

VOOM: Robust Visual Object Odometry and Mapping using Hierarchical Landmarks

Yutong Wang¹, Chaoyang Jiang¹, Xieyuanli Chen²

Abstract—In recent years, object-oriented simultaneous localization and mapping (SLAM) has attracted increasing attention due to its ability to provide high-level semantic information while maintaining computational efficiency. Some researchers have attempted to enhance localization accuracy by integrating the modeled object residuals into bundle adjustment. However, few have demonstrated better results than feature-based visual SLAM systems, as the generic coarse object models, such as cuboids or ellipsoids, are less accurate than feature points. In this paper, we propose a Visual Object Odometry and Mapping framework VOOM using high-level objects and low-level points as the hierarchical landmarks in a coarse-to-fine manner instead of directly using object residuals in bundle adjustment. Firstly, we introduce an improved observation model and a novel data association method for dual quadrics, employed to represent physical objects. It facilitates the creation of a 3D map that closely reflects reality. Next, we use object information to enhance the data association of feature points and consequently update the map. In the visual object odometry backend, the updated map is employed to further optimize the camera pose and the objects. Meanwhile, local bundle adjustment is performed utilizing the objects and points-based covisibility graphs in our visual object mapping process. Experiments show that VOOM outperforms both object-oriented SLAM and feature points SLAM systems such as ORB-SLAM2 in terms of localization. The implementation of our method is available at <https://github.com/yutongwangBIT/VOOM.git>.

I. INTRODUCTION

Accurate localization and mapping are essential in various applications, including autonomous navigation, robot manipulation, and augmented reality. Cameras are cost-effective and easy to integrate while offering rich geometry and semantic information, making them practical for real-world scenarios. Consequently, visual odometry or simultaneous localization and mapping (SLAM) has attracted increasing attention from researchers.

However, traditional visual odometry or SLAM frameworks that use map points as landmarks lack high-level information, resulting in limited robustness and applicability [1]–[3], which has led to a growing interest in object-level odometry and mapping algorithms [4]–[6], especially in

*This work was supported by the National Key Research and Development Project of China (No. 2020YFC1512503), Fund for key Laboratory of Space Flight Dynamics Technology (Num 2022-JYAPAF-F1028), and Young Elite Scientists Sponsorship Program by CAST (No. 2023QNRC001). (Corresponding authors: Chaoyang Jiang and Xieyuanli Chen.)

¹Y. Wang and C. Jiang are with the School of Mechanical Engineering, Beijing Institute of Technology, China, and the Yangtze Delta Region Academy of Beijing Institute of Technology, Jiaxing, China. yutongwang@bit.edu.cn, cjiang@bit.edu.cn

²X. Chen is with the College of Intelligence Science and Technology, National University of Defense Technology, China xieyuanli.chen@nudt.edu.cn




Fig. 1: Top: Groundtruth map. Middle: ORB SLAM2 with loop closure disabled. Bottom: VOOM. The colored ellipsoids represent the 3D objects, blue cones illustrate the camera poses, and green lines represent covisibility relationships between keyframes. The map generated by ORB-SLAM2 has redundant point clouds and blurry boundaries, while our VOOM builds a more lightweight yet semantic enhanced map, resulting in more accurate camera pose estimation.

some lightweight approaches using cuboid or dual quadrics as landmarks [7], [8]. Current algorithms mainly focus on constructing more accurate object landmarks [9], [10], and improving localization accuracy using object information [11]–[13]. Although these methods incorporate object residuals into bundle adjustment, they rely on an accurate initial guess of camera poses from feature-based techniques such as ORB-SLAM2 [2]. Their localization results are similar to or sometimes worse than feature-based methods. This indicates that current object-based methods have not fully utilized valuable object information to improve localization accuracy. How to further exploit object information to improve visual SLAM results remains challenging.

