arXiv:2305.02743v2 [cs.CV] 6 May 2023

Incremental 3D Semantic Scene Graph Prediction from RGB Sequences

Shun-Cheng Wu[1] Keisuke Tateno[2] Nassir Navab[1] Federico Tombari[1,2]
[1]Technische Universität München [2]Google

Abstract

3D semantic scene graphs are a powerful holistic representation as they describe the individual objects and depict the relation between them. They are compact high-level graphs that enable many tasks requiring scene reasoning. In real-world settings, existing 3D estimation methods produce robust predictions that mostly rely on dense inputs. In this work, we propose a real-time framework that incrementally builds a consistent 3D semantic scene graph of a scene given an RGB image sequence. Our method consists of a novel incremental entity estimation pipeline and a scene graph prediction network. The proposed pipeline simultaneously reconstructs a sparse point map and fuses entity estimation from the input images. The proposed network estimates 3D semantic scene graphs with iterative message passing using multi-view and geometric features extracted from the scene entities. Extensive experiments on the 3RScan dataset show the effectiveness of the proposed method in this challenging task, outperforming state-of-the-art approaches. Our implementation is available at https://shunchengwu.github.io/MonoSSG.

1. Introduction

Scene understanding is a cornerstone in many computer vision applications requiring perception, interaction, and manipulation, such as robotics, AR/VR and autonomous systems [17, 54–56]. Semantic Scene Graphs (SSGs) go beyond recognizing individual entities (objects and stuff) by reasoning about the relationships among them [61, 66]. They also proved to be a valuable representation for complex scene understanding tasks, such as image captioning [26, 67], generation [13, 24], scene manipulation [10, 11], task planning [27], and surgical procedure estimation [42, 43]. Given the benefits of such representations, scene graph estimation received increasing attention in the computer vision community.

While earlier methods mainly estimate SSGs from images [18, 19, 33, 66, 72], recent approaches have also investigated estimating them from 3D data. Compared to 2D scene graphs, which describe a single image, 3D scene graphs depict the entire 3D scenes, enabling applications requiring a holistic understanding of the whole scene, such as path planning [47], camera localization, and loop closure detection [23]. However, existing 3D methods either require dense 3D geometry of the scenes to estimate 3D scene graphs [1, 23, 61, 64], which limits the use case since dense geometry is not always available, or constraints the scene graph estimation at the image-level [15, 27, 66], which tend to fail inferring relationships among objects beyond the individual viewpoints. A method that estimates 3D scene graphs relies on sparse scene geometry and reasoning about relationships globally has not been explored yet.

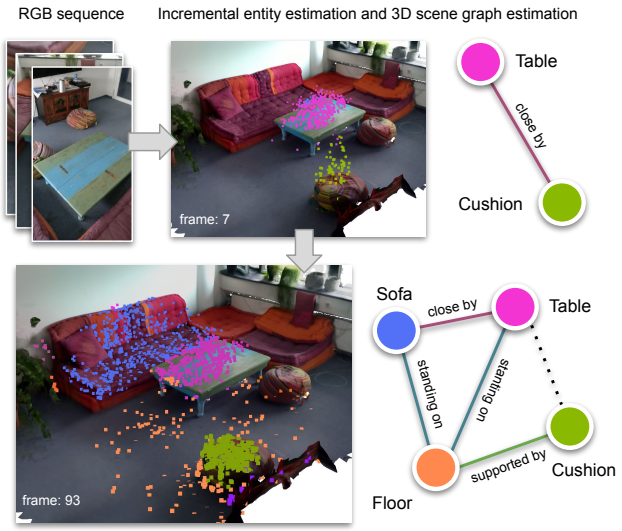

Figure 1. We propose a real-time 3D semantic scene graph estimation method that relies on an abstract understanding of a scene geometry built with RGB input. Our method estimates scene graphs incrementally by continuously estimating scene graphs and fusing local predictions into a global 3D scene graph.

In this work, we propose a real-time framework that incrementally estimates a global 3D SSG of a scene simply requiring an RGB sequence as input. The process is illustrated in Fig. 1. Our method simultaneously reconstructs a segmented point cloud while estimating the SSGs of the current map. The estimations are bound to the point map, which allows us to fuse them into a consistent global scene

graph. The segmented map is constructed by fusing entity estimation from images to the points estimated from a sparse Simultaneous Localization and Mapping (SLAM) method [3]. Our network takes the entities and other properties extracted from the segmented map to estimate 3D scene graphs. Fusing entities across frames is non-trivial. Existing methods often rely on dense inputs [38, 58] and struggle with sparse inputs since the points are not uniformly distributed. Estimating scene graphs with sparse input points is also challenging. Sparse and ambiguous geometry renders the node representations unreliable. On the other hand, directly estimating scene graphs from 2D images ignores the relationship beyond visible viewpoints. We aim to overcome the aforementioned issues by proposing two novel approaches. First, we propose a confidence-based fusion scheme which is robust to variations in the point distribution. Second, we present a scene graph prediction network that mainly relies on multi-view images as the node feature representation. Our approach overcomes the need for exact 3D geometry and is able to estimate relationships without view constraints. In addition, our network is flexible and generalizable as it works not only with sparse inputs but also with dense geometry.

