

MAP-NBV: Multi-agent Prediction-guided Next-Best-View Planning for Active 3D Object Reconstruction

Harnaik Dhami*

Vishnu Dutt Sharma*

Pratap Tokekar

Abstract—Next-Best View (NBV) planning is a long-standing problem of determining where to obtain the next best view of an object from, by a robot that is viewing the object. There are a number of methods for choosing NBV based on the observed part of the object. In this paper, we investigate how predicting the unobserved part helps with the efficiency of reconstructing the object. We present, Multi-Agent Prediction-Guided NBV (*MAP-NBV*), a decentralized coordination algorithm for active 3D reconstruction with multi-agent systems. Prediction-based approaches have shown great improvement in active perception tasks by learning the cues about structures in the environment from data. However, these methods primarily focus on single-agent systems. We design a decentralized next-best-view approach that utilizes geometric measures over the predictions and jointly optimizes the information gain and control effort for efficient collaborative 3D reconstruction of the object. Our method achieves 19% improvement over the non-predictive multi-agent approach in simulations using AirSim and ShapeNet. We make our code publicly available through our project website: <http://raaslab.org/projects/MAPNBV/>.

I. INTRODUCTION

Visual surveying and inspection with robots has been studied for a long time for a wide range of applications such as inspection of civil infrastructure [1], [2] and large vehicles [3], [4], precision agriculture [5], and digital mapping for real estate [6], [7]. Using robots in these applications is highly advantageous as they can access hard-to-reach areas with greater ease and safety compared to humans.

This work focuses on one such long-studied problem of 3D object reconstruction [8], where the objective is to digitally reconstruct the object of interest by combining observations from multiple vantage points. While it could be easier to achieve this in an indoor environment by carefully placing sensors around the object, the same can not be achieved in the outdoors and open areas. For the latter, the sensor(s), must be moved around the object to capture information from different viewpoints. This can be realized with sensors such as cameras and LiDARs mounted on Unmanned Aerial Vehicles (UAVs).


A typical goal is to be selective in choosing observation locations for the UAV to achieve an accurate 3D reconstruction as fast as possible. The trade-off between reconstruction accuracy and task duration in unknown environments is commonly addressed through Next-Best-View (NBV) planning [9], wherein a robot determines the optimal location for the next observation to maximize information gain.

*Equal contribution. Names are listed alphabetically.


Authors are with the Department of Computer Science, University of Maryland, U.S.A. {dhami, vishnuds, tokekar}@umd.edu.

This work is supported by the ONR (grant number N00014-18-1-2829).

Employing a team of UAVs for this problem is an intuitive solution as multiple UAVs can simultaneously cover multiple viewpoints. NBV planning can be used to move the UAVs, gather new observations, and fill the gaps in the object representation. An efficient reconstruction, however, requires not only a correct estimation of the missing information but also coordination among the UAVs to minimize redundancies in the observations.




(a) C-17 and the robots in AirSim simulation




(b) Initial observations by robot 1 and robot 2




(c) Poses selected by frontier baseline based on observations



(d) Poses selected by MAP-NBV based on observations



(e) Observations after the first MAP-NBV iteration



(f) Full point cloud observed by MAP-NBV after termination

Fig. 1: MAP-NBV uses **predictions** to select better NBVs for a team of robots compared to the non-predictive baseline approach.

As shown in our prior work, Pred-NBV [10], estimating the unseen parts of the objects with point cloud completion networks can improve NBV planning, and hence the reconstruction efficiency, for a single UAV. These findings lead us to ask: **Can prediction improve the efficiency of multi-agent object reconstruction, given that multiple robots can themselves provide good coverage of the object?** Assuming predictions can augment the perception of multi-agent systems as well, a naive extension of methods designed for a single agent may result in significant overlaps in the observations by the team, necessitating coordination among all the robots. Prior works have shown that for target coverage problems (which object reconstruction is)

explicit coordination plays an important role in developing an efficient solution [11]. However, this prior work only focused on scenarios where the coordination used past observation and not predictions. This begs the question: **How does coordination, in perception and planning, affect multi-agent object reconstruction when each robot has access to predictions?**

To answer these questions, we make the following contributions in this work:

- 1) We propose a decentralized, multi-agent, prediction-based NBV planning approach, named *MAP-NBV*, for active 3D reconstruction of various objects with a novel objective combining visual information gain and control effort. *MAP-NBV* uses partial point clouds and predicts what the rest of the point cloud would be (Figure 1) and exploits the submodular nature of the objective to coordinate in a decentralized fashion.
- 2) We show that predictions effectively improve the performance by **19.41%** over non-predictive baselines that use frontier-based NBV planning [12] in AirSim [13] simulations.
- 3) We also show that *MAP-NBV* results in at least **17.12%** better reconstruction than non-cooperative prediction-guided method with experiments using the ShapeNet [14] dataset and performs comparable to a centralized approach.

We share the qualitative results and release the project code from our method on our project website.¹

II. RELATED WORK

The use of robots for data acquisition purposes is an extensively studied topic for various domains. Their usage range from infrastructure inspection [15] and environment monitoring [16], [17] for real-world application to the real-world digitization for research datasets and simulations [6], [7], [18]. When the environment is unknown, active methods such as NBV [9] are used to construct an object model on the fly by capturing additional observations. A majority of the works on NBV planning use information-theoretic measures [19] for selection to account for uncertainty in observations [19]–[21]. The widely used frontier and tree-based exploration approaches also utilize uncertainty about the environment for guiding the robot motion [22]–[25]. Some works devise geometric methods that make inferences about the exact shape of the object of interest and try to align the observations with the inferred model [26]–[28]. Prediction-based NBV approaches have emerged as another alternative in recent years, where a neural network takes the robot and/or the environment state as the input and NBV pose or velocity as the output [10], [29]–[31].