In this paper, we propose a novel Visual Object Odometry and Mapping framework, named VOOM, which applies both

low-level feature points and high-level objects to the entire process of odometry and mapping, as shown in Fig. 1. In odometry, we combine a residual model of dual quadrics based on the Wasserstein distance [14] with an object data association method based on the normalized Wasserstein distance to construct an accurate object-level map. By utilizing the object-level information, the current frame's keypoints are better matched with the previous map points, which results in an enhanced pose estimation for the frame. In terms of visual object mapping, we establish an object-based covisibility graph based on the observation relationships between keyframes and objects. The map points corresponding to the current neighbors in the graph are updated, thereby updating the point-based covisibility graph and the local map. Through a bundle adjustment, the poses of local keyframes and the positions of map points are optimized. In experiments, we have demonstrated that the proposed VOOM with only odometry and mapping components even outperforms the ORB-SLAM2 with loop closure in most sequences. To the best of our knowledge, this is the first paper to implement and show that object-oriented SLAM systems using dual quadrics and feature points can achieve better localization accuracy than state-of-the-art feature-based visual SLAM systems.

To sum up, our main contributions are threefold:

- the development of a novel visual object odometry and mapping framework using both feature points and dual quadrics as landmarks;
- the design of effective algorithms for object optimization, object association, and objects-based map points association for constructing the map with hierarchical landmarks;
- extensive experimental validation demonstrating the superiority of the proposed approach compared to state-of-the-art methods.

II. RELATED WORK

A. Visual Odometry and Mapping

Visual Odometry and mapping are often classified into direct and feature-based methods, depending on how they extract information from images to estimate camera motion and reconstruct maps.

Direct methods, such as LSD-SLAM [1] and Direct Sparse Odometry (DSO) [15], estimate camera motion and reconstruct maps by directly utilizing pixel intensity information from images. Direct Sparse Mapping (DSM) [16] extends DSO by reusing existing map information. Although the direct methods can handle low-textured and blurred images, they are sensitive to camera parameters and exposure.

In contrast, feature-based methods are more robust in variable lighting conditions by tracking extracted features, such as SIFT, BRIEF, and ORB. In early visual SLAM, filtering-based methods were widely used [17], [18] for tracking features thanks to its low computational cost. However, with the improvement of computing power, systems like PTAM [19] began to use bundle adjustment to optimize keyframe poses and divide the whole process into tracking

and mapping, which is more accurate than filtering methods. Based on these ideas, the monocular ORB-SLAM [20] tracks the more robust ORB features and limits the complexity of pose optimization and mapping by building a covisibility graph. ORB-SLAM2 [2] further extends the system for stereo and RGB-D cameras. ORB-SLAM2 has been proven to be versatile, efficient, and relatively accurate in practice. Therefore, in our work, we adopt their methods for points-level processing. Building upon this foundation, we introduce the object-level landmarks and seamlessly integrate them into the entire odometry and mapping process.

B. Object-oriented SLAM

Semantic SLAM can represent a high-level understanding of complex environments. Among different visual semantic SLAM approaches, object-oriented SLAM offers a lightweight solution that can enhance localization and mapping performance by leveraging object relations and geometric constraints.

SLAM++ [21] is an early proposed object-oriented mapping system, which shows its superiority in map capacity compared to dense semantic mapping without sacrificing descriptiveness. SLAM++ is not versatile, as it can only track domain-specific objects with known geometric shapes. In [22], arbitrary objects are built as point cloud models using bounding box detection and 3D segmentation. Based on these mapping systems, Fusion++ [5] further explores the utilization of reconstructed objects to enhance the precision of localization. Using only the 3D object information stored in a 6DoF pose graph, Fusion++ has achieved better accuracy. With the advent of CubeSLAM [7] and QuadricSLAM [11], researchers are motivated to use compact geometry shapes to represent the objects, instead of providing a detailed description of the object's surface. Compared with using cuboids used in [7], using dual quadrics in object-oriented SLAM has the superiority of mathematical completeness in projective geometry [8]. Some variations of quadrics-based approaches focus on constructing precise observation models for mapping [9], [10], [23], [24]. Other researchers have also attempted to improve the localization performance by adding objects residual into the cost function of optimization [11]–[13], but there is only negligible improvement over feature-based methods. OA-SLAM [25] has proposed an object-aided relocalization approach, which uses an object model for coarse pose estimate and points-based methods for refining the localization performance. It is indicated in OA-SLAM that objects excel at providing high-level semantic information and having strong discriminating power, whereas points offer greater accuracy in spatial localization.

Inspired by the coarse-to-fine idea of OA-SLAM, we propose a more logical strategy for enhancing odometry and mapping. Instead of directly introducing object residuals into optimization, we tightly incorporate 3D objects as landmarks to assist feature point associations. This enables us to utilize 3D map points repeatedly, not just for temporally adjacent frames but also for spatially adjacent ones. As a result, we can reduce the accumulation of errors in odometry.