We comprehensively evaluate our method on the 3D SSG estimation task from the public 3RScan dataset [60]. We experiment and compare with three input types, as well as 2D and 3D approaches. Moreover, we provide a detailed ablation study on the proposed network. The results show that our method outperforms all existing approaches by a significant margin. The main contributions of this work can be summarized as follows: (1) We propose the first incremental 3D scene graph prediction method using only RGB images. (2) We introduce an entity label association method that works on sparse point maps. (3) We propose a novel network architecture that generalizes with different input types and outperforms all existing methods.

2. Related Work

2.1. 3D Object Localization from Images

Localizing 3D objects from images aims to predict the position and orientation of objects. Existing methods can be broadly divided into two categories: without and with explicit geometrical reasoning.

In the former category, many works focus on estimating 3D bounding boxes by extending 2D detectors with learned priors [28, 37, 41, 70]. When sequential input is available, single view estimations can be fused to estimate a consistent object map [2, 22, 29, 30]. However, the fused results may not fulfill the multi-view geometric constraints. Multi-view approaches estimate oriented 3D bounding boxes from the given 2D detection of views. They mainly focus on minimizing the discrepancies between the projected 3D representation and the detected 2D bounding boxes.

In the latter category, 3D objects are localized with the help of explicit geometric information. Many existing methods treat object detection as spatial landmarks in a map [21, 35, 40, 53, 65, 69], also known as object-level SLAM. Others focus on fusing dense per-pixel predictions to a reconstructed map [6, 16, 36, 38, 44, 71], which is known as semantic mapping or semantic SLAM.

A major difference between object-level and semantic SLAM is that the former focuses only on foreground objects, while the latter also considers the structural and background information. Specifically, SemanticFusion [36] fuses dense semantic segments from images to a consistent dense 3D map with Bayesian updates. Its map representation provides a dense semantic understanding of a scene ignoring individual instances. PanopticFusion [38] proposes to combine the instance and semantic segmentation from images to a panoptic map. Their approach considers foreground object instances and non-instance semantic information from the background. SceneGraphFusion [64] relies on 3D geometric segmentation [59] and scene graph reasoning to achieve instance understanding of all entities in a scene.

One significant difficulty in instance estimation for semantic SLAM is associating the instances across frames. Existing approaches mainly rely on a dense map to associate predictions by calculating the intersection-over-union (IoU) or the overlapping ratio between the input and the rendered image from the map. However, these methods produce sub-optimal results when the map representation is sparse due to the non-uniform distribution of the map points. We overcome this problem by proposing a confidence-based association.

2.2. 3D Semantic Scene Graph

Estimating 3D scene graphs methods can be divided according to various criteria. From the perspective of the scene graph structure, some methods focus on hierarchical scene graphs [1, 23, 48, 50]. These approaches mainly address the problem of relationship estimation between entities from different hierarchical levels, *e.g.* object, and room level. The other methods focuses on pairwise relationships, *e.g.* support and comparative relationships, between nodes within a scene [27, 61, 64, 66]. From the perspective of input data, some previous methods rely on RGB input [15, 27] by fusing 2D scene graph predictions to a consistent 3D map. On the other hands, other methods rely on 3D input [1, 23, 48, 50, 61, 64] by using known 3D geometries. Nevertheless, most of the existing methods estimate scene graphs offline [1, 15, 48, 49], while a few works [23, 27, 64] predict scene graph in real-time.

Among all existing work, the pioneering work in 3D scene graph estimation is proposed by [15]. The authors extend the 2D scene graph estimation method from [66] with