Directly extending single-robot NBV approaches to multi-robot systems may result in sub-optimal performance due to significant overlap in observations. This issue led to the

development of exploration algorithms specifically for multi-robot systems [32]–[34] with information-theoretic measures for determining NBV.

Some recent works on multi-robot systems have explored the use of predictions for improvement in task efficiency. Almadhoun et al. [35] designed a hybrid planner that switches between a classical NBV approach and a learning-based predictor for NBV selection but uses the partial model obtained by robot observations only. Wu et al. [36] use a point cloud prediction model for plants to use the predicted point cloud as an oracle leading to better results than the traditional approaches. This method uses entropy-based information gain measures for NBV and is designed for plant phenotyping with robotic arms. These methods do not consider the control effort required which is important for UAVs with energy constraints when deployed for observing large objects such as airplanes and ships. Also, these works employ information theoretic NBV approaches. We aim to explore a prediction-based approach for geometric NBV selection.

In this work, we used point cloud predictions similar to Pred-NBV [10] and built a decentralized multi-robot NBV planner. The prediction on the point cloud makes the pipeline modular and interpretable, allowing for improvements by enhancing individual modules. We select NBV based on information gain, as well as control effort, making our approach more grounded in the real world.

III. PROBLEM FORMULATION

We are given a team of n robots, each equipped with an RGB-D camera or LiDAR sensor. This team is tasked with navigating around a closed object of volume $\mathcal{V} \in \mathbb{R}^3$ and observes its surface as a set of 3D points, \mathcal{S} . At any given time t , the set of surface points S_t^i observed by the robot r_i from the view-point $\phi_t^i \in \Phi$ at time t is represented as a voxel-filtered point cloud and the relationship between them is defined as $S_t^i = f(r_i, \phi_t^i)$. Each robot r_i follows a trajectory ξ_i , which consists of a sequence of viewpoints aimed at maximizing the coverage of \mathcal{S} while minimizing redundancy among observed points.

The distance traveled by a robot between two poses ϕ_i and ϕ_j is represented by $d(\phi_i, \phi_j)$. The point cloud observed by the team of robots is the union of the surface points observed by the individual robots over their respective trajectories, i.e., $S_{\bar{\xi}} = \bigcup_{i=1}^n \bigcup_{\phi \in \xi_i} f(r_i, \phi)$ and $\bar{\xi}$ represents the set of trajectories for each robot, i.e., $\bar{\xi} = \{\xi_1, \xi_2, \dots, \xi_n\}$.

The objective is to find a set of feasible trajectories $\bar{\xi}^* = \{\xi_1^*, \xi_2^*, \dots, \xi_n^*\}$, such that the team observes the whole voxel-filtered surface \mathcal{S} , while also minimizing the total distance traveled by the robots on their respective trajectories.

$$\bar{\xi}^* = \arg \min_{\bar{\xi}} \sum_{i=1}^n \sum_{j=1}^{|\xi_i|-1} d(\phi_j^i, \phi_{j+1}^i) \quad (1)$$

$$\text{such that } \bigcup_{i=1}^n \bigcup_{\phi \in \xi_i} f(r_i, \phi) = \mathcal{S} \quad (2)$$

Given a finite set of trajectories, if the object model, \mathcal{S} , is known, we can find the optimal set of trajectories through an

¹<http://raaslab.org/projects/MAPNBV/>

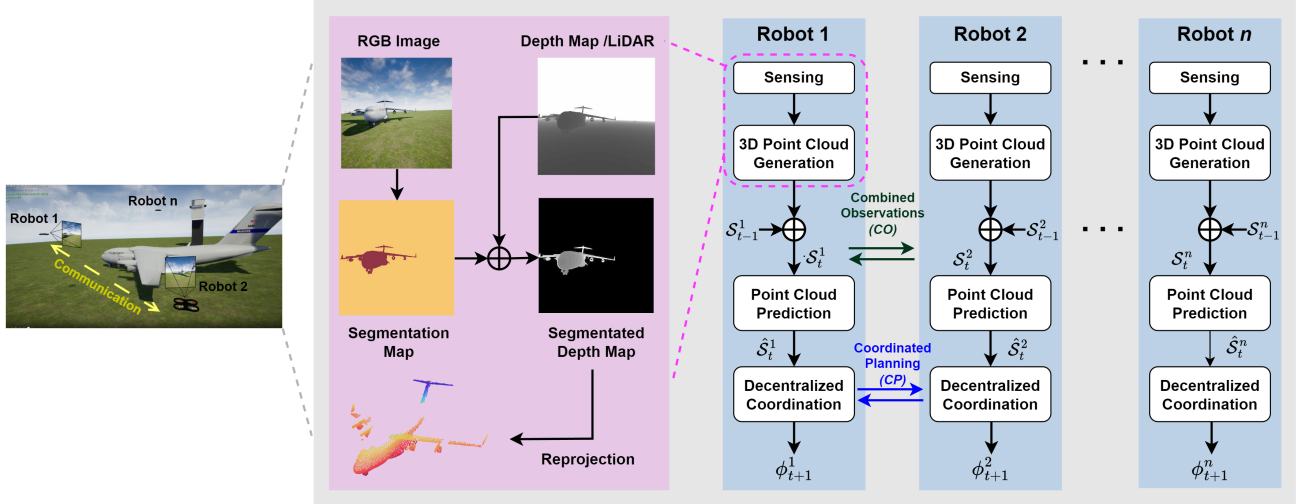


Fig. 2: Algorithm Overview: Each robot runs the same algorithm including perception, prediction, and planning steps. The robots that communicate with each other can share observations and coordinate planning, whereas robots in isolation (e.g., Robot n) perform individual greedy planning.

exhaustive search. As the object model is not known a priori in an unknown environment, the optimal solution can not be found beforehand. Thus, each robot needs to determine the NBV based on the partial observations of the team to reconstruct the object’s surface. Here we assume that each robot can observe the object at the start of the mission, which can be accomplished by moving the robots till they see the object. While a centralized server can help find an optimal assignment solution, a limited communication range can make a centralized solution infeasible. Thus, we define this problem as a decentralized one; each robot solves this objective but the communicating robots can collaborate and coordinate with their neighbors.