Fig. 2: The proposed VOOM framework uses both low-level feature-based map points and high-level objects as landmarks. The color yellow indicates the process has only feature points participated, the color mint reveals that it is a pure object-based process, and the color blue means both landmarks contribute to the procedure.

III. VOOM

A. System Overview

The proposed VOOM takes an RGB-D image as input and outputs a 3D map composed of map points and 3D objects. The overall process is divided into two main parts: Visual Object Odometry and Visual Object Mapping, as shown in Fig.2. In the odometry part, the frontend processes feature points and information from instance segmentation results implemented by YOLOv8 [26]. We use the method in ORB-SLAM2 to process ORB feature points and pose prediction. In the object-level frontend, we estimate ellipses using the direct least squares method [27] based on a set of counter points obtained from instance segmentation. We determine whether objects are associated based on the Wasserstein distance between the estimated ellipse and the projection of the 3D ellipsoid onto the image plane, which we present in Sec. III-B and III-C in detail. In the odometry backend, we use associated objects to update the data association of map points based on the optimized frame pose, introduced in the first part of Sec. III-D. Once the frame is determined as a keyframe, it is passed to the Visual Object Mapping thread, introduced in the second part of Sec. III-D.

B. Observation Model of Objects

We perform instance segmentation to detect objects in each incoming image. It has been demonstrated in QISO-SLAM [13] that using contours of instance segmentation can achieve higher accuracy on object-oriented mapping and localization than using bounding boxes. In line with that, we estimate an observation ellipse E_{fk}^{obs} of a k -th object on f -th frame by minimizing the sum distance between each segmentation contour point and the ellipse outline.

The object can be represented by a constrained dual quadric $\mathbf{Q}_k^* \in \mathbb{R}^{4 \times 4}$. It is projected onto the image plane of f -th frame as a dual conic:

$$\mathbf{C}_{fk}^* = \mathbf{P}_f \mathbf{Q}_k^* \mathbf{P}_f^T, \quad (1)$$

where $\mathbf{P}_f = \mathbf{K}_f [\mathbf{R}_f | \mathbf{t}_f^T] \in \mathbb{R}^{3 \times 4}$ is the camera projection matrix that contains the matrix of the intrinsic parameters $\mathbf{K}_f \in \mathbb{R}^{3 \times 3}$, rotation matrix $\mathbf{R}_f \in \mathbb{R}^{3 \times 3}$ and translation vector $\mathbf{t}_f^T \in \mathbb{R}^{3 \times 1}$. \mathbf{C}_{fk}^* represents an estimated ellipse E_{fk}^{est} . An error function can thus be constructed with extracted ellipses of E_{fk}^{obs} and E_{fk}^{est} .

Inspired by OA-SLAM [25], we also utilize a 2D Gaussian distributions $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}^{-1})$ to interpret the ellipses:

$$\boldsymbol{\mu} = \begin{pmatrix} p_x \\ p_y \end{pmatrix}, \quad \boldsymbol{\Sigma}^{-1} = \mathbf{R}^T(\theta) \begin{pmatrix} \frac{1}{\alpha^2} & 0 \\ 0 & \frac{1}{\beta^2} \end{pmatrix} \mathbf{R}(\theta), \quad (2)$$

where $[p_x, p_y]^T$, $[\alpha, \beta]^T$ and θ describes the center point, the axes length, and the rotation angle of an ellipse, respectively. The error between E_{fk}^{obs} and E_{fk}^{est} is measured by the 2nd order Wasserstein distance between two Gaussian distributions $\mathcal{N}_{fk}^{\text{obs}}(\boldsymbol{\mu}_1, \boldsymbol{\Sigma}_1^{-1})$ and $\mathcal{N}_{fk}^{\text{est}}(\boldsymbol{\mu}_2, \boldsymbol{\Sigma}_2^{-1})$ is computed by:

$$W_2^2(\mathcal{N}_{fk}^{\text{obs}}, \mathcal{N}_{fk}^{\text{est}}) = \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 + \|\boldsymbol{\Sigma}_1^{\frac{1}{2}} - \boldsymbol{\Sigma}_2^{\frac{1}{2}}\|_F^2, \quad (3)$$

where $\|\cdot\|_F$ is the Frobenius norm. When the k -th object is observed by a set \mathcal{F}_L of frames, the parameters of the corresponding ellipsoid are estimated by minimizing:

$$J = \sum_{f \in \mathcal{F}_L} W_2^2(\mathcal{N}_{fk}^{\text{obs}}, \mathcal{N}_{fk}^{\text{est}}). \quad (4)$$

