

V3D-SLAM: Robust RGB-D SLAM in Dynamic Environments with 3D Semantic Geometry Voting

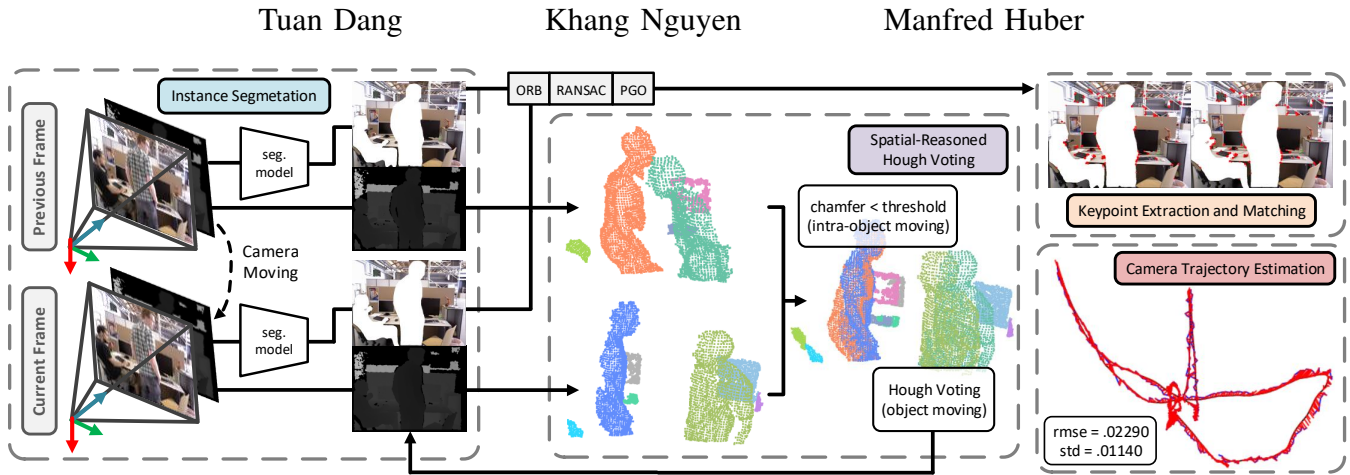


Fig. 1: Overview of *V3D-SLAM*: improving the robustness of RGB-D SLAM in dynamic indoor environments, including instance segmentation coupled with RGB-based feature extraction (Sec. IV-A), sensor noises and segmentation outlier rejection (Sec. IV-B), and spatial-reasoned Hough voting mechanism for dynamic 3D objects (Sec. IV-C), resulting in camera trajectory estimation (Sec. V).

Abstract—Simultaneous localization and mapping (SLAM) in highly dynamic environments is challenging due to the correlation complexity between moving objects and the camera pose. Many methods have been proposed to deal with this problem; however, the moving properties of dynamic objects with a moving camera remain unclear. Therefore, to improve SLAM’s performance, minimizing disruptive events of moving objects with a physical understanding of 3D shapes and dynamics of objects is needed. In this paper, we propose a robust method, *V3D-SLAM*, to remove moving objects via two lightweight re-evaluation stages, including identifying potentially moving and static objects using a spatial-reasoned Hough voting mechanism and refining static objects by detecting dynamic noise caused by intra-object motions using Chamfer distances as similarity measurements. Through our experiment on the TUM RGB-D benchmark on dynamic sequences with ground-truth camera trajectories, the results show that our methods outperform most other recent state-of-the-art SLAM methods. Our source code is available at <https://github.com/tuandang/v3d-slam>.

I. INTRODUCTION

Visual simultaneous localization and mapping (vSLAM), an important study of robotics, essentially relies on visual information from the camera to localize itself and build a map of the environment. With the emergence of low-cost cameras, vSLAM captures tremendous attention from the research community, especially RGB-D-based SLAM, which is one of the most popular choices for its availability and appropriateness for indoor scenes. A number of notable works have been proposed, such as Dense-SLAM [1] and ORB-SLAM [2]; however, like traditional SLAM methods

[3], [4], these approaches implicitly assume that the environment is static, where frame-to-frame extracted keypoints are matched, often failing to deal with dynamic scenes when objects and the camera are moving simultaneously.

Although vSLAM has been extensively studied with help from recent advancements in computer vision, especially with deep neural networks, some issues have not been well-addressed due to the dynamics of the environment. DS-LAM [5] uses a segmentation model and optical flow to detect moving objects and consider these moving objects as noises to be removed. TRS-SLAM [6] deals with objects on the training dataset while segmenting the depth image via k-means clustering to handle unknown objects, which does not require prior knowledge. CFP-SLAM [7] detects moving objects and uses Kalman filter with Hungarian algorithm to compensate for object misdetections. These methods obtain high accuracy in dynamic environments, and their results are considered state-of-the-art on the TUM RGB-D benchmark [8]. While the DefSLAM [9] tracks deformable objects with a parametric template, which limits the method to a small number of objects that are matched with the template, we are seeking a non-parametric method to detect deformable objects, paving the way to detect the intra-moving objects. Nevertheless, their flip side is to represent moving objects as center points of segmentation blobs or bounding boxes, which are easily distorted by different views. Another scenario in which these methods fail is when only a part of the object moves. Still, its center remains unchanged (*i.e.*, a person wobbles his head without displacing him from one place to another or rotating a revolving chair). These motions within an object cause significant noise to the feature extractor, leading to feature inconsistency between frames

All authors are with the Learning and Adaptive Robotics Laboratory, Department of Computer Science and Engineering, University of Texas at Arlington, Arlington, TX 76013, USA. (emails: tuandang@uta.edu, khang.nguyen8@mavs.uta.edu, huber@cse.uta.edu)



Fig. 2: Segmentation of hypothesized moving objects with keypoints on static objects and background, and point clouds of object instances.

and, eventually, errors in camera pose estimation.