IV. PROPOSED APPROACH

In this section, we present *Multi-Agent Pred-NBV (MAP-NBV)*, a prediction-guided NBV approach for a team of robots. Figure 2 shows the overview of our process, which consists of two parts: (1) *3D Model Prediction*, where we combine the observations from the neighboring robots to build a partial model of the object and use PoinTr-C [10], a 3D point cloud completion network, to predict the full shape of the objects, and (2) *Decentralized Coordination* which combines the observations from the communicating robots and solves an NBV objective to maximize information gain and minimize the control effort with sequential greedy assignment. Robots that do not communicate with anyone effectively run an individual greedy algorithm. This approach is detailed in Algorithm 1. Apart from being feasible and scalable, this approach also reduces the computation complexity resulting in fast runtime.

A. 3D Model Prediction (Line 5-6)

To start, we use the RGB images to segment out the object and the depth sensors to generate the point cloud from the current observations for each robot, giving us segmented point clouds. This allows the algorithm to focus on only

Algorithm 1 MAP-NBV Algorithm (for robot r_i)

- 1: **Inputs:** Initial positions ϕ_0^i ; Stopping threshold τ ; Information gain threshold λ
- 2: **Output:** 3D Point Cloud \mathcal{S}
- 3: **Initialization:** $novelty \leftarrow 0$; $t \leftarrow 0$; $\mathcal{S}_{-1}^i \leftarrow \emptyset$, $\xi_i \leftarrow \{\phi_0^i\}$
- 4: **while** $novelty \leq \tau$ **do**
- 5: $\mathcal{S}_t^i \leftarrow \bigcup_{k \in Neighbors} f(r_i, \phi_t^i) \cup \mathcal{S}_{t-1}^i$
- 6: $\hat{\mathcal{S}}_t^i \leftarrow \text{getPredictionFromPoinTr-C}(\mathcal{S}_t^i)$
- 7: $\mathcal{C} \leftarrow \text{generateCandidatePoses}(\hat{\mathcal{S}}_t^i, \mathcal{S}_t^i)$
- 8: $\mathcal{I}_{seen} \leftarrow \bigcup_{k \in Neighbors; k < i} \mathcal{I}(\xi_k)$
- 9: $\hat{\mathcal{I}} \leftarrow \{\mathcal{I}(\xi_i \cup \phi) - \mathcal{I}_{seen}; \forall \phi \in \mathcal{C}\}$
- 10: $\phi_{t+1}^i \leftarrow \arg \min_{\phi \in \mathcal{C}} d(\phi, \phi_{t+1}^i), \text{ s.t. } \frac{\hat{\mathcal{I}}(\phi)}{\max \hat{\mathcal{I}}} \geq \lambda$
- 11: $\xi_i \leftarrow \xi_i \cup \phi_{t+1}^i$
- 12: Broadcast(ξ_i)
- 13: $novelty \leftarrow \frac{|\mathcal{S}_t^i|}{|\mathcal{S}_{t-1}^i|}$
- 14: $t \leftarrow t + 1$
- 15: **end while**
- 16: $\mathcal{S} \leftarrow \bigcup_{i=1}^n \mathcal{S}_{t-1}^i$

the target infrastructure as opposed to also including other obstacles. For the robots that can communicate with each other (e.g., Robot 1 and Robot 2 in Fig 2), each segmented point cloud per robot is transformed into a global reference frame and concatenated together into a single point cloud (Line 5). This point cloud represents the entire communication subgraph’s observations of the target object at the current timestamp. The point cloud concatenation can be replaced with a registration algorithm [37], but we use concatenation due to its ease of use. Lastly, this current timestamp’s point cloud is then concatenated with previous observations to get an up-to-date observation point cloud.

In order to get an approximation $\hat{\mathcal{S}}$ of the full model \mathcal{S} , we use PoinTr-C [10]. The combined observed point cloud of the object at time t , \mathcal{S}_t^i goes as input to PoinTr-C and it predicts the full object point cloud $\hat{\mathcal{S}}_t^i$ (Line 6). PoinTr-C

requires isolating the object point clouds from the scene. This can be realized with the help of distance-based filters and state-of-the-art segmentation networks [38] without any fine-tuning. An example of an input point cloud and a predicted point cloud, both from individual observations and combined observations, is shown in Figure 1.

B. Decentralized Coordination (Line 7-12)

Next Best View Planning. We use the predicted point cloud \mathcal{S}_t as an approximation of the ground truth point cloud for NBV planning. For this, we first generate a set of candidate poses around the partially observed object (Line 7). From these, we select a set of poses as NBVs for each robot, based on information gain and control effort. For a set of k candidate viewpoints, we define the information gain, $\hat{\mathcal{I}}$, as the expected number of new, unique points the robots will observe after moving to these viewpoints. The control effort is defined as the total distance traversed by the robots to reach the viewpoints.