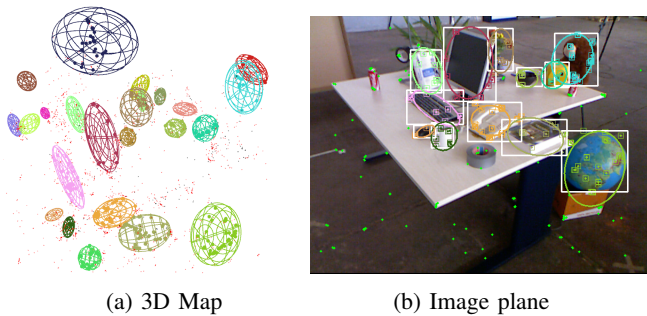


Fig. 3: Illustration of associating map points based on the associated objects.

C. Object-level Data Association

In previous works on object-oriented SLAM based on dual quadrics, the Intersection over Union (IoU) between reprojected boxes and detected bounding boxes was used as the criterion for object data association. However, there are multiple issues by using such a criterion. First, a bounding box enclosing an ellipse has a larger area than the ellipse itself, which increases the overlap among irrelevant objects. Second, we observe that the data association performance of tiny objects is severely affected by camera pose error, as the IoU-based metrics are highly sensitive to it. As we have employed Wasserstein distance to construct the residual model, it can also serve as a measurement for object data association. Normalized Wasserstein distance was proposed in [14] for tiny object detection. Based on (3), the normalized Wasserstein distance is written as:

$$W_n(\mathcal{N}_{fk}^{\text{est}}, \mathcal{N}_{fk}^{\text{obs}}) = \exp\left(-\frac{\sqrt{W_2^2(\mathcal{N}_{fk}^{\text{est}}, \mathcal{N}_{fk}^{\text{obs}})}}{C}\right). \quad (5)$$

Here, C is a constant. This metric incorporates the center point, length axes, and orientation information of the ellipses, which is more comprehensive than IoU.

D. Coarse-to-fine Odometry and Mapping

We present our method for visual object odometry and mapping in a coarse-to-fine manner. The main idea is to use the robustness of 3D object-level landmarks to enhance the efficiency and accuracy of feature-based map points to improve the odometry and mapping performance.

1) Coarse-to-fine Visual Object Odometry Backend: Given a coarse pose predicted by the motion model from the feature-level frontend and the associated objects from the object-level frontend, we search the matches between all map points corresponding to each object (marked as the colored large points inside each ellipsoid in Fig. 3a) and extracted feature keypoints belonging to the associated instance (marked as the colored rectangles inside each bounding box in Fig. 3b) by distance of ORB descriptors [20]. This allows us to achieve more efficient and precise map points association by narrowing down the search area, and to retrieve more map points from earlier frames, rather than from recent ones. With the matches between 3D map points




Fig. 4: Illustration of covisibility graphs.

and keypoints, the frame pose can be refined for the first time by minimizing the reprojection error. Next, we update the local map for the current frame by searching the previous keyframes that observe the associated map points and all map points that are related to those keyframes to match the current keypoints. Then we perform a finer pose optimization for the second time to obtain the current frame pose. Using the same criteria as in ORB-SLAM2, a keyframe is decided to be created and passed to the visual object mapping thread.

2) Coarse-to-fine Visual Object Mapping: After obtaining a reliable pose estimation and map points association for the current keyframe, we update the local map based on the object covisibility graph in visual object mapping. The object covisibility graph records the covisibility relations between different objects in different keyframes, as is shown in Fig. 4a. It can help us select a subset of keyframes most relevant to the current frame. In contrast to the traditional covisibility graph based on map points (illustrated in Fig. 4b), our object covisibility graph has a smaller data size and higher robustness. With the object-based covisibility graph, we can retrieve earlier keyframes that have previously observed objects to construct a local map for local bundle adjustment. Due to the robustness of objects, it implicitly achieves some of the effects of loop detection, even with only local mapping. In local BA, we use the retrieved keyframes and the corresponding map points, same as in [2]. Residuals of objects do not participate in BA because the object model itself is less accurate than the feature-based one, as shown in many previous works [11]–[13]. After BA, the keyframe poses and positions of map points are refined, and outliers are processed. The parameters of each object are optimized according to the poses of its observed keyframes separately.