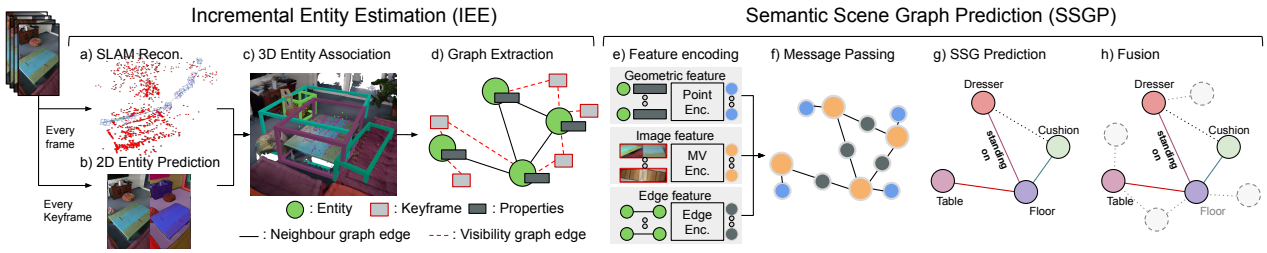

Figure 2. Given a sequence of RGB images, we use every frame to reconstruct a sparse point map (a) and every keyframe to estimate the 2D entities (b). (a) and (b) are associated and merged into a single 3D map (c). We asynchronously extract graph properties from the entity map (d) to estimate SSG. Our network computes the geometric, multi-view, and edge features (e). These features are propagated to each other with message passing (f), then used to predict a SSG (g). Then, periodic SSGs are fused to a global 3D SSG (h).

temporal consistency across frames and use geometric features from ellipsoids. Kim *et al*. [27] propose an incremental framework to estimate 3D scene graphs from 2D estimations. Armeni *et al*. [1] are the first work to estimate 3D scene graphs through the hierarchical understanding of the scene. Rosinol *et al*. [48] build on top of [1] to capture moving agents. This work is subsequently extended to a SLAM system [50]. Wald *et al*. [61] propose the first 3D scene graph method based on the relationship between objects at the same level, along with 3RScan: a richly annotated 3D scene graph dataset. Wu *et al*. [64] extend [61] to real-time scene graph estimation with a novel feature aggregation mechanism. Hughes *et al*. [23] propose to reconstruct a 3D hierarchical scene graph in real-time incrementally. Our method incrementally estimates a flat scene graph with multi-view RGB input and a sparse 3D geometry, which differentiates our work from the previous methods relying on 3D input [61, 64], and approaches without geometric understanding [15, 27].

3. Method

The proposed framework is illustrated in Fig. 2, which shows how, given a sequence of RGB images, it can estimate a 3D semantic scene graph incrementally. The Incremental Entity Estimation (IEE) front end makes use of the images to generate segmented sparse points. Those are merged into 3D entities and used to generate both an entity visibility graph and a neighbour graph. The Semantic Scene Graph Prediction (SSGP) network uses the entities and both graphs to estimate multiple scene graphs and then fuse them into a consistent 3D SSG.

We define a SSG as $\mathcal{G}_s = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ and $\mathcal{E}$ denote a set of entity nodes and directed edges. Each node $v_i \in \mathcal{V}$ is assigned an entity label $l_i \in L$, a set of points $\mathcal{P}_i$, an Oriented Bounding Box (OBB) b_i and a node category $c_i^{\text{node}} \in \mathcal{C}^{\text{node}}$. Each edge $e_{i \to j} \in \mathcal{E}$, connecting node v_i to v_j where $i \neq j$, consists of an edge category $c_{i \to j}^{\text{edge}} \in \mathcal{C}^{\text{edge}}$. L, $\mathcal{C}^{\text{node}}$, and $\mathcal{C}^{\text{edge}}$ denote all entity labels, a node category set, and an edge category set, respectively. An OBB b_i is a gravity-aligned 3D bounding box consisting of a boundary dimension $\mathbf{b}_i \in \mathbb{R}^3$, a center $\mathbf{o}_i \in \mathbb{R}^3$, and an angle that encodes the rotation along the gravity axis. The OBBs are used to build both graphs and features. The entity visibility graph models the visibility relationship of the entities as a bipartite graph $\mathcal{G}_c = (\mathcal{V}, \mathcal{K}, \mathcal{E}_c)$ where $\mathcal{K}$, $\mathcal{E}_c$ denote a set of keyframes and visibility edges, respectively. $\mathcal{G}_c$ gives the knowledge of the visibility of entity nodes in keyframes, which is used in computing multi-view visual features in SSGP. The neighbour graph encodes the proximity relationship of the entities as an undirected graph $\mathcal{G}_p = (\mathcal{V}, \mathcal{E}_p)$, where $\mathcal{E}_p$ is the set of proximity edges. The neighbour graph also serves as the initial graph for the message propagation step in SSGP.

3.1. Incremental Entity Estimation

During the first step of the IEE front end pipeline, a set of labeled 3D points are estimated from the sequence of RGB images (Sec. 3.1.1). The entity labels are determined using an entity segmentation method on selected keyframes (Sec. 3.1.2). Then, they are associated and fused into a sparse point map (Sec. 3.1.3). Finally, the entities and their properties are extracted using the labeled 3D points (Sec. 3.1.4).

3.1.1 Sparse Point Mapping

We use ORB-SLAM3 [3] to simultaneously estimate the camera poses and build a sparse point map by matching estimated keypoints from sequential RGB frames. To guarantee real-time performance, an independent thread is used to run the local mapping process using the stored keyframes. The same thread additionally takes care of running the entity detector and performing the label mapping process. For each point $\mathbf{p}_m \in \mathbf{P}$ in the map, we store its 3D coordinates, an entity label l_m, and its confidence score $w_m \in \mathbb{R}_{\geq 0}$.