To address the above issues, we propose *V3D-SLAM*, which differentiates moving objects from static objects by projecting depths into point clouds and analyzing their 3D shapes and geometry. *V3D-SLAM* first identifies potentially moving objects by a novel Hough voting mechanism via the topology of 3D objects in one frame, resulting in two set objects (static objects and moving objects), thus analyzing the moving parts within static objects to seek moving parts via measuring the similarity using Chamfer distances [10].

II. RELATED WORK

RGB-D SLAM in Dynamic Environments: Estimating poses of a moving camera has been an interesting topic in robotic vision, with the RGB-based earliest example using Tomasi–Kanade factorization [11]. This has laid the foundations for later research addressing dynamic objects in SLAM, including associating correspondences between RGB images [1], [12], [13], weighting edge-like features for tracking [14], identifying objects’ movements via differences of consecutive frames [15], differentiating static and dynamic objects via feature correlations [16], [17]. However, these methods only deal with feature extraction on consecutive frames but lack prior semantic knowledge of the scene. To enhance this, recent works in vSLAM leverage object detection [18], [19], [7] and segmentation [5], [20], [21], [6], [22], [23] models to extract RGB-based knowledge of the scene and use depth frames to reconstruct the scene’s structure. Indeed, these methods only concentrate on 2D object movement via epipolar geometry while not guaranteeing 3D intra-object movement, which may be visually unchanged on RGB images. Furthermore, DefSLAM [9] uses the precomputed template to detect deformable objects [24], leading to a lack of generality to other categories. To address this issue, we employ the non-parametric method, which re-associates depth frames to information extracted from colored frames and computes the similarity between the current and subsequent frames.

Spatial-Reasoned Hough Voting for Dynamic Objects: The voting mechanism in Hough transform [25], searching for the parameter that casts the most votes where each sample votes for another in parameter spaces, has been used widely for object detection [26] with explicit and implicit features. DEHV [27] provides the Hough-based probabilistic approach, in which each object class and location cast the votes from depth scales. Max-margin Hough transform [28] is also introduced to indicate the important weights for possible local locations of the object center. PoseCNN [29] uses Hough voting to predict object centers via network-based extracted features. Similarly, VoteNet [30] samples point

clouds to extract point features where votes are gained from the point feature domain. In this work, we devise a similar approach where the 3D dynamic objects are voted through the topology with other presented 3D objects. Instead of learning implicitly by using a neural network, each object will cast votes directly from pure geometry information.

III. OVERVIEW OF *V3D-SLAM*

The overview of *V3D-SLAM*, improving the robustness of dynamic RGB-D SLAM, is illustrated in Fig. 1. In between two consecutive frames, we first mask out potential moving objects from the scene [31], [32] to obtain static objects with backgrounds, followed by reconstructing point clouds of segmented instances. Noises induced by the segmentation model’s uncertainty create fragments in point clouds of instance (Fig. 3), which spatially do not belong to the objects, slightly shift the center of the point clouds, and influence the votes for potential dynamic objects significantly. To avoid these ambiguities, we statistically remove outliers at the instance level to eliminate segmentation-induced artifacts (Sec. IV-B) before identifying instances’ spatial centers. Thus, based on the topology of 3D objects between two frames, the voting mechanism (Fig. 4) is implemented to identify moving objects with geometric and spatial information from themselves for other presented entities (Sec. IV-C).

In many indoor dynamic environments, the centers of large objects do not change from frame to frame. Still, some parts may be displaced, introducing another type of noise in the sequence. To identify these dynamic intra-objects, we track and calculate 3D objects’ similarities via object deformation instead of objects’ relative displacement on image planes due to their errors and distortion in different views (Sec. IV-D). Within this tracking procedure, based on the physical constraints of objects’ movements, we assume that the objects keep moving in the same direction in an infinitesimal period (between two frames in a 30-fps sequence) after identifying each moving object to compensate for the object misdetection or object that is out of field-of-view. Pixel-level features are also extracted on unmasked regions using ORB [33] for camera pose estimation followed by camera trajectory optimization using pose graph optimization (PGO) [34] (Sec. IV-A).

IV. METHODOLOGY

A. Instance Segmentation & Feature Extraction

We first segment the objects in RGB images using YOLOv8 [31], then map the masked regions into depth images to make point clouds of instances. Hypothesizing as potentially moving, the objects, which are classified as ‘person’, are temporarily excluded from the scene to be

re-evaluated for their movements in 3D space (Sec. IV-C). Thus, we detect edge- and corner-like features using ORB [33] on unmasked regions containing static objects and the scene background. Not only for its segmenting capability, the segmentation model also gives us the number of objects, which benefits our subsequent noise removal procedure.

B. Sensor Noises & Segmentation Outlier Rejection

In the context of SLAM, it is implicitly assumed that either the object or the robot is still. However, in dynamic environments, the fast-moving robot with the RGB-D camera produces perception uncertainty, which is best seen via RGB images, where the pixels are blurred, and depth maps, where the depth pixel cannot be interpolated from the previous frame [35]. Furthermore, segmentation models add uncertainty when recognizing objects with RGB-D perception. To alleviate this, we categorize errors into two types: (1) depth map noises and (2) segmentation-induced mapping errors.

