

FM-Fusion: Instance-aware Semantic Mapping Boosted by Vision-Language Foundation Models

Chuhao Liu¹, Ke Wang^{2,*}, Jieqi Shi¹, Zhijian Qiao¹ and Shaojie Shen¹

Abstract—Semantic mapping based on the supervised object detectors is sensitive to image distribution. In real-world environments, the object detection and segmentation performance can lead to a major drop, preventing the use of semantic mapping in a wider domain. On the other hand, the development of vision-language foundation models demonstrates a strong zero-shot transferability across data distribution. It provides an opportunity to construct generalizable instance-aware semantic maps. Hence, this work explores how to boost instance-aware semantic mapping from object detection generated from foundation models. We propose a probabilistic label fusion method to predict close-set semantic classes from open-set label measurements. An instance refinement module merges the over-segmented instances caused by inconsistent segmentation. We integrate all the modules into a unified semantic mapping system. Reading a sequence of RGB-D input, our work incrementally reconstructs an instance-aware semantic map. We evaluate the zero-shot performance of our method in ScanNet and SceneNN datasets. Our method achieves 40.3 mean average precision (mAP) on the ScanNet semantic instance segmentation task. It outperforms the traditional semantic mapping method significantly.

I. Introduction

Instance-aware semantic mapping in indoor environments is a key module for an autonomous system to achieve a higher level of intelligence. Based on the semantic map, a mobile robot can detect loop more robust [1] and efficiently [2]. The current methods rely on supervised object detectors like Mask R-CNN [3] to detect semantic instances and fuse them into an instance-level semantic map. However, the supervised object detectors are trained in specific data distribution and lack generalization ability. In deploying them in other real-world scenarios without fine-tune the networks, their performance is seriously degenerated. As a result, the reconstructed semantic map is also of poor quality in the target environment.

On the other hand, foundation models have been developing rapidly in vision-language modality [4] [5]. Multiple foundation models are combined to detect and segment objects. GroundingDINO [6], the latest State-of-the-Arts (SOTA) open-set object detection network, reads a text prompt and performs vision-language modal

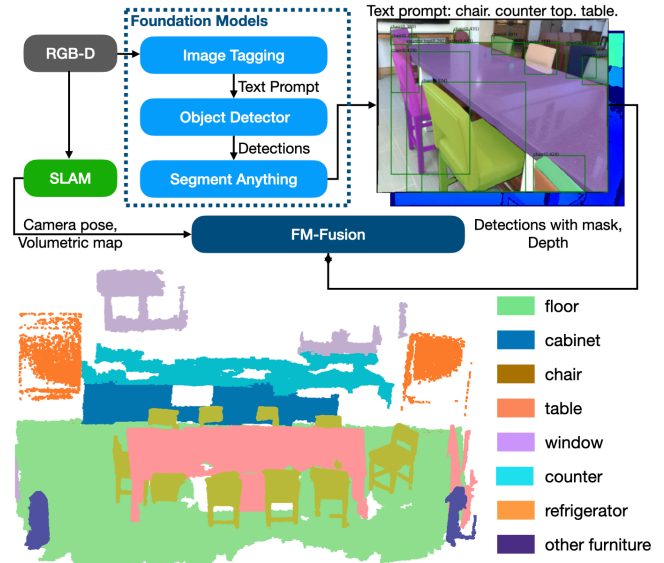


Fig. 1: Our system reads a sequence of RGB-D frames. The vision-language foundation models detect objects in open-set labels and high-quality masks. The SLAM modules generate a camera pose and a global volumetric map. Our method incrementally fuses the object detections from foundation models into an instance-aware semantic map. A reconstructed semantic map from ScanNet scene0011_01 is shown.

fusion. It detects objects with bounding boxes and open-set labels. The open-set labels are open vocabulary semantic classes. GroundingDINO has achieved 52.5 mAP on the zero-shot COCO object detection benchmark. It is higher than most of the supervised object detectors. Moreover, the image tagging model recognizes anything (RAM) [7] predicts semantic tags from an image. The tags can be encoded as a text prompt and sent to GroundingDINO. Vision foundation model segment anything (SAM) [4] generates precise zero-shot image segmentation results from geometric prompts, including a bounding box prompt. SAM can generate high-quality masks for detection results from GroundingDINO.

RAM, GroundingDINO, and SAM can be combined to detect objects in open-set labels and high-quality masks. All of these foundation models are trained using large-scale data and demonstrate strong zero-shot generalization ability in various image distributions. They provide a new approach for the autonomous system to reconstruct a generalizable instance-aware semantic map. This paper explores how to fuse object detection from

¹Authors are with the Department of Electronic and Computer Engineering, the Hong Kong University of Science and Technology, Hong Kong, China. {cliuci,jshias,zqiaoac}@connect.ust.hk, eeshaojie@ust.hk

²Author is with the Department of Information Engineering, Chang'an University, China. kwangdd@chd.edu.cn

* Corresponding author.

foundation models into an instance-aware semantic map.