The number of new points varies with each iteration as robots move to new locations, observing more of the object’s surface. While PoinTr-C predicts the point cloud for the whole object, the robots can observe only points on the surface close to it. Hence, before counting the number of new points, we apply hidden point removal [39] to the predicted point cloud. We represent this relationship between the number of points observed and the trajectories traversed till time for a robot r_i as $I(\xi_i)$, where $\xi = \{\phi_0^i, \phi_t^i, \dots, \phi_t^i\}$ represents the trajectory for the robot r_i till time t consisting of the set of viewpoints the robot has traversed. To balance the information gain and control effort, we use a hyperparameter λ . First, we find the estimated information gain for each candidate pose $\hat{\mathcal{I}}$ (Line 9) by treating the $\hat{\mathcal{S}}$ as the actual full object model. Each robot selects the candidate pose closest to the current pose where it achieves at least $\lambda\%$ of the maximum possible information gain over all the candidate poses (Line 10).

To find the control effort, we use RRT-Connect [40] to find the path between the robot r_i ’s current location ϕ_t^i to each candidate pose ϕ and use its length as $d(\phi, \phi_t^i)$. The candidate poses are generated similar to Pred-NBV [10], i.e., on circles at different heights around the center of the predicted object point cloud (Line 7). One circle is at the same height as the predicted object center with radius $1.5 \times d_{max}$, where d_{max} is the maximum distance of a point from the center of the predicted point cloud. The other two circles are located above and below this circle $0.25 \times z\text{-range}$ away, with a radius of $1.2 \times d_{max}$. The viewpoints are located at steps of 30° on each circle.

Sequential Greedy Assignment. To effectively and efficiently use multiple robots, we coordinate among them for NBV assignment. We devise a decentralized coordination strategy to accommodate the dynamic communication graph effectively. We also leverage the submodular nature of the problem and use a sequential greedy assignment which results in a near-optimal solution [41]. With each communications subgraph, the robots select their NBVs in




Fig. 3: Flight paths of the two robots during C-17 simulation.

the order determined by their IDs (lower to higher). Each robot first considers the trajectories of its neighbors with lower IDs to find the information gain attained by their movement \mathcal{I}_{seen} (Line 8). For the robot with the lowest ID, \mathcal{I}_{seen} is empty. Then we greedily choose the candidate position that would result in the information gain above a threshold λ while minimizing the distance traveled. This candidate position for the next iteration ϕ_{t+1}^i is the NBV for r_i . We add this pose to the robot’s trajectory (Line 11) and broadcast this information with the neighbors (Line 12) to minimize overlaps.

In our experiments, we consider multi-hop communication, thus each robot can coordinate with every robot on its communication subgraph. Some robots may not be within any other robot’s communication range (e.g., Robot n in Fig 2). For such robots, this strategy effectively turns into a greedy prediction-guided assignment similar to Pred-NBV [10]. At each iteration, we calculate *novelty*, i.e., the ratio of the number of points observed over subsequent iterations (Line 13). The robot stops if *novelty* is above a predefined threshold τ (we set $\tau = 0.95$ in our experiments). In the end, all the robots assemble at the same location and combine their observations as the object model \mathcal{S} (Line 16).

V. EXPERIMENTS AND EVALUATION

We design experiments to answer the two key research questions: (1) can point cloud prediction improve multi-agent object reconstruction? and (2) how does coordination between agents affect the reconstruction? To answer the first question, we compare *MAP-NBV* with a non-predictive frontier-based baseline. For the second question, we compare *MAP-NBV* with a centralized and a non-coordinated variation of *MAP-NBV*. We first describe the experiment setups used to answer these questions and then discuss the results obtained.

A. Setup

AirSim: We use AirSim [42] simulator as it allows us to load desired object models and get photo-realistic inputs while also supporting multiple robots. We use Robot Operating System (ROS) Melodic to run the simulations on Ubuntu 18.04. We spawn multiple UAVs close to each other looking towards the object. We equip each UAV with a depth camera and an RGB camera. Each UAV publishes a segmented image using AirSim’s built-in segmentation. This segmented image is used along with a depth map to remove the background and isolate the depth map for the object. We then convert these segmented depth images into 3D point clouds. For collision-free point-to-point planning, we use the MoveIt [43] package implementing the work done by Köse [44].

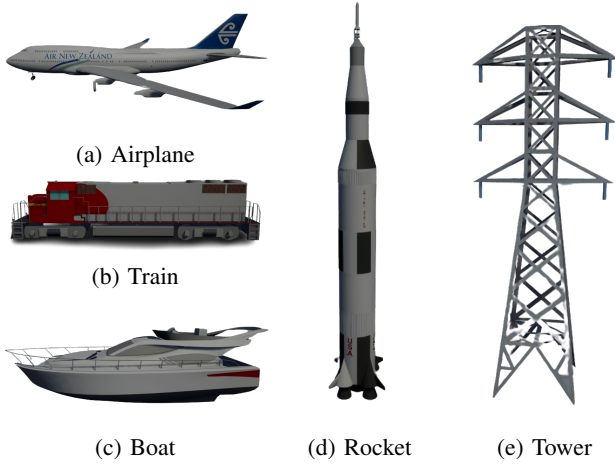


Fig. 4: Examples of the 5 simulation model classes.

ShapeNet: The AirSim setup helps us evaluate the coverage in a realistic setting, but it lacks a ground truth point cloud to compare the reconstruction quality. Hence, we run simulations with the ShapeNet [14] dataset to obtain a ground truth and compare the reconstruction quality across different methods. We simulate the UAVs as free-floating cameras which can teleport to any location. These cameras do not have a roll and pitch and can only change their yaw. We use the Open3D [45] library to load an object mesh (normalized) and select a random starting position around it. The UAVs are placed close to each other at this location. We render the depth image from the object mesh and reproject it to the 3D point cloud. Unlike the previous setup, we do not need segmentation as there is no background in this setup. To evaluate the reconstruction quality, we convert the mesh to a point cloud by uniformly sampling points on it and use it as the ground truth to quantify the reconstruction quality with Directional Chamfer Distance with ℓ_2 -norm ($CD-\ell_2$) (from ground truth point cloud to the observed point cloud). We use Euclidean distance as the distance metric $d(\cdot, \cdot)$ in this setup.