1) *Depth Map Noises*: We re-estimate the depth value, $d(i, j)$, on depth maps via a $k \times k$ Gaussian kernel, $G(i, j)$, centered at that pixel to prevent potential errors caused by captured devices when calculating objects' center locations. The 2D-3D re-projection discretization is re-fined as below:

$$z = \int_{i=u-k}^{u+k} \int_{j=v-k}^{v+k} d(i, j) \cdot G(u, v)$$

$$= \sum_{i=u-k}^{u+k} \sum_{j=v-k}^{v+k} d(i, j) \cdot \frac{1}{2\pi\sigma^2} e^{-(u^2+v^2)/2\sigma^2}$$

$$x = \frac{u - c_x}{f_x} \cdot z \text{ and } y = \frac{v - c_y}{f_y} \cdot z$$

where (u, v) represents the image coordinates of pixels, (x, y, z) are their 3D coordinates, (c_x, c_y, f_x, f_y) are defined as the camera's optical center and focal lengths, and σ^2 is the variance of k^2 pixels locating in the $k \times k$ Gaussian kernel.

2) *Segmentation-Induced Mapping Errors*: As shown in Fig. 2, RGB-based segmentation introduces pixel-level sensitivity when projecting image blobs via the corresponding depth maps to create point clouds of instances. Therefore, we first downsample point clouds through voxelization to efficiently remove these artifacts to guarantee balance and fairness during noise removal between dense and spare regions, as shown in Fig. 3 (right). Thus, we use density-based spatial clustering of applications with noises (DBSCAN) algorithm [36] to group points into different clusters and hence filter out clusters by taking the n -first 3D large blobs with n is the number of objects recognized provided by the segmentation model in Sec. IV-A, resulting in a finer point clouds of 3D semantic objects, as depicted in Fig. 3 (right).

C. Spatial-Reasoned Hough Voting for Dynamic Objects

Identifying dynamic objects in 2D images is complicated and sensitive to the dynamics of the environment, especially when the camera and objects move together since the object-to-object displacements are not correctly spatially interpreted with different projections from a moving camera. For example, objects can move at the same velocity as cameras,



Fig. 3: Outlier removal with semantic perception on point clouds, resulting in unchanged dynamic interpretations based solely on RGB images.

Algorithm 1: Hough Voting for Dynamic Objects

Input : \mathbf{O}_p := objects in previous frame
 \mathbf{O}_c := objects in current frame
 $\mathcal{T}_d, \mathcal{T}_v$:= distance and voting counts thresholds
Output: *DynamicObjects* := lists of voted dynamic objects

```

1 function VoteDynamicObjects( $\mathbf{O}_p, \mathbf{O}_c, \mathcal{T}_d, \mathcal{T}_v$ )
2    $\mathbf{keys}_p, \mathbf{dist}_p = \text{PairwiseCenterDist}(\mathbf{O}_p)$ 
3    $\mathbf{keys}_c, \mathbf{dist}_c = \text{PairwiseCenterDist}(\mathbf{O}_c)$ 
4   accumulator = {}
5   for  $k_p \in \mathbf{keys}_p$  do
6     for  $k_c \in \mathbf{keys}_c$  do
7       if  $k_p = k_c$  and  $\text{dist}(d_p[k_p], d_c[k_c]) \geq \mathcal{T}_d$  then
8         id = ExtractObjectId( $\mathbf{key}_c$ )
9         accumulator[id] += 1
10  DynamicObjects = accumulator.where(id  $\geq \mathcal{T}_v$ )
11  return DynamicObjects

```

To alleviate these concerns, we interpret scene dynamics by computing Euclidean distances between centroids of 3D instances in the current frame and estimating the displacement of objects in the next frame. If the displacement is larger than the pre-defined threshold, the accumulator $V(\mathbf{O}_i|\mathbf{O}_j, r)$ casts the vote from object j to object i for the high frame rate, $r \approx 30$; on the other hand, low-fps sequences neither can track and recognize objects.

$$V(\mathbf{O}_i|\mathbf{O}_j, r) = \sum_j \mathbf{I}[\text{dist}(\mathbf{O}_i, \mathbf{O}_j|r)] \quad (1)$$

where $\text{dist}(\cdot)$ represents the Euclidean distance between 3D objects' centroids and $\mathbf{I}(\cdot)$ denotes the indicator function and is as follows with the defined distance threshold, \mathcal{T}_d :

$$\mathbf{I}[\text{dist}(\mathbf{O}_i, \mathbf{O}_j|r)] = \begin{cases} 1, & \text{if } \text{dist}(\mathbf{O}_i, \mathbf{O}_j|r) \geq \mathcal{T}_d \\ 0, & \text{otherwise.} \end{cases}$$

Alg. 1 takes in objects, \mathbf{O} , in previous (p) and current (c) frames as inputs and outputs the keys and IDs of objects that are voted as moving. Specifically, a class and object ID are assigned for each presented object instance in the frame. This information is tracked until the objects are no longer given in the scene. In contrast, the 3D pairwise distances of objects' centers are calculated among in-frame objects, and votes are cast among objects based on Eq. 1.