To fuse object detection from foundation models, two challenges should be addressed. Firstly, the foundation models generate open-set tags or labels. However, the semantic mapping task requires each constructed instance to be classified in close-set semantic classes. A label fusion method is required to predict an instance’s semantic class from a sequence of observed open-set labels. Secondly, SAM is operating on a single image. In dense indoor environments, SAM frequently generates inconsistent instance masks at changed viewpoints. It results in over-segmented and noisy instance volumes. Refining instance volumes integrated from inconsistent instance segmentation results is the challenge. However, these challenges have not been considered in traditional semantic mapping works. If foundation models are directly used in a traditional semantic mapping system, they reconstruct semantic instances in a less satisfied quality.

To address such challenges, we propose a probabilistic label fusion method following the Bayes filter algorithm. Meanwhile, we refine the instance volume via merging over-segmentation and fuse instance volume with the global volumetric map. The label fusion and instance refinement modules are incrementally run in our system. As shown in Figure 1, reading a sequence of RGB-D frames, FM-Fusion fuses the detections from foundation models and runs simultaneously with a traditional SLAM system. Our main contributions are:

- An approach to fuse the object detections from vision-language foundation models into an instance-aware semantic map. The foundation models are used without fine-tune.
- A probabilistic label fusion method that predicts close-set semantic classes from open-set label measurements.
- Instances are refined to address inconsistent masks at changed viewpoints.
- The method is zero-shot evaluated in ScanNet [8]. It outperforms the traditional semantic mapping method significantly. We further evaluate it in SceneNN [9] to demonstrate its robustness in other image distributions.

II. Related Works

A. Vision-Language Foundation Models

The image tagging foundation model RAM [7], recognizes the semantic categories in the image and generates related tags. The open-set object detector, such as GLIP [10] and GroundingDINO [6], reads a text prompt to detect the objects. The text prompt can be a sentence or a series of semantic labels. It extracts the regional image embeddings and matches the image embedding to the phrase of the text prompt through a grounding scheme. The network is trained using contrastive learning to align the image embeddings and text embeddings. The detection results contain a bounding box and a set

of open-set label measurements. SAM [4] can precisely segment any object with a geometric prompt. It is trained with 11M images and evaluated in zero-shot benchmarks. SAM demonstrates strong generalization ability across data distribution without fine-tune. The combined foundation models read an image and detect objects with open-set labels and masks. We denote them as RAM-Grounded-SAM.

The foundation models have been applied in a series of downstream tasks without fine-tuning. Without semantic prediction, SAM3D [11] projects the image-wise segmentation from SAM to a 3D point cloud map. It further merges the segments with geometric segments generated from graph-based segmentation [12]. SAM is also combined with a neural radiance field to generate a novel view of objects [13]. On the other hand, combining the SAM or other foundation models with semantic mapping is still an open area.

B. Semantic Mapping

SemanticFusion [14] is a pioneer work in semantic mapping. It trains a lightweight CNN-based semantic segmentation network [15] on the NYUv2 dataset. SemanticFusion incrementally fuses the semantic labels, ignoring the instance-level information, into each surfel of the global volumetric map. In Bayesian fusing the label measurement, the semantic probability is directly provided by the object detector. Relying on a pre-trained Mask R-CNN on the COCO dataset, Kimera [16] uses similar methods to fuse semantic labels into a voxel map. It clusters the nearby voxels with identical semantic labels into instances. Kimera further constructs a scene graph, which is a hierarchical map representation. Based on Kimera, Hydra [2] utilizes the scene graph to detect loops more efficiently.

On the other hand, Fusion++ [17] directly detects semantic instances on images and fuses them into instance-wise volumetric maps. It further demonstrates that semantic landmarks can be used in loop detection. Later methods use similar methods to construct semantic instance maps but utilize the semantic landmarks in novel methods to detect loops [1] [18].

Rather than a pure dense map such as a surfel map or voxel map, Voxblox++ [19] first generates geometric segments on each depth frame [20]. If the object detection masks the complete region of an instance, it can merge those broken segments generated from geometric segmentation. Then, the merged segments with their labels are fused into a global segment map through a data association strategy.

The main limitation of the current semantic mapping methods is the lack of ability to generalize. The supervised object detection networks are trained with limited source data. Considering the majority of target SLAM scenarios do not provide annotated semantic data, object detection can not be fine-tuned on the target distribution. To avoid the issue of generalization,

Kimera has to experiment in a synthetic dataset [16], including some experiments that rely on ground-truth segmentation. Lin et al. [1] sets up an environment with sparsely distributed objects to reconstruct a semantic map. Voxblox++ evaluates a few of the 9 semantic classes in 10 scans. Although they propose novel semantic SLAM methods, the semantic mapping module prevents their methods from being used in other real-world scenes.

To enhance robustness in the distribution shift, our method fuses the object detections from foundation models to reconstruct the instance-aware semantic map. We evaluate its zero-shot performance on the ScanNet semantic instance segmentation benchmark. It involves 20 classes in the NYUv2 label-set and evaluates their average precision (AP) in 3D space. We also show the qualitative results in several SceneNN scans, which have been used by the previous semantic mapping works.

