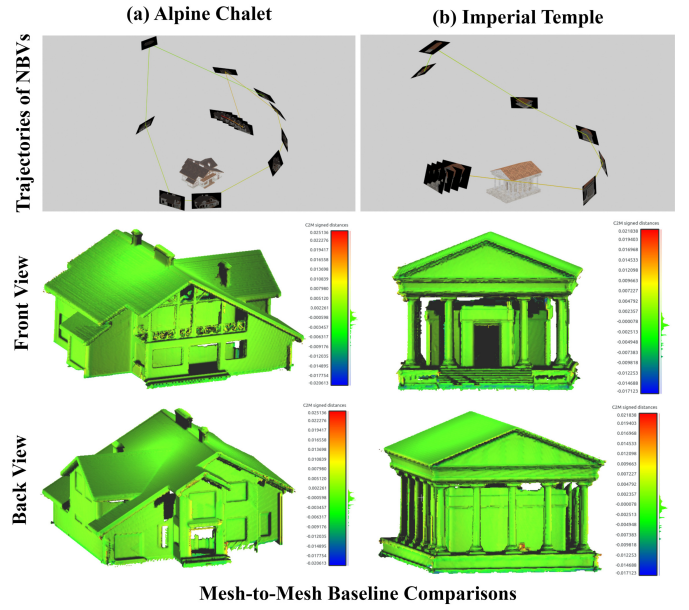


Fig. 7. Testing the generalizability of our RL-based NBV planner on two samples from our dataset. The resulting Hausdorff distance and Chamfer discrepancy of the *Alpine Chalet* reconstruction were calculated to be 9.47 cm and 1.01 cm while those of the *Imperial Temple* were 6.35 cm and 0.82 cm, respectively.

utility assignment to the views without tracking a full or partial reconstruction of the target through tracking unfamiliar visual features with its vocabularies contained within what we called a Bag-of-Views (BoV). We also developed the View Planning Toolbox (VPT), which offers a comprehensive solution for training, evaluation, and custom dataset generation in the context of view planning and 3D reconstruction. Our first set of experiments focused on exploring the size of BoV and the vocabularies within it on reconstruction quality. Using the reconstruction results from a complete coverage scan of the target as a baseline, we found that the BoV model achieved a remarkable reduction of views used for reconstruction (70.6% decrease) while simultaneously reducing the reconstruction error (33.5% decrease). These outcomes showcased the efficacy of our model in identifying optimal views for reconstruction. Building upon this proof of concept, we extended the application of the BoV model to shape the reward of a reinforcement learning (RL) agent trained using the Soft Actor-Critic (SAC) algorithm for online NBV planning. Once again, our model yielded high-quality reconstructions with a significantly low number of views (down to 5% of the number of baseline views). Furthermore, the RL model exhibited substantial generalizability to unseen targets.

TABLE II

STUDYING THE INFLUENCE OF THE SIZE OF BAG-OF-VIEWS ON THE DATASET SIZE AND THE RECONSTRUCTION RESULTS FROM THE TSDF FUSION OF RGB AND DEPTH IMAGES

Size of BoV	5	5	5	7	7	7	9	9	9	11	11	11
Number of Words per View Range	10	20	30	10	20	30	10	20	30	10	20	30
Medieval Tavern	Number of Selected Views	143	67	37	142	63	34	140	62	35	146	42
	Hausdorff Distance (cm)	2.1509	1.8134	1.8901	1.8045	2.3907	2.2756	1.8803	2.3693	1.5216	1.4825	1.7841
	Chamfer Discrepancy (cm)	0.1183	0.1935	0.4542	0.1206	0.4345	0.3742	0.1254	0.1906	0.2998	0.1109	0.1953
Forest Ruin House	Number of Selected Views	102	38	23	92	33	24	98	42	32	99	36
	Hausdorff Distance (cm)	2.5810	2.2692	2.5231	2.0577	2.1886	3.1446	2.1244	2.2862	2.3698	2.04848	2.6412
	Chamfer Discrepancy (cm)	0.2374	0.4554	0.6548	0.2416	0.5375	0.5775	0.1869	0.6057	0.5842	0.3115	0.5359