B. Qualitative Example

We evaluate *MAP-NBV* on the same 20 objects that were used in *Pred-NBV* to allow a direct comparison. The 20 objects consist of five different ShapeNet classes: airplane, rocket, tower, train, and watercraft. Examples of each class are shown in Figure 4. These classes represent diverse shapes and infrastructures that are regularly inspected. Figure 3 shows the path followed by two UAVs as given by *MAP-NBV* in the C-17 airplane simulation. This environment includes other obstacles that are not of interest but still need to be accounted for in collision-free path planning. *MAP-NBV* finds a collision-free path for both UAVs while targeting the maximum coverage of the C-17 airplane.

C. Can Point Cloud Prediction Improve Reconstruction?

We compared the performance of *MAP-NBV* with a modified baseline NBV method [12] designed for multi-agent use on 20 objects, as listed in Table I. The baseline method employs frontiers to select the next-best views. Frontiers are points located at the edge of the observed space near

unknown areas. We utilized the same modifications described in *Pred-NBV* [10]. Specifically, we used our segmented point cloud to choose frontiers near the target object. To ensure that the UAVs always face the target object, the orientation of all poses selected by the baseline aligns with the center of the observed target object point clouds.

We further adapted this baseline method to function in a multi-agent setting. The pose for the first UAV is selected in the exact same manner as in the single-agent baseline. For each subsequent UAV, the remaining best pose is chosen, as long as it does not fall within a certain distance threshold compared to the previously selected poses in the current iteration of the algorithm.

Both *MAP-NBV* and the baseline algorithm employ the same stopping criteria. The algorithm terminates if the total points observed in the previous step exceed 95% of the total points observed in the current step. We run both the algorithms in the AirSim setup. Our evaluation, presented in Table I, demonstrates that *MAP-NBV* observes, on average, 19.41% more points than the multi-agent baseline for object reconstruction across all 20 objects from the five different model classes. In our simulations, we utilized 2 UAVs for both algorithms.

Furthermore, the *MAP-NBV* algorithm can be readily extended to accommodate more than just 2 robots. By incorporating additional UAVs, the algorithm can effectively leverage the collaborative efforts of a larger multi-agent system to improve object reconstruction performance and exploration efficiency. However, in our current evaluation, we utilized 2 UAVs for both algorithms due to limited computational resources. The simulations were computationally intensive, and our computer experienced significant slowdowns with just 2 robots in the simulation. Despite this limitation, the promising results obtained with 2 UAVs suggest that scaling up the algorithm to include more robots has the potential to yield even more significant performance improvements.

Additionally, Figure 5 illustrates that *MAP-NBV* observes more points per step than the multi-agent baseline.

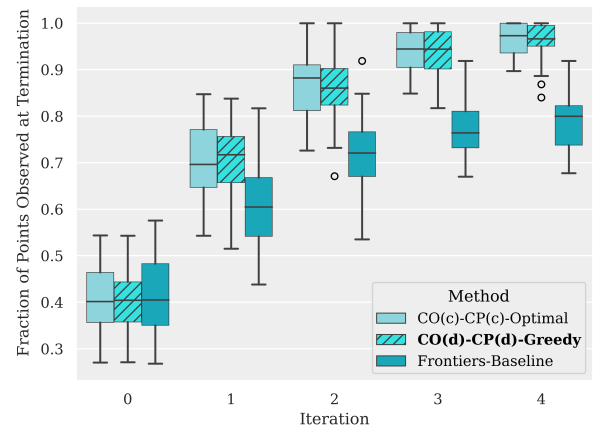


Fig. 5: *MAP-NBV* (CO (d) –CP (d) –Greedy) performs comparably to the optimal solution (CO (c) –CP (c) –Optimal; Section V-D), and much better than the frontiers-baseline in AirSim experiments.

TABLE I: *MAP-NBV* results in a better coverage compared to the multi-agent baseline NBV method [12] for all models in AirSim upon algorithm termination.

Class	Model	Points Seen		Improvement
		<i>MAP-NBV</i>	MA Baseline	
Airplane	747	16140	13305	19.26%
	A340	10210	8156	22.37%
	C-17	13278	10150	26.70%
	C-130	6573	5961	9.77%
	Fokker 100	14986	13158	12.99%
Rocket	Atlas	2085	1747	17.64%
	Maverick	3625	2693	29.50%
	Saturn V	1041	877	17.10%
	Sparrow	1893	1664	12.88%
	V2	1255	919	30.91%
Tower	Big Ben	4294	3493	20.57%
	Church	7884	6890	13.46%
	Clock	3163	2382	28.17%
	Pylon	2986	2870	3.96%
	Silo	5810	4296	29.96%
Train	Diesel	4013	3233	21.53%
	Mountain	5067	4215	18.36%
Watercraft	Cruise	5021	3685	30.69%
	Patrol	4078	3683	10.18%
	Yacht	11678	10341	12.14%

D. How does coordination affect reconstruction?

For studying the effect of coordination we randomly select 25 objects, 5 from each of the ShapeNet classes mentioned in Section V-C for the ShapeNet setup (Section V-A).