3.1.2 2D Entity Detection

We estimate an entity label mask $\tilde{M}_t(u) \in \tilde{L}_t$ and a confidence mask $\tilde{W}_t(u) \in [0,1] \subset \mathbb{R}$ with every given keyframe $k_t \in \mathcal{K}$, where $u \in \mathbb{R}^2$ denotes the image coordinates and $\tilde{L}_t$ all entity labels in k_t. Both masks are estimated using a class-agnostic segmentation network which further improves other instance segmentation methods [4, 7, 31] by enabling the discovery of unseen entities [14, 25, 46]. Although segmentation networks provide accurate masks, the estimations are independent across frames. Thus, a label association stage is required to build a consistent label map.

3.1.3 Label Association and Fusion

Inspired by [35, 38, 51, 58], we use the reference map approach to handle the label inconsistency. It relies on a map reconstruction to solve label consistency by comparing input label mask to rendered mask. Then fuse the associated mask to the global point map.

Label Association. We start by building a reference entity mask $M_t(u) \in L$ by projecting point entity labels from the sparse point map using the pose of k_t. The consistency-resolved entity mask $\hat{M}'_t(u)$ is estimated by evaluating the corresponding labels on the image mask $\tilde{M}_t(u)$ and the reference mask $M_t(u)$. This evaluation can be performed by different methods such as using intersection over union statistics [38] or the maximum overlapping ratio between label masks [58]. However, both methods assume that points are uniformly distributed; a premise that fails in most sparse point reconstruction tasks. In such cases, these methods become unstable, as shown in the example provided in the supplementary material. To overcome this problem, we propose to use the maximum mean confidence as the criteria to find the best candidate. First, a confidence mask $W_t(u)$ is built by projecting the point label confidence using the pose of k_t, then the mean confidence score of a label $\tilde{l} \in \tilde{M}_t(u)$ and a reference label $l \in M_t(u)$ is computed by

$$\bar{\mathbf{M}}(\tilde{l}, l) = \frac{\sum_{u' \in \mathbf{\Pi}(\tilde{l}, l)} W_t(u')}{\#\left(\mathbf{\Pi}(\tilde{l}, l)\right)}, \tag{1}$$

where $\mathbf{\Pi}(\tilde{l}, l)$ gives a set of image coordinates $u' \in \mathbb{R}^2$ where $\tilde{l}$ and l overlap: $\{u' \mid (\tilde{M}_t(u') = \tilde{l}) \wedge (M_t(u') = l)\}$, and $\#(\cdot)$ is the cardinality operator. Then, the mask $\hat{M}'_t(u)$ is generated by replacing the per-pixel entity label $\tilde{l} \in \tilde{M}_t(u)$ with either a reference label l or a new label $l_{\text{new}} \notin L$ depending on:

$$\tilde{l} = \begin{cases} \arg\max_l \bar{\mathbf{M}}(\tilde{l}, l) & \text{if } \max_l \frac{\#\left(\mathbf{\Pi}(\tilde{l}, l)\right)}{\#(M_t(u) = l)} > \tau \\ l_{\text{new}} & \text{otherwise} \end{cases}, \tag{2}$$

where we filter out a match if the number of overlapped pixel has a low coverage over the total number of label l on the reference mask with a threshold τ. In addition, similar to [38], a reference entity label is assigned to only one input entity label. If a reference label has been assigned, we use descending order to search for the next best candidate.

Label Fusion. After the association process, the associated entity labels $\hat{M}'_t(u)$ are fused to the sparse point map $\mathbf{P}$. Since each label on $\hat{M}'_t(u)$ sources from a map point, the label and confidence value of a point are updated by

$$w_{\psi(u)} = \begin{cases} w_{\psi(u)} + \tilde{W}_t(u) & \text{if } M_t(u) = \hat{M}'_t(u) \\ w_{\psi(u)} - \tilde{W}_t(u) & \text{otherwise} \end{cases}, \tag{3}$$

where $\psi(u)$ is the corresponding point index that is projected on the pixel location u on both $M_t(u)$ and $W_t(u)$. In particular, when $w_{\psi(u)} < 0$, we set the entity label $l_{\psi(u)}$ to $\hat{M}'_t(u)$, and the weight $w_{\psi(u)}$ to $\tilde{W}_t(u)$.

3.1.4 Extraction

We use the points belonging to each entity label to compute the 3D OBB b_i of an entity $v_i \in \mathcal{V}$. We perform statistical outlier removal (from PCL [52]) to filter out points that could lead to distorted boxes. For the computation, we make use of the minimum volume estimation method [5] assuming gravity alignment.