D. Intra-Object Movements

Objects like revolving chairs or people sitting without shifting their locations can cause noises by moving their heads or rotating the chair while talking with others. We

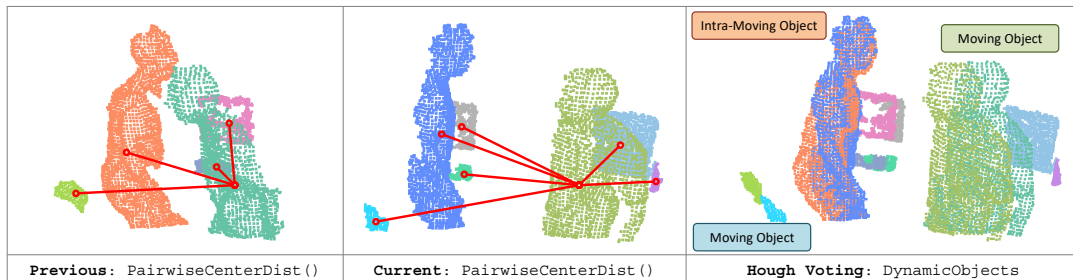


Fig. 4: Spatial-reasoned Hough voting mechanism for moving objects (right) after computing accumulator array in previous (left) and current (middle) frames. Red lines illustrate Euclidean distances between one of the ‘person’ objects and other presented entities. The moving ‘person’ object and the ‘chair’ object are identified using Alg. 1; meanwhile, the other ‘person’ object is only identified as an intra-moving object as he wobbles his head and his 3D centroid does not exceed the distance threshold for deformable objects.

treat these objects as deformable objects that must be considered when extracting their keypoints. To solve this, we use Chamfer distance to measure the similarity between two point clouds of segmented moving objects, $\mathcal{D}(\mathcal{P}_1, \mathcal{P}_2)$, in consecutive frames:

$$\mathcal{D}(\mathcal{P}_1, \mathcal{P}_2) = \frac{1}{|\mathcal{P}_1|} \left[\sum_{\mathbf{p}_i \in \mathcal{P}_1} \min_{\mathbf{p}_j \in \mathcal{P}_2} \|\mathbf{p}_i - \mathbf{p}_j\|_2^2 \right] + \frac{1}{|\mathcal{P}_2|} \left[\sum_{\mathbf{p}_j \in \mathcal{P}_2} \min_{\mathbf{p}_i \in \mathcal{P}_1} \|\mathbf{p}_j - \mathbf{p}_i\|_2^2 \right] \quad (2)$$

where \mathcal{P}_1 and \mathcal{P}_2 represent two clouds, respectively, \mathbf{p}_i and \mathbf{p}_j are i^{th} and j^{th} points in \mathcal{P}_1 and \mathcal{P}_2 , respectively, and $|\mathcal{P}_i|$ indicates the number of points in the point cloud \mathcal{P}_i .

Using Eq. 2, if an object is calculated as the deformable object as such definition, we also included features extracted from their RGB frame for camera pose estimation.

E. Camera Pose Estimation

After identifying dynamic objects, we mask them out of regions of interest and consider them as noises on images, leaving images constructed by static objects by changing the camera views. By extracting and tracking features on static objects and the scene background when the camera moves, we thus estimate the camera poses and trajectory. Inspired by the robustness of this task from ORB-SLAM2 [12] and ORB-SLAM3 [37], we apply feature extraction on unmasked RGB images to extract keypoints on two consecutive frames, then matching the corresponding features between two keypoint sets. At this point, we assume two point sets are static point sets with different camera views. The transformation matrices are computed using RANSAC while removing outliers on

these pointsets that do not satisfy the triangulation constraint of epipolar geometry of a pair RGBD. The two new point sets are used to look at depth to reconstruct 3D point sets, which are used to estimate the camera pose, including translation and rotation, using the RANSAC method. The estimated camera pose is added to the trajectory for closure detection and trajectory optimization using PGO [34].

V. EVALUATION ON TUM RGB-D BECHMARK

A. Testing Sequences & Evaluation Metrics

To test the performance of our proposed technique with spatial-reasoned votes for dynamic 3D objects, we evaluate Alg. 1 on the TUM RGB-D benchmark, containing eight dynamic sequences in terms of Absolute Trajectory Error (ATE) and Relative Pose Error (RPE), including translational and rotational drifts, in both root mean square error (RMSE) and standard deviation (SD). Across these metrics, we also compare our method against state-of-the-art vSLAM techniques: ORB-SLAM2 [12], DS-SLAM [5], DynaSLAM [20], TRS [6], Blitz-SLAM [22], and CFP-SLAM [7].

B. Quantitative Results

Table I shows the quantitative comparisons between our method and other methods for the metric of ATE: for most of the sequences, we are able to obtain lower RMSE and SD compared to the recent state-of-the-art CFP-SLAM, except *fr3/w/xyz* and *fr3/w/rpy*. Table II and Table III depict the RPE in terms of the camera’s translational and rotational drifts, respectively. We achieve better estimates for translational drift, but the *fr3/s/half* sequence, where we underperform Blitz-SLAM and CFP-SLAM. For rotational drift, we obtain finer results in *fr3/s/xyz*, *fr3/s/rpy*, *fr3/w/half*, and *fr3/w/static* sequences, which

TABLE I: Comparisons of ATE between RGB-D SLAM techniques.