In a multi-agent setting, the communicating robots can combine their observation (CO) or they may choose to rely on their observations (IO) for point cloud predictions. After the prediction, they may choose to collaborate for planning (CP) or make individual decisions without relying on others (IP). To highlight the effect of these coordination strategies on reconstruction quality, for a team of n robots and k candidate viewpoints, we compare MAP-NBV with two contrasting coordination approaches:

- **CO(c)–CP(c)–Optimal**: This approach relies on combined observation and coordinated planning in a centralized manner. Here, a central server aggregates the observations, performs prediction, and finds the NBVs for each robot. For selecting NBVs, the algorithm evaluates all robot-candidate pose assignments and chooses the optimal setting, i.e., which results in maximum joint information gain. Among all possible permutations of robots that result in the optimal setting, we select the one that minimizes the maximum displacement for any robot. This process has a runtime complexity of $\mathcal{O}(k^n)$.
- **IO–IP**: In this approach, each robot operates individually and does not share observations or coordinate with others regardless of the communication range. This is effectively a naive extension of Pred-NBV to a multi-agent setup. The runtime complexity here is $\mathcal{O}(k)$.

Following the notations above, **MAP-NBV** is **CO(d)–CP(d)–Greedy**, i.e., combined observations (decentralized) and coordinated planning (decentralized) with greedy assignment. *Decentralized* indicates that the

observations and planning are shared among only those robots that form a communication subgraph. *MAP-NBV* thus has a runtime complexity of $\mathcal{O}(n \cdot k)$ in the worst case only, i.e., when all the robots are connected. We study these algorithms for teams of 2, 4 and 6 robots.

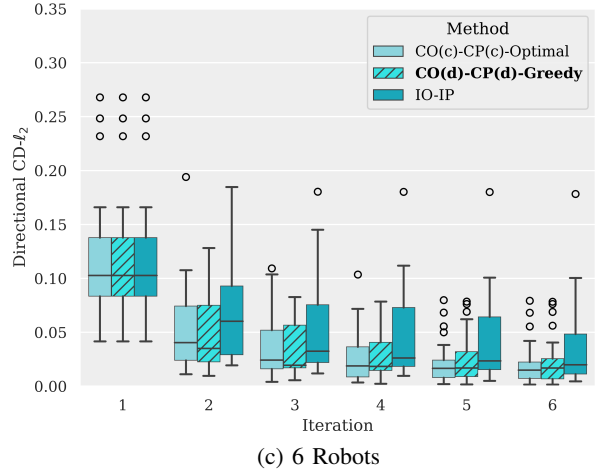

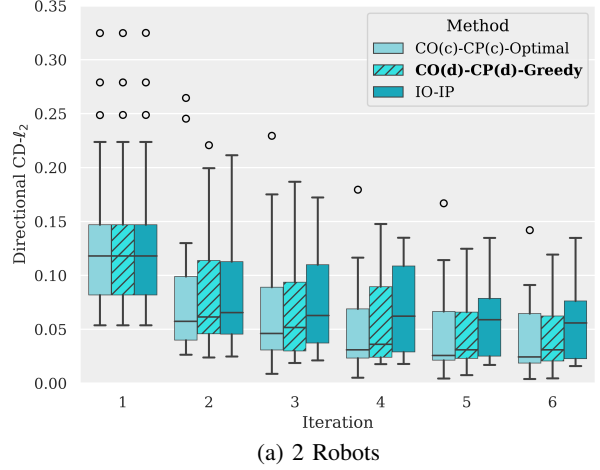


Fig. 6: Directional $CD\text{-}l_2$ for teams of 2, 4, and 6 robots on ShapeNet models [14] with different coordination strategies.

At each iteration, we combine the observations from all robots and compute Directional $CD\text{-}l_2$ over it. This thus represents the reconstruction quality if the mission terminated

at that iteration. We found that 6 iterations were enough for all robots to reach the stopping criteria in each setting. Our findings are shown in Figure 6.

We observed that coordination and information sharing play a crucial role in improving the reconstruction quality. $CO(c)-CP(c)-Optimal$ exhibits the best reconstruction over time, as expected, owing to shared observations which lead to better estimation of the partial point cloud, and coordinated planning, which minimizes overlap in information gain. In early iterations, the limited observations may lead to an imprecise prediction, which may result in inefficient NBV assignments. Still, over time the observation coverage increases leading to better predictions and performance over the other algorithms. We observed similar trends for the AirSim setup as well as shown in Figure 5.

However, $CO(c)-CP(c)-Optimal$ requires large computation time and does not scale well in comparison with $MAP-NBV$ ($CO(d)-CP(d)-Greedy$). On a Ubuntu 20.04 system with 32-core, 2.10Ghz Xeon Silver-4208 CPU and Nvidia GeForce RTX 2080Ti GPU, $MAP-NBV$ was 2x faster for 4 robots and 10x faster for 6 robots compared to $CO(c)-CP(c)-Optimal$. Even with faster execution, $MAP-NBV$ performs comparably to $CO(c)-CP(c)-Optimal$, making it more attractive than the former. Additionally, the decentralized nature makes $MAP-NBV$ a more feasible algorithm than $CO(c)-CP(c)-Optimal$ which requires centralization. $IO-IP$, where the robots do not share observations or coordinate plans, exhibits the worst improvement over time, highlighting that **the coordination plays a crucial role in multi-agent object reconstruction**. In fact, $MAP-NBV$ achieves a relative improvement of **17-22%** in directional $CD-\ell_2$ over $IO-IP$ after termination.