IV. EXPERIMENTAL RESULTS

In this section, we present the experimental results of our proposed method. Our systems take the results of instance segmentation using YOLOv8 [26] as input, which is processed offline by Python. We only keep the detection with a confidence higher than 0.2 in the experiments. The online system is built upon the open-source C++ system ORB-SLAM2 with loop closure disabled. We demonstrate the localization performance in subsection IV-A, the object data association results in subsection IV-B, and an ablation study in IV-C on TUM RGB-D dataset [28] and LM-Data [29].

A. Localization

We conducted experiments on two datasets to compare the localization accuracy of our proposed method with several state-of-the-art methods. To quantify the localization accuracy, we use the root-mean-square error (RMSE) as the evaluation metric. We compare VOOM with ORB-SLAM2 [2], ORB-SLAM2 No LC, Fusion++ [5] QuadricSLAM [11], a RGB-D object SLAM [12] and QISO-SLAM [13] on TUM RGB-D Dataset, while we only compare with ORB-SLAM2 and ORB-SLAM2 No LC on LM-Data. ORB-SLAM2 is a representative feature-based SLAM method, and ORB-SLAM2 No LC is the same method with loop closure detection disabled. The other four frameworks are object-oriented SLAM, which attempted to improve the localization accuracy results using objects. Among them, Fusion++ uses refined dense object models to improve the TSDF odometry without traditional feature-based tracking. QuadricSLAM, RGB-D [12], and QISO-SLAM rely on the ORB-SLAM2 results as the initial camera pose and only use object information for refinement. In contrast, our method fuses feature points and object information in a unified framework to obtain the localization results directly.

Table I presents the localization performance on five data sequences: Fr1_desk1, Fr1_desk2, Fr2_desk, Fr2_person, and Fr3_office. The first column lists the sequence names, the second column shows the total lengths of the sequences, and the subsequent columns report the RMSE values obtained by each algorithm. The table reveals that our algorithm outperforms the others, especially in three longer sequences, and consistently improves upon ORB-SLAM2 No LC across all sequences. The other three object-oriented methods rely on the results of ORB-SLAM2 with loop closure. Still, they do not exhibit noticeable enhancement in localization accuracy, which is also analyzed and demonstrated in their papers. Another remarkable observation is that on the Fr2_person sequence with several dynamic objects, our localization result achieves a 92.8% error reduction over ORB-SLAM2 No LC. This further demonstrates the potential of our work to make contributions to dynamic SLAM.

Table II reports the results on four LM-Data Diamond sequences: Ground, Walk Head, Walk, and MAV. VOOM exceeds ORB-SLAM2 and ORB-SLAM2 No LC in all sequences. A notable result is obtained on the Diamond Walk sequence, where the camera travels long and revisits some same places. The error of ORB-SLAM2 without loop closure detection is over 5cm, and VOOM declined it by about 3/4. The results above prove that our proposed approach can effectively reduce the odometry drift without loop closure detection, as it utilizes both the robustness of objects and the accuracy of features. Furthermore, because VOOM updates and optimizes local maps with objects and feature points throughout the whole process, unlike loop closure which only optimizes when loops are detected, it even surpasses ORB-SLAM2 with loop closure. With the qualitative mapping results on the Diamond Ground sequence displayed in Fig. 5, we can better understand how VOOM

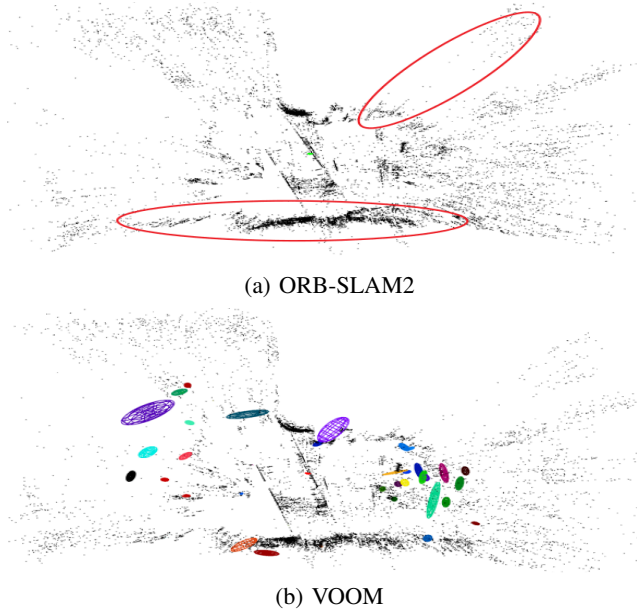


Fig. 5: The qualitative mapping results on the Diamond Ground Sequence. The obvious distorted parts of the ORB-SLAM2 map are circled with red color.