Notably, we discovered that the degree of generalizability depended on the relative visual complexity of the training and testing environments, further validating the effectiveness of our appearance-based view selection approach. We emphasize that our choice of simulated data was deliberate, allowing us to present a comprehensive and novel approach unencumbered by real-world data complexities. Promising future research can include using custom feature extractors and pre-training the visual vocabularies of the BoV for tracking certain visual features associated with structural defects in infrastructure as well as exploring real-world data evaluation. While we used SIFT feature extractor for creating BoV models which can cause vulnerability to environmental and lighting conditions, custom feature extractors can be separately developed to introduce robustness to such conditions.

## REFERENCES

- [1] S. Chen, Y. Li, N. M. Kwok, Active vision in robotic systems: A survey of recent developments, *The International Journal of Robotics Research* 30 (11) (2011) 1343–1377.
- [2] W. R. Scott, Model-based view planning, *Machine Vision and Applications* 20 (1) (2009) 47–69.
- [3] A. Aryan, F. Bosché, P. Tang, Planning for terrestrial laser scanning in construction: A review, *Automation in Construction* 125 (2021) 103551.
- [4] A. Kleiner, C. Dornhege, A frontier-void-based approach for autonomous exploration in 3d, *Advanced Robotics* 27 (11) (2011).
- [5] D. Peralta, J. Casimiro, A. M. Nilles, J. A. Aguilar, R. Atienza, R. Cajote, Next-best view policy for 3d reconstruction, in: A. Bartoli, A. Fusiello (Eds.), *Computer Vision – ECCV 2020 Workshops*, Springer International Publishing, Cham, 2020, pp. 558–573.
- [6] A. Torralba, A. Oliva, M. S. Castelhan, J. M. Henderson, Contextual guidance of eye movements and attention in real-world scenes: the role of global features in object search., *Psychological review* 113 (4) (2006) 766.
- [7] S. Lazebnik, C. Schmid, J. Ponce, Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories, Vol. 2, 2006, pp. 2169 – 2178.
- [8] T. Haarnoja, A. Zhou, P. Abbeel, S. Levine, Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor (01 2018).
- [9] M. Maboudi, M. Homaei, S. Song, S. Malihi, M. Saadatseresht, M. Gerke, A review on viewpoints and path-planning for uav-based 3d reconstruction, arXiv preprint arXiv:2205.03716 (2022).
- [10] 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.
- [11] E. Vidal, N. Palomeras, K. Istenič, J. D. Hernández, M. Carreras, Two-dimensional frontier-based viewpoint generation for exploring and mapping underwater environments, *Sensors* 19 (6) (2019).
- [12] S. Manandhar, R. Sadananda, Effect of hamming distance of patterns on storage capacity of hopfield network, 2002, pp. 253– 256 vol.1.
- [13] J. Vasquez-Gomez, L. Sucar, R. Murrieta-Cid, View/state planning for three-dimensional object reconstruction under uncertainty, *Autonomous Robots* 41 (01) (2017).
- [14] S. Song, S. Jo, Online inspection path planning for autonomous 3d modeling using a micro-aerial vehicle, in: 2017 IEEE International Conference on Robotics and Automation (ICRA), 2017, pp. 6217–6224.
- [15] L. Heng, A. Gotovos, A. Krause, M. Pollefeys, Efficient visual exploration and coverage with a micro aerial vehicle in unknown environments, in: 2015 IEEE International Conference on Robotics and Automation (ICRA), 2015, pp. 1071–1078.
- [16] C. Forster, M. Pizzoli, D. Scaramuzza, Appearance-based active, monocular, dense reconstruction for micro aerial vehicles, 2014.
- [17] S. Potapova, A. Artemov, S. Sviridov, D. Musatkina, D. Zorin, E. Burnaev, Next best view planning via reinforcement learning for scanning of arbitrary 3d shapes, *Journal of Communications Technology and Electronics* 65 (2020) 1484–1490.