III. Fuse Multi-frame Detections

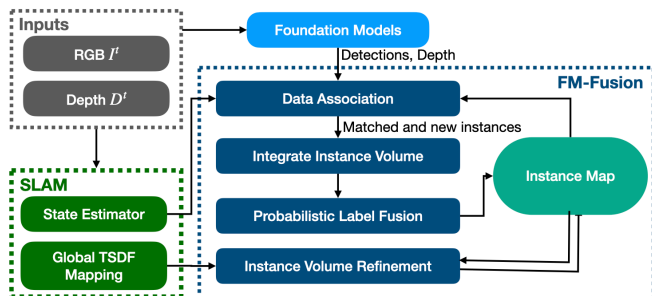


Fig. 2: System overview of FM-Fusion

A. Overview

As shown in Figure 2, FM-Fusion reads an RGB-D sequence and reconstructs a semantic instance map. Each semantic instance is represented as $\mathbf{f} = \{L_s, \mathbf{v}\}$, where L_s is its predicted semantic class and \mathbf{v} is its voxel grid map. L_s is predicted as a label c_n over the NYUv2 label-set \mathcal{L}_c . At each RGB-D frame $\{I^t, D^t\}$ at frame index t , RAM generates a set of possible object tags. The valid tags are encoded into the text prompt q^t . GroundingDINO generates object detections with each of the detection $z_k^t = \{y_i, s_i, q^t\}_i$, where y_i is the predicted open-set label, s_i is the corresponding similarity score and q^t is the frame-wise text prompt. For each z_k^t , SAM generates an object mask m_k .

B. Prepare the object detector

We first construct open-set labels of our interests \mathcal{L}_o . RAM generates various tags. Many of them are not correlated with the pre-defined labels \mathcal{L}_c . The labels of interest can be selected by sampling a histogram of measured labels for each semantic class in \mathbf{L}_c . In the ScanNet experiment, we select 38 open-set labels to construct \mathcal{L}_o . Only the tags belonging to \mathcal{L}_o are encoded into the q^t and sent to GroundingDINO. GroundingDINO matches each detected object with the tags in the text prompt. The tags in q^t and label measurements $\{y_i\}_i$ in each z_k^t are all from the label-set \mathcal{L}_o .

In a single image frame, RAM can miss some objects in its generated tags due to occlusion. The missing tags further cause GroundingDINO to detect objects incorrectly. It is a natural limitation of running foundation models on a single image. To address it, we encode the detected labels in adjacent frames into the text prompt. The augmented text prompt $q^t = \bar{q}^t \cup U^t$, where \bar{q}^t is the valid tags from RAM and U^t is a set of measured labels in previous adjacent frames. All the tags in \bar{q}^t and labels in U^t belong to the \mathcal{L}_o . The text prompt augmentation can reduce the missing tags generated from a single image. More complete tags improve the detection performance of GroundingDINO.

C. Data association and integration

In our system, each instance maintains an individual voxel grid map \mathbf{v} , similar to Fusion++ [17]. Meanwhile, the SLAM module integrates a global TSDF map [21] separately. The advantage of separating semantic mapping and global volumetric mapping is that false or missed object detection can not affect the global volumetric map. So, in each RGB-D frame, all the observed sub-volumes are integrated into the global TSDF map despite the detection variances.

In each detection frame, data association is conducted between detection results and volumes of the existing instances. Specifically, the observed instance voxels are first queried. They can be searched by projecting the depth image into the voxel grid map of all the instances. If an instance is observed, its voxels are projected to the current RGB frame. For a detection z_k^t and a projected instance \mathbf{f}_j , their intersection over union (IoU) can be calculated $\Omega(z_k^t, \mathbf{f}_j) = \frac{m_k \cap r_j}{m_k \cup r_j}$, where m_k is a detection mask and r_j is the projected mask of an existed instance. If $\Omega(z_k^t, \mathbf{f}_j)$ is larger than a threshold, the detection k is associated with instance j .

After data association, we integrate the voxel grid map of matched instances accordingly. Those unmatched detections initiate new instances. An instance voxel grid map \mathbf{v} is integrated using the traditional voxel map fusion method [21]. Specifically, we raycast the masked depth of a detected object and update all of its observed voxels.

D. Probabilistic label fusion

As shown in Figure 3, an object is observed by GroundingDINO across frames. Each generated detection result $z_k^t = \{y_i, s_i, q^t\}_i$ contains multiple label measurements y_i , the corresponding similarity score s_i and a text prompt q^t , where $y_i = o_m, o_m \in \mathcal{L}_o$. Based on the associated detections, we predict a probability distribution $p(L_s^t = c_n)$, where $c_n \in \mathcal{L}_c$ and t is the index of the image frame.

We follow the Bayes filter algorithm [22] to fuse open-set label measurements and propagate them along the image sequence. The input to the Bayesian label fusion is detection result z_k^t , semantic probability distribution

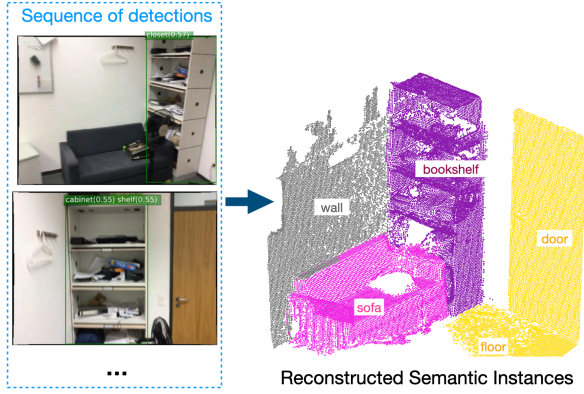


Fig. 3: GroundingDINO detects a bookshelf and generates multiple open-set label measurements across frames. Our label fusion module predicts its semantic class in NYUv2 label-set \mathcal{L}_c from label measurements in \mathcal{L}_o .

at the last frame $p(L_s^{t-1})$, and a uniform control input u^t . And it predicts the latest semantic probability distribution $p(L_s^t)$.