Interestingly, we observe that as the number of robots increases, the gap between these algorithms decreases. This is expected as more robots lead to more coverage at any time. Since the robots start the mission at close distances in our experiments, we found that $IO-IP$ still exhibits worse improvement as the robots suffer significant overlaps in observations. Thus coordination still provides benefits and faster runtime of $MAP-NBV$ makes it a more suitable choice than the centralized alternative, even if centralization is feasible. We also performed experiments with a centralized, greedy selection variant of $MAP-NBV$ and found it performs only marginally better than $MAP-NBV$'s performance. We share these findings on our [project webpage](#) due to lack of space.

E. Qualitative Real-World Experiment

We also conducted a real-world experiment to evaluate the feasibility of implementing $MAP-NBV$ on hardware. A single iteration of the $MAP-NBV$ pipeline for 2 robots was run using a ZED camera. Depth images and point clouds were captured from the ZED of a Toyota RAV4 car. The car was extracted from the point clouds and used as input for the $MAP-NBV$ pipeline. Then, depth images and point clouds were captured from the poses $MAP-NBV$ outputted. Due to the lower quality of the ZED camera, an iPhone 12

Pro Max with built-in LiDAR was used to capture data to create a high-resolution reconstruction shown in Figure 7d.

VI. CONCLUSIONS

We present a multi-agent, decentralized, prediction-guided NBV planning approach for active 3D reconstruction. This method can be helpful in a variety of applications including civil infrastructure inspection. We show that our method can faithfully reconstruct the object point clouds efficiently compared to non-predictive multi-agent methods and other prediction-guided approaches. Our NBV planning objective considers both information gain and control effort, making it more suitable for real-world deployment given the flight time limit imposed on UAVs by their battery capacity.

We are currently working on a bandwidth-aware extension of $MAP-NBV$ and the preliminary studies show encouraging results (shared on our [webpage](#)). In this work, we focus solely on geometric measures for information gain. Many existing works on NBV have developed sophisticated information theoretic measures. We will explore combining both types of measures in our future work.

REFERENCES

- [1] P. Shanthakumar, K. Yu, M. Singh, J. Oreillo, E. Bianchi, M. Hebdon, and P. Tokekar, "View planning and navigation algorithms for autonomous bridge inspection with uavs," in *Proceedings of the 2018 International Symposium on Experimental Robotics*. Springer, 2020, pp. 201–210.
- [2] H. Dhami, K. Yu, T. Williams, V. Vajipey, and P. Tokekar, "GATSBI: An online GTSP-based algorithm for targeted surface bridge inspection," in *2023 International Conference on Unmanned Aircraft Systems (ICUAS)*. IEEE, jun 2023. [Online]. Available: <https://doi.org/10.1109%2Ficuas57906.2023.10156013>
- [3] A. Kim and R. Eustice, "Pose-graph visual slam with geometric model selection for autonomous underwater ship hull inspection," in *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2009, pp. 1559–1565.
- [4] L. Ropek, "This startup is using drones to conduct aircraft inspections," Mar 2021. [Online]. Available: <https://gizmodo.com/this-startup-is-using-drones-to-conduct-airplane-inspec-1846527002>
- [5] H. Dhami, K. Yu, T. Xu, Q. Zhu, K. Dhakal, J. Friel, S. Li, and P. Tokekar, "Crop height and plot estimation for phenotyping from unmanned aerial vehicles using 3d lidar," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [6] T. Zhou, R. Tucker, J. Flynn, G. Fyffe, and N. Snavely, "Stereo magnification: Learning view synthesis using multiplane images," *ACM Trans. Graph. (Proc. SIGGRAPH)*, vol. 37, 2018. [Online]. Available: <https://arxiv.org/abs/1805.09817>
- [7] S. K. Ramakrishnan, A. Gokaslan, E. Wijmans, O. Maksymets, A. Clegg, J. M. Turner, E. Undersander, W. Galuba, A. Westbury, A. X. Chang, M. Savva, Y. Zhao, and D. Batra, "Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for embodied AI," in *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*, 2021. [Online]. Available: <https://openreview.net/forum?id=v4OuQNs5P>
- [8] R. Bajcsy, Y. Aloimonos, and J. K. Tsotsos, "Revisiting active perception," *Autonomous Robots*, vol. 42, pp. 177–196, 2018.
- [9] C. Connolly, "The determination of next best views," in *Proceedings. 1985 IEEE international conference on robotics and automation*, vol. 2. IEEE, 1985, pp. 432–435.
- [10] H. Dhami, V. D. Sharma, and P. Tokekar, "Pred-nbv: Prediction-guided next-best-view planning for 3d object reconstruction," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 7149–7154.
- [11] M. Corah, C. O'Meadhra, K. Goel, and N. Michael, "Communication-efficient planning and mapping for multi-robot exploration in large environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1715–1721, 2019.




Fig. 7: Real-World MAP-NBV experiment. (a) RGB Image. (b) Observations, Predictions, and MAP-NBV poses. (c) Initial, Drone 1, and Drone 2 points after MAP-NBV iteration. (d) Reconstruction.