improves the performance¹. It is observed that ORB-SLAM2 cannot close the loop due to the limited view range of the ground robot, resulting in a redundant and distorted point cloud map. Nonetheless, our system is capable of correcting the distortion and building a more structured map when the camera re-observes certain objects.

B. Object Data Association Results

Both the observation model and object association influence the object association results. Thus, we evaluated and compared four object association methods and two observation models. The metric we used was the total number of reconstructed objects, reflecting the object-level map accuracy. The closer the number is to the ground truth, the better the data association performance is. Unlike many previous works, such as QuadricSLAM [11], OA-SLAM [25], and QISO-SLAM [13], we kept all the 3D objects added to the map without culling objects with low occurrence frequency to demonstrate the real effect of different methods in object data association. We used three sequences from two datasets: Fr2_desk from TUM RGB-D, Diamond Walk (D_Walk), and Diamond MAV (D_MAV) from LM-Data.

Table III summarizes the results. We compare four object data association methods: DA1, DA2, DA3, DA4. Among them, DA1 is the method used in OA-SLAM [25], DA2 uses IoU as metric, DA3 is the method used in QISO-SLAM2 [13], and DA4 is our proposed approach. The main metric for DA1-DA3 is IoU, with a basic threshold of 0.3. The difference is that DA1 only allows objects of the same category to be associated, while DA2 and D3 associate

¹Videos of this and other experiments can be found at <https://www.bilibili.com/video/BV1w14y1C7Jb/>

TABLE I: Localization RMSE comparison on TUM sequences.

Sequence	Length	ORB-SLAM2	ORB-SLAM2 No LC	Fusion++	QuadricSLAM	RGB-D [12]	QISO-SLAM	VOOM
Fr1_desk1	9.263m	0.0151	0.0193	0.049	0.0167	<u>0.0156</u>	0.0166	0.0189
Fr1_desk2	10.161m	0.0243	0.0264	0.153	0.0245	0.0210	<u>0.0220</u>	0.0245
Fr2_desk	18.880m	<u>0.0087</u>	0.0114	0.114	0.0124	0.0100	0.0099	0.0080
Fr2_person	17.044m	<u>0.0065</u>	0.0866	-	-	-	-	0.0062
Fr3_teddy	19.807m	<u>0.0370</u>	0.0417	-	-	-	-	0.0179
Fr3_office	21.455m	<u>0.0107</u>	0.0113	0.108	0.0230	0.0111	0.0112	0.0100

The best result is highlighted in **bold**, and the second best is underlined.

TABLE II: Localization RMSE comparison on LM-Data Diamond sequences. The best result is highlighted in **bold**.

Sequence	Length	ORB-SLAM2	ORB-SLAM2 No LC	VOOM
Ground	5.87m	0.1658	0.1665	0.1287
Walk Head	13.82m	0.0135	0.0193	0.0125
Walk	22.56m	0.0213	0.0596	0.0150
MAV	33.84m	0.0132	0.0195	0.0114

TABLE III: Object Data Association Results.

Dataset	DA1		DA2		DA3		DA4 (Ours)		GT
	M1	M2	M1	M2	M1	M2	M1	M2	
Fr2_desk	131	119	68	64	55	54	52	49	44
D_Walk	73	74	71	59	58	55	59	53	50
D_MAV	74	80	73	78	51	59	50	52	43