Sequences	ORB-SLAM2 [12]		DS-SLAM [5]		DynaSLAM [20]		TRS [6]	Blitz-SLAM [22]		CFP-SLAM ⁻ [7]		CFP-SLAM [7]		V3D-SLAM (Ours)	
	RMSE	SD	RMSE	SD	RMSE	SD	RMSE	RMSE	SD	RMSE	SD	RMSE	SD	RMSE	SD
<i>fr3/s/xyz</i>	0.0092 [§]	0.0047 [§]	–	–	0.0127	0.0060	0.0117	0.0148	0.0069	0.0129	0.0068	0.0090 [†]	0.0042 [*]	0.0087 [*]	0.0043 [†]
<i>fr3/s/half</i>	0.0192	0.0110	–	–	0.0186	0.0086	0.0172	0.0160	0.0076	0.0159 [†]	0.0072 [§]	0.0147 [*]	0.0069 [†]	0.0147 [*]	0.0066 [*]
<i>fr3/s/static</i>	0.0087	0.0042	0.0065	0.0033	–	–	–	–	–	0.0061 [§]	0.0029 [†]	0.0053 [*]	0.0027 [*]	0.0058 [†]	0.0031 [§]
<i>fr3/s/rpy</i>	0.0195 [†]	0.0124 [†]	–	–	–	–	–	–	–	0.0244 [§]	0.0175 [§]	0.0253	0.0154	0.0169 [*]	0.0101 [*]
<i>fr3/w/xyz</i>	0.7214	0.2560	0.0247	0.0161	0.0164	0.0086	0.0194	0.0153 [§]	0.0078	0.0149 [†]	0.0077 [†]	0.0141 [*]	0.0072 [*]	0.0153 [§]	0.0080 [§]
<i>fr3/w/half</i>	0.4667	0.2601	0.0303	0.0159	0.0296	0.0157	0.0290	0.0256	0.0126	0.0235 [†]	0.0114 [*]	0.0237 [†]	0.0114 [*]	0.0229 [*]	0.0114 [*]
<i>fr3/w/static</i>	0.3872	0.1636	0.0081	0.0036	0.0068 [§]	0.0032 [§]	0.0111	0.0102	0.0052	0.0069	0.0032 [§]	0.0066 [†]	0.0030 [†]	0.0065 [*]	0.0028 [*]
<i>fr3/w/rpy</i>	0.7842	0.4005	0.4442	0.2350	0.0354 [*]	0.0190 [*]	0.0371 [§]	0.0356 [†]	0.0220 [†]	0.0411	0.0250	0.03680	0.0230 [§]	0.0781	0.0360

are more significant compared to ORB-SLAM2 and CFP-SLAM. For notations, we use **bold asterisk***, **bold dagger†**,

and **bold section sign§** to highlight the lowest error, the second-lowest error, and the third- error, respectively.