The entity visibility graph $\mathcal{G}_c = (\mathcal{V}, \mathcal{K}, \mathcal{E}_c)$ consists of all nodes $\mathcal{V}$ and keyframes $\mathcal{K}$ connected by visibility edges $\mathcal{E}_c$. A visibility edge $e_{ij} \in \mathcal{E}_c$ exists if entity $v_i \in \mathcal{V}$ is visible in keyframe $k_j \in \mathcal{K}$. Specifically, the visibility is determined by checking if any point in node v_i is visible at k_j.

The neighbour graph $\mathcal{G}_p = (\mathcal{V}, \mathcal{E}_p)$ consists of nodes $\mathcal{V}$ and its proximity edges $\mathcal{E}_p$. A proximity edge $e_{i \to j} \in \mathcal{E}_p \mid v_i, v_j \in \mathcal{V}, i \neq j$ exists if nodes v_i, v_j are close in space, which is determined using a bounding box collision detection method. Since the size of the OBBs is not precise, we extend their dimensions by a margin τ^C to include additional potential neighbours.

3.2. Semantic Scene Graph Prediction

For every of the scene extractions obtained by the IEE front end, SSGP estimates 3D semantic scene graphs using message passing to jointly update initial feature representations and relationships [15, 61, 64, 66]. In the last step, the network fuses all of them into a consistent global 3D SSG. The initial node features are computed with multi-view image features (Sec. 3.2.1), while the initial edge feature is computed with the relative geometric properties of its connected two nodes (Sec. 3.2.2). Both initial features are jointly updated with a GNN along the connectivity given by the neighbour graph (Sec. 3.2.3). The updated node and edge are used to estimate their class distribution (Sec. 3.2.4). We apply a temporal scene graph fusion procedure to combine the predictions into a global 3D SSG (Sec. 3.2.5).

Our network architecture combines the benefits of 2D and 3D scene graph estimation methods by using 2D image features and 3D edge embedding. Image features are generally a better scene representation than 3D features, while using edge embedding in 3D allows performing relationship estimation without the constraint of the field of view. The effects of 2D and 3D features are compared in Sec. 4.2.

3.2.1 Node Feature

For each node $v_i \in \mathcal{V}$, we compute a multi-view image feature $\mathbf{v}_i$ and a geometric feature $\mathbf{g}_i$. We use the former as the initial node feature $\mathbf{f}_i^v = \mathbf{v}_i$ and include the latter with a learnable gate in the message passing step (Sec. 3.2.3).

The multi-view image feature is computed by aggregating multiple observations of v_i on images given by the entity visibility graph. For each view, an image feature is extracted with an image encoding network given the Region-of-Interest (ROI) of the node. The image features are aggregated using a mean operation to the multi-view image feature $\mathbf{v}_i$. Although there are sophisticated methods to compress multi-view image features, such as using gated averaging [15] and learning a canonical representation [63], we empirically found that averaging all the input features [57] yields the best result (see supplementary material). The mean operation also allows incrementally computing the multi-view image feature with a simple moving average. The geometric feature $\mathbf{g}_i$ is computed from the point set $\mathcal{P}_i$ using a simple point encoder [45].

3.2.2 Edge Feature

For each edge $e_{i\to j} \in \mathcal{E}_p$, an edge feature $\mathbf{f}_{i\to j}^e$ is computed using the node properties from its connected two nodes v_i and v_j by

$$\mathbf{f}_{i\to j}^e = g_s\left([\mathbf{o}_j - \mathbf{o}_i, \mathbf{b}_j - \mathbf{b}_i, R_{i\to j}]\right), \tag{4}$$

where $g_s(\cdot)$ is a Multilayer Perceptron (MLP), $[\cdot]$ denotes a concatenation function, and $R_{i\to j}$ is a relative pose descriptor which encodes the relative angle between two entities.

The relative pose descriptor is designed to implicitly encode relative angles between two nodes. Using an explicit one is not optimal since OBB estimations do not return the exact pose of an object, which makes explicit pose descriptor not applicable. We instead use the relative geometry properties on a reference frame constructed by the two nodes to implicitly encode the relative pose, as illustrated in Fig. 3. First, we construct a reference frame with the origin the midpoint of the center of two nodes, the x-axis to $\mathbf{o}_j$, the y-axis to the inverse of the gravity direction, and the z-axis the cross product of the x-axis and y-axis. Second, we take maximum and minimum values on each axis of the reference frame to compute the relative pose descriptor as

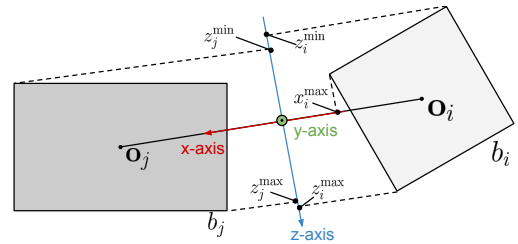

Figure 3. An illustration of our relative pose descriptor. The descriptor describe the relative maximum and minimum value of given two bounding boxes on a reference frame.