- [18] C. Feng, H. Li, F. Gao, B. Zhou, S. Shen, Predrecon: A prediction-boosted planning framework for fast and high-quality autonomous aerial reconstruction, arXiv preprint arXiv:2302.04488 (2023).
- [19] S. Shah, D. Dey, C. Lovett, A. Kapoor, Airsim: High-fidelity visual and physical simulation for autonomous vehicles, in: *Field and Service Robotics: Results of the 11th International Conference*, Springer, 2018, pp. 621–635.
- [20] Epic Games, Unreal engine.  
URL <https://www.unrealengine.com>
- [21] B. O. Community, Blender - a 3D modelling and rendering package, Blender Foundation, Stichting Blender Foundation, Amsterdam (2018).  
URL <http://www.blender.org>
- [22] G. Hardouin, J. Moras, F. Morbidi, J. Marzat, E. M. Mouaddib, Next-best-view planning for surface reconstruction of large-scale 3d environments with multiple uavs, in: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2020, pp. 1567–1574.
- [23] N. Koenig, A. Howard, Design and use paradigms for gazebo, an open-source multi-robot simulator, in: 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566), Vol. 3, 2004, pp. 2149–2154 vol.3.
- [24] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, W. Zaremba, Openai gym, arXiv preprint arXiv:1606.01540 (2016).
- [25] M. Everett, Y. F. Chen, J. P. How, Motion planning among dynamic, decision-making agents with deep reinforcement learning, in: 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), IEEE, 2018, pp. 3052–3059.
- [26] M. Lodell, B. Brito, A. Serra-Gómez, L. Ferranti, R. Babuška, J. Alonso-Mora, Where to look next: Learning viewpoint recommendations for informative trajectory planning, in: 2022 International Conference on Robotics and Automation (ICRA), IEEE, 2022, pp. 4466–4472.
- [27] M. Tucso, S. H. Gazani, K. Gupta, H. Najjaran, 3d reconstruction from 2d images: A two-part autoencoder-like tool, in: 2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC), 2022, pp. 538–543.
- [28] D. G. Lowe, Distinctive image features from scale-invariant keypoints, *International Journal of Computer Vision* 60 (2004) 91–110.
- [29] D. Nister, H. Stewenius, Scalable recognition with a vocabulary tree, in: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06), Vol. 2, 2006, pp. 2161–2168.
- [30] S. Lazebnik, C. Schmid, J. Ponce, Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories, in: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06), Vol. 2, 2006, pp. 2169–2178.
- [31] K. Arai, A. Barakbah, Hierarchical k-means: An algorithm for centroids initialization for k-means, *Reports of the Faculty of Science and Engineering* 36 (2007) 25–31.
- [32] G. Van Rossum, F. L. Drake, *Python 3 Reference Manual*, CreateSpace, Scotts Valley, CA, 2009.
- [33] C. Heindl, L. Brunner, S. Zambal, J. Scharinger, Blendtorch: A real-time, adaptive domain randomization library, in: *Pattern Recognition. ICPR International Workshops and Challenges: Virtual Event, January 10–15, 2021, Proceedings, Part IV*, Springer, 2021, pp. 538–551.
- [34] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, S. Chintala, Pytorch: An imperative style, high-performance deep learning library, in: *Advances in Neural Information Processing Systems* 32, Curran Associates, Inc., 2019, pp. 8024–8035.
- [35] P. R. Wolf, B. A. Dewitt, B. E. Wilkinson, *Elements of Photogrammetry with Applications in GIS*, McGraw-Hill Education, 2014.
- [36] A. Zeng, S. Song, M. Nießner, M. Fisher, J. Xiao, T. Funkhouser, 3dmatch: Learning local geometric descriptors from rgb-d reconstructions, in: *CVPR*, 2017.
- [37] W. E. Lorensen, H. E. Cline, Marching cubes: A high resolution 3d surface construction algorithm, *ACM siggraph computer graphics* 21 (4) (1987) 163–169.
- [38] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel, et al., Soft actor-critic algorithms and applications, arXiv preprint arXiv:1812.05905 (2018).
- [39] Cloudcompare.  
URL <https://www.cloudcompare.org>