Algorithm 1: Bayes Filter for Label Fusion

Input: $p(L_s^{t-1})$, $z_k^t = \{y_i, s_i, q^t\}_{i \in [0:J]}$, $u^t = 1$

Output: $p(L_s^t)$

for $c_n \in \mathcal{L}_c$ do

Prediction:

$$\bar{p}(L_s^t = c_n) = p(L_s^t = c_n | L_s^{t-1} = c_n, u^t) p(L_s^{t-1} = c_n) \quad (1)$$

$$\bar{p}(L_s^t = c_n) = p(L_s^{t-1} = c_n) \quad (2)$$

Measurement Update:

$$p(L_s^t = c_n) = \eta p(z_k^t | L_s^t = c_n) \bar{p}(L_s^t = c_n) \quad (3)$$

$$p(L_s^t = c_n) = \eta \prod_{i=0}^{J-1} p(y_i, s_i, q^t | L_s^t = c_n) \bar{p}(L_s^t = c_n) \quad (4)$$

$$p(y_i, s_i, q^t | L_s^t = c_n) = p(s_i | y_i, q^t, L_s^t = c_n) p(y_i, q^t | L_s^t = c_n) \quad (5)$$

$$p(y_i, q^t | L_s^t = c_n) = p(y_i = o_m, \exists o_m \in q^t | L_s^t = c_n) \quad (6)$$

end

The key part in our Bayesian label fusion module is the likelihood function $p(y_i, s_i, q^t | L_s^t = c_n)$, as shown in equation (5). The score likelihood $p(s_i | y_i, q^t, L_s^t = c_n)$ is given by GroundingDINO, while label likelihood $p(y_i, q^t | L_s^t = c_n)$ should be statistic summarized. Since GroundingDINO can only detect a label y_i if it is given in the text prompt, the label likelihood can be transmitted as equation (6). $\exists o_m \in q^t$ denotes the detected label o_m exists in the text prompt q^t .

Here, we further expand the label likelihood in equation (6) into two conditional probabilities,

$$p(y_i = o_m, \exists o_m \in q^t | L_s^t = c_n) = p(y_i = o_m | \exists o_m \in q^t, L_s^t = c_n) p(\exists o_m \in q^t | L_s^t = c_n) \quad (7)$$

The first term is a detection likelihood while the second term is a tagging likelihood. They can be statistically summarized using the detection results from GroundingDINO and tagging results from RAM. We follow the equation (7) to construct a label likelihood matrix over $o_m \in \mathcal{L}_o$ and $c_n \in \mathcal{L}_c$. In the ScanNet training set, we sample 35,000 image frames with tagging results, detection results, and ground-truth annotation to summarize the statistics. In the Bayesian update step in equation (5), the label likelihood between each pair of $\{o_m, c_n\}$ can be queried from the constructed label likelihood matrix.




Fig. 4: The label likelihood matrix $p(y_i = o_m, \exists o_m \in q^t | L_s = c_n)$ summarized in ScanNet is shown on the left. Each column represents a specific true semantic class c_n , while each row represents a measured open-set label o_m . On the right, it is a manually assigned likelihood matrix.

As shown in Figure 4(a), parts of the constructed label likelihood matrix are visualized, while the complete likelihood matrix involves the entire \mathcal{L}_o and \mathcal{L}_c . For comparison, we construct a manually assigned label likelihood matrix similar to Kimera. As shown in Figure 4, the statistic summarized likelihood matrix is quite different from the manually assigned one. In the statistical label likelihood, each semantic class can be detected by its similar open-set labels at various probabilities. Those cells beyond the diagonal can also have likelihood values, indicating the probability of falsely measured labels. The summarized likelihood matrix following equation (7) describes the probability distribution of label measurements reasonably.

In actual implementation, the multiplicative measurement update in equation (3) frequently generates over-confident probability distribution, which is also reported in Fusion++ [17]. It causes $p(L_s^t)$ can be easily dominated by the latest measurement z_k^t even if previous label measurements are all different with z_k^t . As a result, in the measurement update, we propagate the probability distribution by weighted addition.

$$p(L_s^t) = \frac{p(z_k^t | L_s^t) + (t-1)\bar{p}(L_s^t)}{t} \quad (8)$$

Then, the predicted semantic class for each instance at frame t is $\arg \max_{c_n} p(L_s^t = c_n)$.




Fig. 5: An example of an inconsistent instance mask generated from SAM. In each of the three frames, different areas of the bed are segmented.

IV. Instance refinement

A. Merge over-segmentation

Although SAM has demonstrated promising segmentation on a single image, it generates inconsistent instance masks at changed viewpoints, as shown in Figure 5. The inconsistent masks prevent a correct data association between detections and observed instances. Those mismatched detections are initialized as new instances and cause over-segmentation, as shown in Figure 6(a).