objects of any category. Moreover, DA3 also considers shape similarity and lowers the IoU threshold to 0.1 when the object labels are the same. As for our proposed DA4, We empirically set the constant C in (5) as $C = 10$ and a very low threshold of 0.005 as we found it is not as sensitive as the IoU threshold. We also compare two different observation models: M1 and M2. M1 is the model used in OA-SLAM, and M2 is our proposed one. As can be seen, our proposed DA4 outperforms the other methods in most cases. It is observed that DA1 has the most number of redundant objects among all data association methods. This is because it is very common for an object to be classified into different categories in real cases. DA3 performs better than DA2, as it incorporates more criteria for object data association. As for the observation model, M2 shows better performance than M1 overall. Particularly on Fr2_desk and Diamond Walk sequences, all data related to M2 are superior to M1. The only exception is the Diamond MAV sequence, where M2 is slightly inferior to M1. It may be because objects are estimated with larger sizes using M1, which helps object association in faster-speed scenarios. In general, the combination of DA4 and M2, which we proposed, is suitable for most situations, while in some cases, using DA4 and M1 can also achieve sufficient results.

C. Ablation Study

We conducted an ablation study to evaluate the contribution of different parts of VOOM. Besides the full VOOM (F), we also show the results of three variants: VOOM (O),

TABLE IV: Ablation Study. The best result is highlighted in **bold**, and the second best is underlined.

Dataset	ORB No LC	VOOM (O)	VOOM (M)	VOOM (A)	VOOM (F)
TUM Fr2_desk	0.0114	0.0086	<u>0.0084</u>	0.0105	0.0080
TUM Fr2_person	0.0866	0.0065	<u>0.0063</u>	0.0071	0.0062
Diamond Walk	0.0596	0.0205	<u>0.0172</u>	0.0249	0.0150

VOOM (M), and VOOM (A). VOOM (O) represents that we use objects only in odometry, while VOOM (M) means objects only contribute to mapping. In VOOM (A) we replaced the proposed data association method and observation model with the ones used in OA-SLAM to demonstrate their impact on localization performance.

The RMSE of the ablation study is presented in Table IV. We chose Fr2_desk and Fr2_person from the TUM RGB-D dataset, and Diamond Walk from the LM-Data. The baseline is ORB-SLAM2 No LC, which is the basis of our algorithm. First, all variants surpass the original ORB-SLAM2 No LC, which demonstrates the effectiveness of each improvement in our framework. Furthermore, the results show that the full VOOM algorithm achieves the best performance, followed by VOOM (M), which indicates that local bundle adjustment with objects contributes slightly more than pose estimation in odometry. Additionally, VOOM (A) is worse than the full VOOM, which suggests that the superior data association method and observation model we proposed can also enhance localization accuracy.

V. CONCLUSION

Our experimental results have demonstrated that by incorporating object information, VOOM is able to achieve better localization performance than the state-of-the-art feature-based visual SLAM methods. This level of improvement has not been previously observed in object-oriented SLAM approaches using dual quadrics as landmarks, to the best of our knowledge. Our system has the odometry and mapping components, with a primary focus on short-term and mid-term performance. In future work, we plan to build upon this foundation by introducing a loop closure and relocalization algorithm that integrates both objects and points efficiently, with the goal of achieving improved long-term SLAM performance.