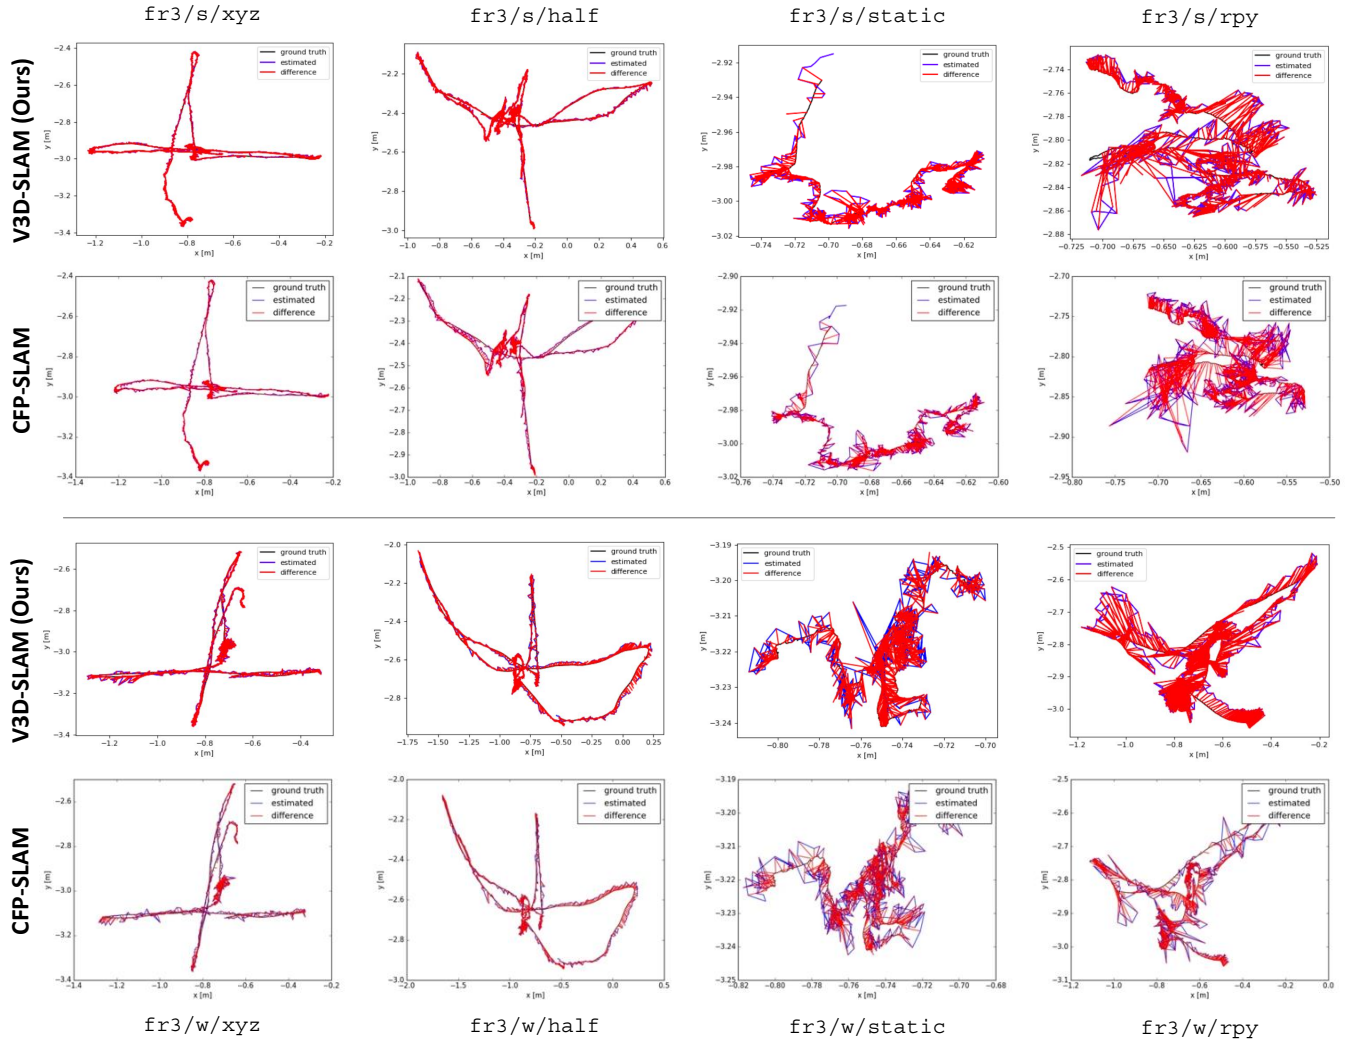


Fig. 5: Qualitative results of camera trajectories of TUM RGB-D dynamic sequences estimated by our method and CFP-SLAM [7]. The ground truth, the estimated trajectory, and their differences are encoded as black lines, blue lines, and red lines, respectively.

Sequences	ORB-SLAM2 [12]		DS-SLAM [5]		DynaSLAM [20]		TRS [6]	Blitz-SLAM [22]		CFP-SLAM ⁻ [7]		CFP-SLAM [7]		V3D-SLAM (Ours)	
	RMSE	SD	RMSE	SD	RMSE	SD	RMSE	RMSE	SD	RMSE	SD	RMSE	SD	RMSE	SD
fr3/s/xyz	0.0117[§]	0.0060[§]	-	-	0.0142	0.0073	0.0166	0.0144	0.0071	0.0149	0.0081	0.0114[†]	0.0055[†]	0.0105*	0.0051*
fr3/s/half	0.0231	0.0163	-	-	0.0239	0.0120	0.0259	0.0165[†]	0.0073*	0.0214	0.0099	0.0162*	0.0079[†]	0.0184[§]	0.0088[§]
fr3/s/static	0.0090	0.0043	0.0078[§]	0.0038	-	-	-	-	-	0.0078[§]	0.0034*	0.0072[†]	0.0035[†]	0.0068*	0.0037[§]
fr3/s/rpy	0.0245[†]	0.0144[†]	-	-	-	-	-	-	-	0.0322	0.0217	0.0316[§]	0.0186[§]	0.0221*	0.0127*
fr3/w/xyz	0.3944	0.2964	0.0333	0.0229	0.0217	0.0119	0.0234	0.0197	0.0096*	0.0196[§]	0.0099[§]	0.0190*	0.0097[†]	0.0193[†]	0.0100
fr3/w/half	0.3480	0.2859	0.0297	0.0152	0.0284	0.0149	0.0423	0.0253[†]	0.0123[†]	0.0274	0.0130	0.0259[§]	0.0128[§]	0.0242*	0.0115*
fr3/w/static	0.2349	0.2151	0.0102	0.0048	0.0089	0.0044	0.0117	0.0129	0.0069	0.0092[§]	0.0043[§]	0.0089[†]	0.0040[†]	0.0078*	0.0035*
fr3/w/rpy	0.4582	0.3447	0.1503	0.1168	0.0448*	0.0262*	0.0471[†]	0.0473	0.0283[†]	0.0540	0.0350	0.0500[§]	0.0306[§]	0.0726	0.0480

TABLE II: Comparisons of Translational Drift in RPE between RGB-D SLAM techniques.

Sequences	ORB-SLAM2 [12]		DS-SLAM [5]		DynaSLAM [20]		TRS [6]	Blitz-SLAM [22]		CFP-SLAM ⁻ [7]		CFP-SLAM [7]		V3D-SLAM (Ours)	
	RMSE	SD	RMSE	SD	RMSE	SD	RMSE	RMSE	SD	RMSE	SD	RMSE	SD	RMSE	SD
fr3/s/xyz	0.4890[§]	0.2713	-	-	0.5042	0.2651	0.5968	0.5024	0.2634[†]	0.5126	0.2793	0.4875[†]	0.2640[§]	0.4825*	0.2577*
fr3/s/half	0.6015[§]	0.2924[§]	-	-	0.7045	0.3488	0.7891	0.5981[†]	0.2739*	0.7697	0.3718	0.5917*	0.2834[†]	0.6543	0.3420
fr3/s/static	0.2850	0.1241	0.2735	0.1215	-	-	-	-	-	0.2749[§]	0.1192[†]	0.2654*	0.1183*	0.2658[†]	0.1199[§]
fr3/s/rpy	0.7772[§]	0.3999[§]	-	-	-	-	-	-	-	0.8303	0.4653	0.7410[†]	0.3665[†]	0.6957*	0.3405*
fr3/w/xyz	7.7846	5.8335	0.8266	0.5826	0.6284	0.3848	0.6368	0.6132[§]	0.3348*	0.6204	0.3850	0.6023*	0.3719[†]	0.6079[†]	0.3757[§]
fr3/w/half	7.2138	5.8299	0.8142	0.4101	0.7842[§]	0.4012	0.9650	0.7879	0.3751[§]	0.7853	0.3821	0.7575[†]	0.3743[†]	0.6995*	0.3350*
fr3/w/static	4.1856	3.8077	0.2690	0.1182	0.2612	0.1259	0.2872	0.3038	0.1437	0.2535[§]	0.1130[§]	0.2527[†]	0.1051[†]	0.2356*	0.1019*
fr3/w/rpy	8.8923	6.6658	3.0042	2.3065	0.9894*	0.5701[†]	1.0587	1.0841[§]	0.6668[§]	1.0521[†]	0.5577*	1.1084	0.6722	1.2773	0.8049

TABLE III: Comparisons of Rotational Drift in RPE between RGB-D SLAM techniques.