$$R_{i\to j} = \log\left(\left|\left[\mathbf{p}_i^{\max} \oslash \mathbf{p}_j^{\max}, \mathbf{p}_i^{\min} \oslash \mathbf{p}_j^{\min}\right]\right|\right), \tag{5}$$

where $\oslash$ is the Hadamard division, $\mathbf{p}_\square^{\max}, \mathbf{p}_\square^{\min} \in \mathbb{R}^3$ are the maximum and minimum points on the reference frame for $\square \in (i, j)$. We use an absolute logarithm ratio to improve the numerical stability.

3.2.3 Message Passing

Given an initial node feature $\mathbf{f}_i^v$ and an edge feature $\mathbf{f}_{i\to j}^e$, we aggregate the messages from the neighbors for both nodes and edges to enlarge the receptive field and leverage the spatial understanding composition of the environment. We follow [66] by aggregating the messages with a respective GRU unit shared for all nodes and edges. Following, we explain the process taking place in each of the message-passing layers.

First, we incorporate the geometric feature to each node feature using a learnable gate:

$$\hat{\mathbf{f}}_i^v = \mathbf{f}_i^v + \sigma\left(\mathbf{w}^T[\mathbf{f}_i^v, \mathbf{g}_i]\right)\sigma(\mathbf{g}_i), \tag{6}$$

where $\hat{\mathbf{f}}_i^v$ is the enhanced node feature, σ denotes a sigmoid function, and $\mathbf{w}^T$ are learnable parameters. The geometric feature may be unreliable, especially when the input geometry is ambiguous or unstable. Thus, we use the learnable gate to learn if the feature should be included. A node message m_i and an edge message $m_{i\to j}$ are computed by

$$m_i = g_v\left(\left[\hat{\mathbf{f}}_i^v, \max_{j\in\mathcal{N}(i)}\left(\text{FAN}\left(\hat{\mathbf{f}}_i^v, \mathbf{f}_{i\to j}^e, \hat{\mathbf{f}}_j^v\right)\right)\right]\right), \tag{7}$$

$$m_{i\to j} = g_e\left([\hat{\mathbf{f}}_i^v, \mathbf{f}_{i\to j}^e, \hat{\mathbf{f}}_j^v]\right), \tag{8}$$

where $g_v(\cdot)$ and $g_e(\cdot)$ are MLPs, $\mathcal{N}(i)$ is the set of indices representing the neighbouring nodes of i, FAN is the feature-wise attention network [64] which weights all neighbour node feature $\hat{\mathbf{f}}_j^v$ using input query $\hat{\mathbf{f}}_i^v$, and key $\mathbf{f}_{i\to j}^e$ given $j \in \mathcal{N}(i)$.

3.2.4 Class Prediction and Loss Functions

We use the softmax function to estimate the class distribution on both nodes and edges. For multiple predicate estimations, we use the sigmoid function (with a threshold of 0.5) to estimate whether a predicate exists. The network is trained with a cross-entropy loss for classifying both entities and edges. The loss for the edge class is replaced with binary cross entropy for multiple predicates estimation [61].

3.2.5 Fusion

Multiple predictions on the same nodes and edges are fused to ensure temporal consistency. We use the running average approach [8] to fuse predictions [64]. For each entity and edge, we store the full estimated probability estimation $\boldsymbol{\mu}^t$ and a weight $\rho^t \in \mathbb{R}_{\geq 0}$ at time t. Given a new prediction, we update the previously stored $\boldsymbol{\mu}^{t-1}$ and ρ^{t-1} as

$$\boldsymbol{\mu}^t = \frac{\boldsymbol{\mu}^t \cdot \rho^t + \boldsymbol{\mu}^{t-1} \cdot \rho^{t-1}}{\rho^t + \rho^{t-1}}, \tag{9}$$

$$\rho^t = \min\left(\rho_{\max}, \rho^t + \rho^{t-1}\right), \tag{10}$$

where $\rho_{\max}$ is the maximum weight value.

4. Evaluation

We evaluate our method on the task of 3D semantic scene graph estimation (Sec. 4.2) and incremental label association (Sec. 4.3). In addition, we provide ablation studies on the proposed network (Sec. 4.4), and a runtime analysis of our pipeline (Sec. 4.5).