REFERENCES

- [1] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," in *Proc. of the Europ. Conf. on Computer Vision (ECCV)*. Springer, 2014, pp. 834–849.
- [2] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE Trans. on Robotics (TRO)*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [3] T. Qin, P. Li, and S. Shen, "Vins-mono: A robust and versatile monocular visual-inertial state estimator," *IEEE Trans. on Robotics (TRO)*, vol. 34, no. 4, pp. 1004–1020, 2018.
- [4] M. Runz, M. Buffier, and L. Agapito, "MaskFusion: Real-Time Recognition, Tracking and Reconstruction of Multiple Moving Objects," in *Proc. of the Intl. Symposium on Mixed and Augmented Reality (ISMAR)*, 2018, pp. 10–20.
- [5] J. McCormac, R. Clark, M. Bloesch, A. Davison, and S. Leutenegger, "Fusion++: Volumetric Object-Level SLAM," in *Proc. of the Intl. Conf. on 3D Vision (3DV)*, 2018, pp. 32–41.
- [6] A. Rosinol, M. Abate, Y. Chang, and L. Carlone, "Kimera: an open-source library for real-time metric-semantic localization and mapping," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*. IEEE, 2020, pp. 1689–1696.
- [7] S. Yang and S. Scherer, "Cubeslam: Monocular 3-d object slam," *IEEE Trans. on Robotics (TRO)*, vol. 35, no. 4, pp. 925–938, 2019.
- [8] C. Rubino, M. Crocco, and A. Del Bue, "3D Object Localisation from Multi-View Image Detections," *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 40, no. 6, pp. 1281–1294, 2018.
- [9] Z. Liao, Y. Hu, J. Zhang, X. Qi, X. Zhang, and W. Wang, "So-slam: Semantic object slam with scale proportional and symmetrical texture constraints," *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 2, pp. 4008–4015, 2022.
- [10] X. Lin, Y. Yang, L. He, W. Chen, Y. Guan, and H. Zhang, "Robust improvement in 3d object landmark inference for semantic mapping," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*. IEEE, 2021, pp. 13 011–13 017.
- [11] L. Nicholson, M. Milford, and N. Sünderhauf, "Quadricslam: Dual quadrics from object detections as landmarks in object-oriented slam," *IEEE Robotics and Automation Letters (RA-L)*, vol. 4, no. 1, pp. 1–8, 2018.
- [12] Z. Liao, W. Wang, X. Qi, and X. Zhang, "Rgb-d object slam using quadrics for indoor environments," *Sensors*, vol. 20, no. 18, p. 5150, 2020.
- [13] Y. Wang, B. Xu, W. Fan, and C. Xiang, "Qiso-slam: Object-oriented slam using dual quadrics as landmarks based on instance segmentation," *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 4, pp. 2253–2260, 2023.
- [14] J. Wang, C. Xu, W. Yang, and L. Yu, "A normalized gaussian wasserstein distance for tiny object detection," *arXiv preprint*, 2021.
- [15] J. Engel, V. Koltun, and D. Cremers, "Direct sparse odometry," *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 40, no. 3, pp. 611–625, 2017.
- [16] J. Zubizarreta, I. Aguinaga, and J. M. M. Montiel, "Direct sparse mapping," *IEEE Trans. on Robotics (TRO)*, vol. 36, no. 4, pp. 1363–1370, 2020.
- [17] A. J. Davison, I. D. Reid, N. D. Molton, and O. Stasse, "Monoslam: Real-time single camera slam," *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 29, no. 6, pp. 1052–1067, 2007.
- [18] J. Civera, O. G. Grasa, A. J. Davison, and J. M. Montiel, "1-point ransac for extended kalman filtering: Application to real-time structure from motion and visual odometry," *Journal of Field Robotics (JFR)*, vol. 27, no. 5, pp. 609–631, 2010.
- [19] G. Klein and D. Murray, "Parallel tracking and mapping for small ar workspaces," in *Proc. of the Intl. Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2007, pp. 225–234.
- [20] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "Orb-slam: a versatile and accurate monocular slam system," *IEEE Trans. on Robotics (TRO)*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [21] R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. J. Kelly, and A. J. Davison, "SLAM++: Simultaneous Localisation and Mapping at the Level of Objects," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013, pp. 1352–1359.
- [22] N. Sünderhauf, T. T. Pham, Y. Latif, M. Milford, and I. Reid, "Meaningful maps with object-oriented semantic mapping," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2017, pp. 5079–5085.
- [23] Y. Wu, Y. Zhang, D. Zhu, Y. Feng, S. Coleman, and D. Kerr, "Eao-slam: Monocular semi-dense object slam based on ensemble data association," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS)*, 2020, pp. 4966–4973.
- [24] Y. Wu, Y. Zhang, D. Zhu, Z. Deng, W. Sun, X. Chen, and J. Zhang, "An object slam framework for association, mapping, and high-level tasks," *IEEE Trans. on Robotics (TRO)*, 2023.
- [25] M. Zins, G. Simon, and M.-O. Berger, "Oa-slam: Leveraging objects for camera relocalization in visual slam," in *Proc. of the Intl. Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 2022, pp. 720–728.
- [26] G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," Jan. 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
- [27] A. Fitzgibbon, M. Pilu, and R. B. Fisher, "Direct least square fitting of ellipses," *IEEE Trans. on Pattern Analysis and Machine Intelligence (TPAMI)*, vol. 21, no. 5, pp. 476–480, 1999.
- [28] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *Proc. of the International Conference on Intelligent Robot Systems (IROS)*, Oct. 2012.
- [29] S. Saeedi, E. D. Carvalho, W. Li, D. Tzoumanikas, S. Leutenegger, P. H. Kelly, and A. J. Davison, "Characterizing visual localization and mapping datasets," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA)*. IEEE, 2019, pp. 6699–6705.