Fig. 6: Moving sequences of the Baxter robot (top) and the resulting reconstructed 3D scene with the estimated camera trajectory (down). The red line indicates the estimated camera trajectory with optimization, and orange frustums mark camera poses along the trajectory.

C. Qualitative Results

Fig. 5 shows the ATE along the camera trajectories for further assessment of qualitative results from our method (*first and third lines*) and CFP-SLAM (*second and fourth lines*). Fig. 5 shows the ground truth (black), the estimated trajectory (blue), and their difference (red). Overall, our estimated trajectories have better ground-truth coverage than CFP-SLAM in most cases, with improvements in noticeable regions, but only worse for the *fr3/w/ rpy* sequence.

D. Ablation Study

We use ORB-SLAM2 as a baseline where no moving objects and intra-objects are detected. The baseline code extracts features from images and generates spatial coordinates for mapping while adding keyframes and performing bundle adjustments using GPO. Without Hough Voting, which predicts and removes the moving objects, performance gains by 8% compared to the baseline but drops four times compared to full *V3D-SLAM*. Without intra-object moving, it does not consider deformable parts/objects but moving objects, resulting in a dropping performance of 6% compared to *V3D-SLAM*.

VI. REAL-ROBOT EXPERIMENTS & DEPLOYABILITY

A. Experimental Setup

The Intel RealSense D435i RGB-D camera is mounted on the display of the Baxter robot, and the scene of various indoor objects with bean bags, backpacks, chairs, books, and cups is set up, as shown in Fig. 6. The robot perceives objects within a 3-meter radius, the ideal range given by the camera, while arbitrarily moving to estimate its camera poses and trajectory and reconstruct the scene simultaneously with the acquired RGB-D stream.

B. Performance of Proposed Method

We deploy our method on the Intel NUC5i3RYH PC, which runs on its native onboard CPU without any dedicated

GPUs. The entire pipeline sufficiently achieves on-robot deployable performance. Fig. 6 shows the arbitrary movements of the Baxter robot to perceive the setup scene up front at the top row. Our arbitrary movements supported by four mecanum wheels on the integrated Dataspeed mobility base allow the robot to slide to the left and right side, move toward and backward, and rotate in a clockwise and anti-clockwise manner, closely mimicking the camera’s movements in the TUM RGB-D benchmark [8]. The qualitative results of the estimated camera trajectory in 3D and the 3D reconstructed scene are also shown in Fig. 6 at the bottom row with the red line indicating the estimated camera trajectory with optimization, and the camera frustums are marked in orange.

C. Demonstration

The demonstration video shows the deployability of our method on the Baxter mobile robot with an employed Intel RealSense D435i RGB-D camera and is made available at <https://youtu.be/K4RcKrASpqI>.

VII. CONCLUSIONS

In this paper, we present *V3D-SLAM*, a technique that reliably estimates and reconstructs the camera trajectory by removing noises induced by the dynamic nature of the environment. The dynamic objects are identified at the pixel level using state-of-the-art object segmentation on RGB images and refined using geometrical information in the 3D domain. Besides using Hough voting to identify moving objects in 3D, we also detect deformable objects using Chamfer distance to exclude their intra-object changes, which noise the feature extractor. To verify the robustness of our proposed method, we conduct experiments on the TUM RGB-D benchmark and compare our proposed method to recent state-of-the-art SLAM techniques. The experimental results show that our method mostly outperforms others regarding ATE and rotational and translational RPE metrics. Through deployment, our method enables the Baxter robot to perform RGB-D SLAM and its operations.