4.1. Implementation Details

In all experiments, we use the default ORB-SLAM3 [3] setup provided by the authors[1] for our IEE front end. For the 2D entity detection, we use EntitySeg [46] with a ResNet50 [20] backbone pretrained on COCO [32] and fine-tuned on the 3RScan [60] training split. For multi-view feature extraction, we use a ResNet18 [20] pretrained on ImageNet [9] without fine-tuning. The point encoder is the vanilla PointNet without learned feature transformation [45]. Regarding hyperparameters, we set τ to 0.2, τ^C to 0.5 meters, $\rho_{\max}$ to 100, and the number of message passing layers to 2.

We use the ground truth pose to guide the scene reconstruction because (i) our focus lands on entity detection and scene graph estimation, and (ii) the provided image sequence from 3RScan [60] has a low frame rate (10 Hz), severe image blur, and jittery motion.

4.2. 3D Semantic Scene Graph estimation

For the input types, we compare all methods with the input of ground truth segmentation [61] (*GT*), geometric segmentation [59] (*Dense*) and sparse segmentation (*Sparse*). For the baseline methods, we compare ours with two 2D methods (IMP [66], and VGfM [15]), and two 3D methods (3DSSG [61] and SGFN [64]).

Baseline Methods. We will briefly discuss baseline methods here. Check supplementary for further details. IMP [66] computes a node feature using the image feature cropped from the ROI of the node in an image and computes an edge feature using the union of two ROIs from its connected nodes. Both features are jointly updated with prime-dual message passing and learnable message pooling. VGfM [15] extends IMP by adding geometric features and temporal message passing to handle sequential estimation. 3DSSG [61] extends the ROI concept in IMP [66] in

	Method	*Recall*(%)			*mRecall*(%)	
		Rel.	Obj.	Pred.	Obj.	Pred.
GT	IMP [66]	49.8	70.1	94.3	53.0	38.1
	VGfM [15]	49.3	69.4	94.8	57.5	44.6
	3DSSG [61]	34.6	58.0	95.2	46.8	58.7
	SGFN [64]	41.8	63.8	94.3	57.7	65.5
	Ours	**66.1**	**81.2**	**95.6**	**77.4**	**71.5**
Dense	IMP [66]	25.8	51.8	90.4	30.0	23.0
	VGfM [15]	28.3	53.3	**90.7**	31.6	24.4
	3DSSG [61]	17.5	41.4	88.2	31.9	26.6
	SGFN [64]	31.4	56.7	89.6	38.3	30.5
	Ours	**34.1**	**58.1**	89.9	**43.0**	**33.3**
Sparse	IMP [66]	7.9	27.5	90.7	20.6	14.0
	VGfM [15]	8.2	26.9	**90.8**	17.6	15.4
	3DSSG [61]	0.9	9.7	87.9	5.9	15.1
	SGFN [64]	1.7	12.6	88.9	8.3	14.4
	Ours	9.9	29.5	90.4	23.5	**16.5**
	Ours (i)	**10.7**	**30.2**	90.4	**24.5**	15.9

Table 1. We compare our method with four baseline methods on the task of scene graph prediction on 3RScan [60] dataset with 20 objects and 8 predicate classes. The results from Ours are obtained by using our network to obtain predictions, while Ours (i) contains the results from using the incremental pipeline.

Method	*Recall*(%)			*mRecall*(%)	
	Rel.	Obj.	Pred.	Obj.	Pred.
IMP [66]	44.5	35.9	9.0	18.7	4.9
VGfM [15]	44.5	37.9	14.7	17.9	6.5
3DSSG [61]	46.8	29.6	**68.8**	11.7	**25.5**
SGFN [64]	45.2	29.4	42.8	11.8	13.5
Ours	**52.7**	**56.7**	50.4	**27.2**	23.9

Table 2. Evaluation on scene graph prediction with 160 object and 26 predicate classes using ground truth segmentation and fully connected neighbor graph.

[1] https://github.com/UZ-SLAMLab/ORB_SLAM3.git

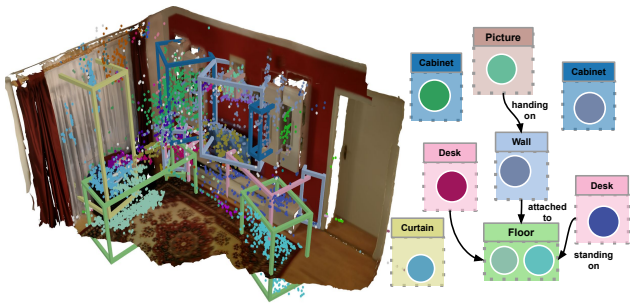

Figure 4. Qualitative evaluation of our scene graph prediction framework. Each 3D bounding box represents a detected entity on the left, and the color is the predicted label. On the right side, we visualize the estimated scene graph on this scene. We only select representative entities on the scene graph visualization for visualization purposes.