Fig. 6: The visualization shows instance voxel grid map (a) before and (b) after the merge.

The inconsistent instance mask is a natural limitation for image-based segmentation networks, including SAM and Mask R-CNN. To address it, we utilize spatial overlap information to merge the over-segmentation. For a pair of instances $\{\mathbf{f}_a, \mathbf{f}_b\}$ at detection frame t , where \mathbf{f}_a is volumetric larger than \mathbf{f}_b , their semantic similarity $\sigma(\mathbf{f}_a, \mathbf{f}_b)$ and 3D IoU $\Omega(\mathbf{f}_a, \mathbf{f}_b)$ are calculated,

$$\sigma(\mathbf{f}_a, \mathbf{f}_b) = \bar{p}(L_s^t(a)) \cdot \bar{p}(L_s^t(b)) \quad (9)$$

$$\Omega(\mathbf{f}_a, \mathbf{f}_b) = \frac{\hat{\mathbf{v}}_a \cup \mathbf{v}_b}{\mathbf{v}_b} \quad (10)$$

where $\bar{p}^t(L_s(a))$ is the normalized semantic distribution, \mathbf{v} is an instance voxel grid map and $\hat{\mathbf{v}}$ is the inflated voxel grid map. The voxel inflation is designed to enhance 3D IoU for instances with sparse volume. It can be directly generated by scaling the length of each voxel in \mathbf{v} . If the semantic similarity and 3D IoU are both larger than the corresponding thresholds, \mathbf{f}_b is integrated into the voxel grid map of \mathbf{f}_a and further cleared from the instance map. As shown in Figure 6(b), over-segmented instances caused by inconsistent object masks are merged.

B. Instance-geometry fusion

The instance-wise voxel grid map can contain voxel outliers due to noisy depth images being integrated. On the other hand, the global TSDF map is a precise 3D geometric representation. It is because the global TSDF map integrates all the observed volumes in each RGB-D frame, while instance volume only integrates a masked




Fig. 7: Illustration of the instance-geometry fusion. Geometric points are extracted from the global map.

RGB-D frame if the corresponding instance is correctly detected. To filter voxel outliers, we fuse instance-wise voxel grid map \mathbf{v} with the point cloud \mathbf{P} extracted from the global TSDF map. As shown in Figure 7, those voxels in \mathbf{v} that are not occupied by any point in \mathbf{P} are outliers and have been removed. The fused voxel grid map represents the instance volume precisely.

V. Experiment

We chose the public dataset ScanNet and SceneNN to evaluate the semantic mapping quality. In ScanNet, 30 scans from its validation set are used. We evaluated its semantic instance segmentation by average precision (AP). In another experiment, we selected 5 scans from SceneNN and evaluated the generalization ability of our method. The SceneNN scans are also used by the previous method [19]. In all the experiments, camera poses are provided by the dataset.

We compared our method with Kimera and a self-implemented Fusion++. To enable Kimera to read open-set labels \mathcal{L}_o , we converted each label in \mathcal{L}_o to a semantic class in NYUv2 label-set \mathcal{L}_c . The hard association between $\{\mathcal{L}_o, \mathcal{L}_c\}$ are decided by an academic ChatGPT. Then, Kimera can reconstruct a point cloud with semantic labels. To further generate instance-aware point cloud, we employed the geometric segmentation method known as "Cluster-All" [23]. It clusters the nearby points with identical semantic labels into an instance. Cluster-All is applied as a post-processing step on the reconstructed semantic map from Kimera. Notice that Cluster-All is very similar to the post-processing module provided by Kimera. But we use Cluster-All for convenient implementation. Meanwhile, Fusion++ is implemented based on our system modules. Compared with the original Fusion++ method, the main difference is that our implemented version does not maintain a foreground probability for each voxel. Instead, we updated the voxel's weight and filter background voxels using their weights. In experiments with traditional object detection, we used a Mask R-CNN backbone with FPN101 image backbone. We evaluated a pre-trained Mask R-CNN and a fine-tuned Mask R-CNN. The pre-trained one is trained in COCO instance segmentation dataset, while we also fine-tuned it using ScanNet dataset.

In implementation, we utilized Open3D [24] toolbox to construct the global TSDF map and instance-wise voxel grid map. The global TSDF map is integrated for every RGB-D frame, while our method and all baselines run in