REFERENCES

- [1] C. Kerl, J. Sturm, and D. Cremers, "Dense visual slam for rgb-d cameras," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 2100–2106.
- [2] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "Orb-slam: a versatile and accurate monocular slam system," *IEEE transactions on robotics*, vol. 31, no. 5, pp. 1147–1163, 2015.
- [3] S. Thrun, D. Fox, W. Burgard, and F. Dellaert, "Robust monte carlo localization for mobile robots," *Artificial intelligence*, vol. 128, no. 1-2, pp. 99–141, 2001.
- [4] G. Grisetti, C. Stachniss, and W. Burgard, "Improved techniques for grid mapping with rao-blackwellized particle filters," *IEEE transactions on Robotics*, vol. 23, no. 1, pp. 34–46, 2007.
- [5] C. Yu, Z. Liu, X.-J. Liu, F. Xie, Y. Yang, Q. Wei, and Q. Fei, "Ds-slam: A semantic visual slam towards dynamic environments," in *2018 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2018, pp. 1168–1174.
- [6] T. Ji, C. Wang, and L. Xie, "Towards real-time semantic rgb-d slam in dynamic environments," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 11 175–11 181.
- [7] X. Hu, Y. Zhang, Z. Cao, R. Ma, Y. Wu, Z. Deng, and W. Sun, "Cfp-slam: A real-time visual slam based on coarse-to-fine probability in dynamic environments," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 4399–4406.
- [8] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, 2012, pp. 573–580.
- [9] J. Lamarca, S. Parashar, A. Bartoli, and J. M. M. Montiel, "DefSLAM: Tracking and Mapping of Deforming Scenes From Monocular Sequences," *IEEE Transactions on Robotics*, vol. 37, no. 1, pp. 291–303, Feb. 2021, conference Name: IEEE Transactions on Robotics. [Online]. Available: <https://ieeexplore.ieee.org/document/9201190/?arnumber=9201190>
- [10] M. A. Butt and P. Maragos, "Optimum design of chamfer distance transforms," *IEEE Transactions on Image Processing*, vol. 7, no. 10, pp. 1477–1484, 1998.
- [11] C. Tomasi and T. Kanade, "Shape and motion from image streams under orthography: a factorization method," *International journal of computer vision*, vol. 9, pp. 137–154, 1992.
- [12] R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE transactions on robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- [13] R. Scona, M. Jaimez, Y. R. Petillot, M. Fallon, and D. Cremers, "Staticfusion: Background reconstruction for dense rgb-d slam in dynamic environments," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 3849–3856.
- [14] S. Li and D. Lee, "Rgb-d slam in dynamic environments using static point weighting," *IEEE Robotics and Automation Letters*, vol. 2, no. 4, pp. 2263–2270, 2017.
- [15] Y. Sun, M. Liu, and M. Q.-H. Meng, "Improving rgb-d slam in dynamic environments: A motion removal approach," *Robotics and Autonomous Systems*, vol. 89, pp. 110–122, 2017.
- [16] D.-H. Kim and J.-H. Kim, "Effective background model-based rgb-d dense visual odometry in a dynamic environment," *IEEE Transactions on Robotics*, vol. 32, no. 6, pp. 1565–1573, 2016.
- [17] W. Dai, Y. Zhang, P. Li, Z. Fang, and S. Scherer, "Rgb-d slam in dynamic environments using point correlations," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 1, pp. 373–389, 2020.
- [18] F. Zhong, S. Wang, Z. Zhang, and Y. Wang, "Detect-slam: Making object detection and slam mutually beneficial," in *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 2018, pp. 1001–1010.
- [19] L. Xiao, J. Wang, X. Qiu, Z. Rong, and X. Zou, "Dynamic-slam: Semantic monocular visual localization and mapping based on deep learning in dynamic environment," *Robotics and Autonomous Systems*, vol. 117, pp. 1–16, 2019.
- [20] B. Bescos, J. M. Fácil, J. Civera, and J. Neira, "Dynaslam: Tracking, mapping, and inpainting in dynamic scenes," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 4076–4083, 2018.
- [21] X. Yuan and S. Chen, "Sad-slam: A visual slam based on semantic and depth information," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 4930–4935.
- [22] Y. Fan, Q. Zhang, Y. Tang, S. Liu, and H. Han, "Blitz-slam: A semantic slam in dynamic environments," *Pattern Recognition*, vol. 121, p. 108225, 2022.
- [23] Y. Wang, K. Xu, Y. Tian, and X. Ding, "Drg-slam: A semantic rgb-d slam using geometric features for indoor dynamic scene," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1352–1359.
- [24] K. Nguyen, T. Dang, and M. Huber, "Online 3d deformable object classification for mobile cobot manipulation," in *ISR Europe 2023; 56th International Symposium on Robotics*, 2023, pp. 139–146.
- [25] P. V. Hough, "Machine analysis of bubble chamber pictures," in *International Conference on High Energy Accelerators and Instrumentation, CERN, 1959*, 1959, pp. 554–556.
- [26] J. Gall and V. Lempitsky, "Class-specific hough forests for object detection," *Decision forests for computer vision and medical image analysis*, pp. 143–157, 2013.
- [27] M. Sun, G. Bradski, B.-X. Xu, and S. Savarese, "Depth-encoded hough voting for joint object detection and shape recovery," in *Computer Vision—ECCV 2010: 11th European Conference on Computer Vision, Heraklion, Crete, Greece, September 5-11, 2010, Proceedings, Part V 11*. Springer, 2010, pp. 658–671.
- [28] S. Maji and J. Malik, "Object detection using a max-margin hough transform," in *2009 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 2009, pp. 1038–1045.
- [29] Y. Xiang, T. Schmidt, V. Narayanan, and D. Fox, "Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes," *arXiv preprint arXiv:1711.00199*, 2017.
- [30] C. R. Qi, O. Litany, K. He, and L. J. Guibas, "Deep hough voting for 3d object detection in point clouds," in *proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019, pp. 9277–9286.
- [31] G. Jocher, A. Chaurasia, and J. Qiu, "YOLO by Ultralytics," Jan. 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
- [32] C.-Y. Wang and H.-Y. M. Liao, "YOLOv9: Learning what you want to learn using programmable gradient information," 2024.
- [33] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski, "Orb: An efficient alternative to sift or surf," in *2011 International conference on computer vision*. Ieee, 2011, pp. 2564–2571.
- [34] S. Teller, E. Olson, and J. Leonard, "Fast iterative optimization of pose graphs with poor initial estimates," in *Proc. of IEEE International Conference on Robotics and Automation*, 2006, pp. 2262–2269.
- [35] D. Scharstein and R. Szeliski, "A taxonomy and evaluation of dense two-frame stereo correspondence algorithms," *International journal of computer vision*, vol. 47, pp. 7–42, 2002.
- [36] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, *et al.*, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *kdd*, vol. 96, no. 34, 1996, pp. 226–231.
- [37] C. Campos, R. Elvira, J. J. G. Rodríguez, J. M. Montiel, and J. D. Tardós, "Orb-slam3: An accurate open-source library for visual, visual-inertial, and multimap slam," *IEEE Transactions on Robotics*, vol. 37, no. 6, pp. 1874–1890, 2021.