- [12] J. Aleotti, D. L. Rizzini, R. Monica, and S. Caselli, "Global registration of mid-range 3d observations and short range next best views," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 3668–3675.
- [13] S. Shah, D. Dey, C. Lovett, and 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.
- [14] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "ShapeNet: An Information-Rich 3D Model Repository," Stanford University — Princeton University — Toyota Technological Institute at Chicago, Tech. Rep. arXiv:1512.03012 [cs.GR], 2015.
- [15] T. Ozaslan, S. Shen, Y. Mulgaonkar, N. Michael, and V. Kumar, "Inspection of penstocks and featureless tunnel-like environments using micro uavs," in *International Conference on Field and Service Robotics*, 2013.
- [16] M. Dunbabin and L. Marques, "Robots for environmental monitoring: Significant advancements and applications," *IEEE Robotics and Automation Magazine*, vol. 19, no. 1, pp. 24–39, Mar 2012.
- [17] Y. Sung and P. Tokekar, "A competitive algorithm for online multi-robot exploration of a translating plume," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, may 2019.
- [18] P. Ammirato, P. Poirson, E. Park, J. Košecá, and A. C. Berg, "A dataset for developing and benchmarking active vision," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1378–1385.
- [19] J. Delmerico, S. Isler, R. Sabzevari, and D. Scaramuzza, "A comparison of volumetric information gain metrics for active 3d object reconstruction," *Autonomous Robots*, vol. 42, no. 2, pp. 197–208, 2018.
- [20] B. Kuipers and Y.-T. Byun, "A robot exploration and mapping strategy based on a semantic hierarchy of spatial representations," *Robotics and autonomous systems*, vol. 8, no. 1-2, pp. 47–63, 1991.
- [21] J. I. Vasquez-Gomez, L. E. Sucar, R. Murrieta-Cid, and E. Lopez-Damian, "Volumetric next-best-view planning for 3d object reconstruction with positioning error," *International Journal of Advanced Robotic Systems*, vol. 11, no. 10, p. 159, 2014.
- [22] 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*. IEEE, 1997, pp. 146–151.
- [23] H. H. González-Banos and J.-C. Latombe, "Navigation strategies for exploring indoor environments," *The International Journal of Robotics Research*, vol. 21, no. 10-11, pp. 829–848, 2002.
- [24] B. Adler, J. Xiao, and J. Zhang, "Autonomous exploration of urban environments using unmanned aerial vehicles," *Journal of Field Robotics*, vol. 31, no. 6, pp. 912–939, 2014.
- [25] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon path planning for 3d exploration and surface inspection," *Autonomous Robots*, vol. 42, no. 2, pp. 291–306, 2018.
- [26] K. A. Tarabanis, P. K. Allen, and R. Y. Tsai, "A survey of sensor planning in computer vision," *IEEE transactions on Robotics and Automation*, vol. 11, no. 1, pp. 86–104, 1995.
- [27] J. E. Banta, L. Wong, C. Dumont, and M. A. Abidi, "A next-best-view system for autonomous 3-d object reconstruction," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 30, no. 5, pp. 589–598, 2000.
- [28] S. Kriegel, M. Brucker, Z.-C. Marton, T. Bodenmüller, and M. Suppa, "Combining object modeling and recognition for active scene exploration," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 2384–2391.
- [29] E. Johns, S. Leutenegger, and A. J. Davison, "Pairwise decomposition of image sequences for active multi-view recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 3813–3822.
- [30] M. Mendoza, J. I. Vasquez-Gomez, H. Taud, L. E. Sucar, and C. Reta, "Supervised learning of the next-best-view for 3d object reconstruction," *Pattern Recognition Letters*, vol. 133, pp. 224–231, 2020.
- [31] R. Zeng, W. Zhao, and Y.-J. Liu, "Pc-nbv: A point cloud based deep network for efficient next best view planning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 7050–7057.
- [32] W. Burgard, M. Moors, C. Stachniss, and F. E. Schneider, "Coordinated multi-robot exploration," *IEEE Transactions on robotics*, vol. 21, no. 3, pp. 376–386, 2005.
- [33] A. A. Amanatiadis, S. A. Chatzichristofis, K. Charalampous, L. Doitsidis, E. B. Kosmatopoulos, P. Tsalides, A. Gasteratos, and S. I. Roumeliotis, "A multi-objective exploration strategy for mobile robots under operational constraints," *IEEE Access*, vol. 1, pp. 691–702, 2013.
- [34] G. Hardouin, J. Moras, F. Morbidi, J. Marzat, and 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.
- [35] R. Almadhoun, T. Taha, L. Seneviratne, and Y. Zweiri, "Multi-robot hybrid coverage path planning for 3d reconstruction of large structures," *IEEE Access*, vol. 10, pp. 2037–2050, 2021.
- [36] C. Wu, R. Zeng, J. Pan, C. C. Wang, and Y.-J. Liu, "Plant phenotyping by deep-learning-based planner for multi-robots," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3113–3120, 2019.
- [37] X. Huang, G. Mei, J. Zhang, and R. Abbas, "A comprehensive survey on point cloud registration," 2021.
- [38] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, et al., "Segment anything," *arXiv preprint arXiv:2304.02643*, 2023.
- [39] S. Katz, A. Tal, and R. Basri, "Direct visibility of point sets," in *ACM SIGGRAPH 2007 papers*, 2007, pp. 24–es.
- [40] J. J. Kuffner and S. M. LaValle, "Rrt-connect: An efficient approach to single-query path planning," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 2. IEEE, 2000, pp. 995–1001.
- [41] G. Calinescu, C. Chekuri, M. Pal, and J. Vondrák, "Maximizing a monotone submodular function subject to a matroid constraint," *SIAM Journal on Computing*, vol. 40, no. 6, pp. 1740–1766, 2011.
- [42] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "Airsim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and Service Robotics*, 2017. [Online]. Available: <https://arxiv.org/abs/1705.05065>
- [43] D. Coleman, I. Sucan, S. Chitta, and N. Correll, "Reducing the barrier to entry of complex robotic software: a moveit! case study," *arXiv preprint arXiv:1404.3785*, 2014.
- [44] "tahsinkose/hector-moveit: Hector quadrotor with moveit! motion planning framework," <https://github.com/tahsinkose/hector-moveit>, (Accessed on 03/03/2021).
- [45] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3D: A modern library for 3D data processing," *arXiv:1801.09847*, 2018.