| Method | cab. | bed | cha. | sof. | tab. | door | win. | bkf. | pic. | cou. | desk | cur. | ref. | show.toi. | sink | bath. | oth. | mAP ₅₀ | |
|------------------|------|------|------|------|------|------|------|------|------|------|------|------|------|-----------|-------|-------|------|-------------------|------|
| M-CNN& Kimera | 0.0 | 6.4 | 10.0 | 25.1 | 17.3 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 24.6 | 0.0 | 10.4 | 4.3 | 0.0 | 0.0 | 5.4 |
| M-CNN& Fusion++ | 0.0 | 27.1 | 3.7 | 14.7 | 4.4 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 0.0 | 23.9 | 0.0 | 46.6 | 20.0 | 0.0 | 0.0 | 7.8 |
| M-CNN*& Kimera | 27.8 | 55.5 | 18.7 | 0.0 | 0.0 | 16.5 | 33.1 | 31.2 | 12.9 | 23.3 | 3.9 | 26.0 | 0.0 | 75.0 | 60.0 | 22.7 | 60.0 | 0.0 | 25.9 |
| M-CNN*& Fusion++ | 7.1 | 22.0 | 31.2 | 0.0 | 13.3 | 12.5 | 15.0 | 11.5 | 28.5 | 0.0 | 0.0 | 18.1 | 0.0 | 0.0 | 0.0 | 40.1 | 0.0 | 0.0 | 11.1 |
| G-SAM & Kimera | 32.5 | 21.0 | 16.8 | 54.5 | 21.7 | 26.0 | 31.5 | 45.3 | 21.9 | 8.7 | 3.9 | 24.8 | 29.6 | 0.0 | 50.0 | 9.4 | 46.2 | 0.0 | 24.7 |
| G-SAM & Ours | 4.6 | 46.2 | 49.2 | 39.6 | 37.3 | 19.5 | 12.0 | 50.4 | 44.0 | 3.8 | 8.9 | 15.5 | 66.7 | 82.2 | 100.0 | 41.5 | 75.0 | 30.7 | 40.3 |

TABLE I: Evaluate Kimera, the implemented Fusion++, and the proposed method on ScanNet using 30 validation scans. We report AP₅₀ at 50% IoU threshold for each semantic class. M-CNN denotes Mask R-CNN, M-CNN* is fine-tuned Mask R-CNN, while G-SAM refers to RAM-Grounded-SAM.



Fig. 8: The reconstructed instance map using RAM-Grounded-SAM in ScanNet scene0011, scene0435 and scene0633 (from top to bottom). The falsely predicted semantic classes in (a) and (b) are highlighted in red circles, while spatial conflicted semantics are in yellow. All semantic maps are colored following the NYUv2 color map and instances are colored randomly.

every 10 frames to integrate the detected instances. In all the experiments, the RGB-D images are in 640×480 dimension and the voxel length is set to be 1.5 cm. The experiment is run on an Intel-i7 computer with Nvidia RTX-3090 GPU in an offline fashion.

A. ScanNet Evaluation

In the instance segmentation benchmark, as shown in Table I, semantic mapping based on Mask R-CNN can only reconstruct a few of the semantic categories. It is because the pre-trained Mask R-CNN is trained using COCO label-set and those new semantic classes in NYUv2 label-set are predicted with 0 AP. Even for those predictable semantic classes, the pre-trained Mask R-CNN suffers from the issue of generalization and achieve low AP₅₀ scores. In experiment with fine-tune Mask R-CNN, although the mean AP is improved, they still reconstruct a few of semantic classes with 0 AP. We also notice that Kimera performs significantly better than the implemented Fusion++. We believe the difference

comes from their different map management methods. Unlike Kimera ignores the instance-wise segmentation, Fusion++ maintains instance-wise volumes and requires data association. But the fine-tuned Mask R-CNN still generate detections with noisy instance masks, causing a large amount of false data association. As a result, Fusion++ generates instances with too many over-segmentation and maintains a low AP score.

The results demonstrated that semantic mapping based on supervised object detection can be easily affected by image distribution, label-set distribution and annotation quality. On the other hand, boosted by pre-trained foundational models RAM-Grounded-SAM, both Kimera and our method reconstructed semantic instances in higher quality than semantic mapping methods based on the supervised object detection.

However, simply replacing object detectors with foundation models could not utilize the maximum potential of the foundation models. Compared with Kimera us-

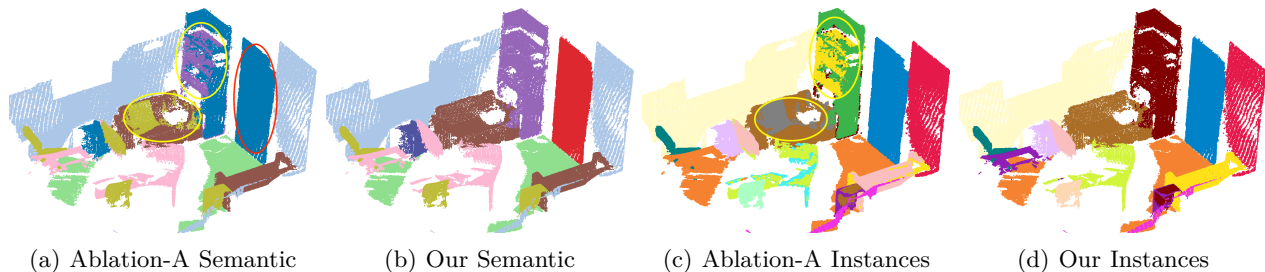


Fig. 9: A visualized example at ScanNet scene0025_01. (a) A falsely predicted instance caused by manually assigned likelihood is highlighted in a red circle, while the spatial conflicted semantic predictions are highlighted in yellow. (b) Proposed semantic result. (c) Over-segmentation is highlighted in yellow. (d) Our refined instances.

ing RAM-Grounded-SAM, our method achieved +15.6 mAP₅₀. The boosted performance comes from two aspects. Firstly, our probabilistic label fusion predicts semantic class in higher accuracy. As shown in Figure 8(a), Kimera predicts semantics of some sub-volumes falsely. Since Kimera updates the label measurements with a manually assigned likelihood probability and ignores the similarity score provided by GroundingDINO, it is easier to be affected by false label measurements. Secondly, Kimera ignores the instance-level segmentation and reconstructs many over-segmented instances. Some of them are predicted with different semantic labels, as shown in Figure 8(a) and 8(c). However, our method is instance-aware. Each instance volume is maintained separately. Our instance refinement module merges over-segmented instances caused by inconsistent instance masks at changed viewpoints. We further fused instance volume with a global volumetric map. Hence, our instances volumes are spatially consistent and relative precise, as shown in Figure 8(d).

| Method | Prm. Aug. | Likelihood | Refine | mAP ₅₀ |
|--------|-----------|---------------|--------|-------------------|
| A | ✓ | Manual Assign | ✓ | 35.9 |
| B | × | statistic | ✓ | 34.1 |
| C | ✓ | statistic | × | 23.4 |
| Ours | ✓ | statistic | ✓ | 40.3 |

TABLE II: Ablation study of FM-Fusion. Prm. Aug. denotes text prompt augmentation.

The rest of the ScanNet experiment focus on evaluating each module of our method through an ablation study. As shown in Table II, the text prompt augmentation, probabilistic label fusion with statistic summarized likelihood, and instance refinement all improve the reconstructed semantic instances.

A visualized example of the Ablation-A is shown in Figure 9. As shown in Figure 9(a), Ablation-A predicts an instance falsely, similar to Kimera. It also predicts overlapped instances with over-confident semantic probability distributions. They can not be merged during refinement due to their low semantic similarity. So, the over-segmented instances can not be merged, as shown in Figure 9(c). On the other hand, our method predicts the corresponding semantic classes correctly. The over-segmented instances are predicted with a similar semantic probability distribution and have been merged successfully, as shown in Figure 9(b) and 9(d).

To sum up, simply replacing traditional object detectors with RAM-Grounded-SAM to construct the semantic map improves the semantic mapping performance significantly. However, false label measurements, inconsistent instance masks, and missed tags in the text prompt still exist in foundation models. They limit the performance of semantic mapping. We consider those limitations of foundation models. Compared with Kimera using RAM-Grounded-SAM, our method further improves mAP₅₀ by +15.6.

B. SceneNN evaluation

In the SceneNN experiment, we kept using the label likelihood matrix $P(y_i = o_m, \exists o_m \in q^t | L_s = c_n)$ summarized in ScanNet and compare it with Kimera.

As shown in Figure 10, Kimera reconstructed some instances with false labels and over-segmentation, similar to its reconstruction in ScanNet. On the other hand, our semantic prediction is more accurate and significantly less over-segmentation. The quantitative results can be found in Table III. Although our statistical label likelihood is summarized using ScanNet data, we have not observed a domain gap in implementing it in SceneNN. One of the reasons is that foundation models preserve strong generalization ability. RAM-Grounded-SAM maintains a similar label likelihood matrix across the image distribution. For example, a door is frequently detected as a cabinet in both ScanNet and SceneNN datasets, which are highlighted in red in Figure 9(a) and Figure 10(a). Hence, our statistical label likelihood can be used across domains.

| | 096 | 206 | 223 | 231 | 255 | All |
|--------|------|------|------|------|------|------|
| Kimera | 52.0 | 29.9 | 28.6 | 34.0 | 37.5 | 41.1 |
| Ours | 63.1 | 69.7 | 37.5 | 38.8 | 25.0 | 49.7 |

TABLE III: SceneNN Quantitative results (mAP₂₅).

C. Efficiency

| | Base | Scaling |
|-------------------|---------------|----------|
| Foundation Models | RAM | 28.5 ms |
| | GroundingDINO | 120.7 ms |
| | SAM | 464.4 ms |
| FM-Fusion | Projection | 307ms |
| | Data Assoc. | 47.1ms |
| | Integration | 71.9ms |
| Total | 1039.6 ms | - |

TABLE IV: Runtime analysis for each frame in ScanNet.

So far, the system run offline. As shown in Table IV, the total time for each frame is 1039.6 ms. Although




Fig. 10: Reconstructions in SceneNN 096. False semantic and over-segmented instances are highlighted in red circles.

it is not a real-time system yet, many modules can be optimized in the future. SAM-related variants have been published to generate instance masks faster [25]. In FM-Fusion, a few modules are implemented with Python, the efficiency can be further improved by deploying it with C++. That would be one of our future works.

VI. Conclusion

In this work, we explored how to boost instance-aware semantic mapping with zero-shot foundation models. With foundation models, objects are detected in open-set semantic labels at various probabilities. The object masks generated at changed viewpoints are inconsistent and cause over-segmentation. The current semantic mapping methods have not considered such challenges. On the other hand, our method uses a Bayesian label fusion module with statistic summarized likelihood and refines the instance volumes simultaneously. Compared with the baselines, our method performs significantly better in ScanNet and SceneNN benchmarks.

References

- [1] S. Lin, J. Wang, M. Xu, H. Zhao, and Z. Chen, "Topology Aware Object-Level Semantic Mapping towards More Robust Loop Closure," *IEEE Robotics and Automation Letters (RA-L)*, vol. 6, pp. 7041–7048, 2021.
- [2] N. Hughes, Y. Chang, and L. Carlone, "Hydra: A real-time spatial perception system for 3D scene graph construction and optimization," in *Proc. of Robotics: Science and System (RSS)*, 2022.
- [3] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proc. of the IEEE intl. Conf. on Comp. Vis. (ICCV)*, 2017, pp. 2961–2969.
- [4] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, "Segment anything," *arXiv:2304.02643*, 2023.
- [5] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark et al., "Learning transferable visual models from natural language supervision," in *International conference on machine learning (ICML)*. PMLR, 2021, pp. 8748–8763.
- [6] S. Liu, Z. Zeng, T. Ren, F. Li, H. Zhang, J. Yang, C. Li, J. Yang, H. Su, J. Zhu et al., "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," *arXiv preprint:2303.05499*, 2023.
- [7] Y. Zhang, X. Huang, J. Ma, Z. Li, Z. Luo, Y. Xie, Y. Qin, T. Luo, Y. Li, S. Liu et al., "Recognize anything: A strong image tagging model," *arXiv preprint:2306.03514*, 2023.
- [8] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3d reconstructions of indoor scenes," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [9] B.-S. Hua, Q.-H. Pham, D. T. Nguyen, M.-K. Tran, L.-F. Yu, and S.-K. Yeung, "Scenenn: A scene meshes dataset with annotations," in *Proc. of the International Conference on 3D Vision (3DV)*, 2016, pp. 92–101.
- [10] L. H. Li, P. Zhang, H. Zhang, J. Yang, C. Li, Y. Zhong, L. Wang, L. Yuan, L. Zhang, J.-N. Hwang et al., "Grounded language-image pre-training," in *Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 10965–10975.
- [11] D. Zhang, D. Liang, H. Yang, Z. Zou, X. Ye, Z. Liu, and X. Bai, "Sam3d: Zero-shot 3d object detection via segment anything model," *arXiv preprint:2306.02245*, 2023.
- [12] P. F. Felzenszwalb and D. P. Huttenlocher, "Efficient graph-based image segmentation," *International journal of computer vision (IJCV)*, vol. 59, no. 2, pp. 167–181, 2004.
- [13] Q. Shen, X. Yang, and X. Wang, "Anything-3d: Towards single-view anything reconstruction in the wild," *arXiv preprint:2304.10261*, 2023.
- [14] J. McCormac, A. Handa, A. Davison, and S. Leutenegger, "Semanticfusion: Dense 3d semantic mapping with convolutional neural networks," in *Proc. of the IEEE Intl. Conf. on Robot. and Autom. (ICRA)*. IEEE, 2017, pp. 4628–4635.
- [15] H. Noh, S. Hong, and B. Han, "Learning deconvolution network for semantic segmentation," in *Proc. of the IEEE intl. Conf. on Comp. Vis. (ICCV)*, 2015, pp. 1520–1528.
- [16] 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 Robot. and Autom. (ICRA)*, 2020.
- [17] J. McCormac, R. Clark, M. Bloesch, A. J. Davison, and S. Leutenegger, "Fusion++: Volumetric object-level SLAM," *Proc. of the International Conference on 3D Vision (3DV)*, pp. 32–41, 2018.
- [18] J. Yu and S. Shen, "Semanticloop: loop closure with 3d semantic graph matching," *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 2, pp. 568–575, 2022.
- [19] M. Grinvald, F. Furrer, T. Novkovic, J. J. Chung, C. Cadena, R. Siegwart, and J. Nieto, "Volumetric instance-aware semantic mapping and 3D object discovery," *IEEE Robotics and Automation Letters (RA-L)*, vol. 4, no. 3, pp. 3037–3044, 2019.
- [20] F. Furrer, T. Novkovic, M. Fehr, A. Gawel, M. Grinvald, T. Sattler, R. Siegwart, and J. Nieto, "Incremental object database: Building 3d models from multiple partial observations," in *Proc. of the IEEE/RSJ Intl. Conf. on Intell. Robots and Syst.(IROS)*. IEEE, 2018, pp. 6835–6842.
- [21] H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, and J. Nieto, "Voxblox: Incremental 3D Euclidean Signed Distance Fields for on-board MAV planning," *Proc. of the IEEE/RSJ Intl. Conf. on Intell. Robots and Syst.(IROS)*, vol. 2017-Sep, pp. 1366–1373, 2017.
- [22] S. Thrun, "Probabilistic robotics," *Communications of the ACM*, vol. 45, no. 3, pp. 52–57, 2002.
- [23] B. Douillard, J. Underwood, N. Kuntz, V. Vlaskine, A. Quadros, P. Morton, and A. Frenkel, "On the segmentation of 3d lidar point clouds," in *Proc. of the IEEE Intl. Conf. on Robot. and Autom. (ICRA)*. IEEE, 2011, pp. 2798–2805.
- [24] Q.-Y. Zhou, J. Park, and V. Koltun, "Open3D: A modern library for 3D data processing," *arXiv:1801.09847*, 2018.
- [25] X. Zhao, W. Ding, Y. An, Y. Du, T. Yu, M. Li, M. Tang, and J. Wang, "Fast segment anything," *arXiv preprint:2306.12156*, 2023.