3D by replacing ROIs to 3D bounding boxes. The node and edge features are computed with PointNet [45]. Both features are jointly updated with a graph neural network with average message pooling. SGFN [64] improves 3DSSG by replacing the initial edge descriptor with the relative geometry properties between two nodes and introducing an attention method to handle dynamic message aggregation, which enables incremental scene graph estimation.

Implementation. For all methods, we follow their implementation details and train on 3RScan dataset [60] from scratch until converge, using a custom training and test split since the scenes in the original test do not have ground truth scene graphs provided. For IMP [66], since it is a single image prediction method, we adopt the voting mechanism as in [27] to average the prediction over multiple frames. Since ours and other 2D baseline methods rely on image input, we generate a set of keyframes by sampling all input frames using their poses for the *GT* and *Dense* inputs (check supplementary material for further details). To ensure diversity in the viewpoint, we filter out a frame if its pose is too similar to any selected frames with the threshold values of 5 degrees in rotation and 0.3 meters in translation.

Evaluation Metric. We report the overall recall (*Recall*) as used in many scene graph work [34, 61, 64, 66, 68] but with the strictest top-k metric with $k = 1$ as in [62]. In addition, we report the mean recall (*mRecall*) which better indicates model performance when the input dataset has a severe data imbalance issue (see supplementary material for the class distribution). Moreover, since different segmentation methods may result in different number of segments, we map all predictions on estimated segmentation back to ground truth. This allows us to compare the reported numbers across different segmentation methods. We report the *Recall* of relationship triplet estimation (Rel.), object class estimation (Obj.) and predicate estimation (Pred.), and the *mRecall* of object class estimation and predicate estimation.

Results. Following the evaluation scheme in [64] and in [61], we report two evaluations in Tbl. 1 and Tbl. 2, respectively. The former one maps the node classes to 20 NYUv2 labels [39] to suppress the severe class imbalance in the data as discussed in [62] and estimates a single predicate out of seven support relationship types plus the "same part" relationship to handle over-segmentation. The latter uses 160 node and 26 edge classes with multiple predicate estimation.

In Tbl. 1, overall, it can be seen that all image-based methods (IMP, VGfM, Ours) outperforms points-based methods (3DSSG, SGFN) in almost all object prediction metrics, while the methods based on 3D edge descriptor (3DSSG, SGFN, Ours) tend to have better predicate estimation. This suggests that the 2D representations from images are more representative than 3D, and 3D edge descriptors are more suitable for estimating support types of predicates. By comparing IMP [66] and VGfM [15], it can be seen that the effect of the geometric feature and the temporal message passing is mainly reflected in the *mRecall* metrics. However, it deteriorates the performance when the input is sparse. A possible reason is that the geometric feature is relatively unstable, which decreases the network performance. It is also reflected in the two 3D methods, *i.e.* 3DSSG and SGFN, where they failed to perform in classifying objects with sparse segmentation while giving a similar performance in the predicate classification. Among all methods, our method outperforms all baselines among all input types and all metrics, apart from the predicate estimation, which has a slightly worse result on some input types. In addition, we report ours using the proposed incremental estimation pipeline, denoted as Ours(i). The incremental estimation process improves slightly in object estimation. The same behavior is also reported in [64]. We show a qualitative result using our full pipeline in Fig. 4.

In Tbl. 2, our method outperforms all other methods in the relationship and object estimation in *Recall* and object estimation in *mRecall*. 3DSSG [61] has the best results in predicate estimations. This suggests that union 3D bounding boxes are more suitable when estimating multiple predicates.

4.3. Incremental Label Association

We evaluate our label association method in the task of incremental entity segmentation, which aims to estimate accurate class-agnostic segmentation given sequential sensor input, with two baseline methods, *i.e.* InSeg [58] and PanopticFusion [38].

Baseline Methods. Both baseline methods use reference map approaches as mentioned in Sec. 3.1.3. InSeg [58] considers only the overlapping ratio between labels on an estimated mask and a reference mask. PanopticFusion uses IoU [38] as the evaluation method and limits one reference label can only be assigned to one query label.

Method	AOS (%)
InSeg [58]	38.6
PanopticFusion [38]	35.9
Ours	**39.6**

Table 3. Evaluation of different label association methods in the task of incremental entity estimation in 3RScan dataset [60].

Implementation. For all methods, we use our IEE pipeline with different label association methods on all training scenes in the 3RScan dataset [60].

Evaluation Metric. We use the Average Overlap Score (AOS) as the evaluation metric [59]. It measures the ratio of the dominant segment of its corresponding ground truth instance. We use the nearest neighbour search to find the ground truth instance label of a reconstructed point. Since our method reconstructs a sparse point map, using the ground truth points as the denominator does not reflect the performance. We instead calculate the score over the sum of all estimated points within the ground truth instance as:

$$\text{AOS} = \sum_i \frac{\max_j \text{Overlap}(\mathcal{S}_i, \mathcal{P}_j)}{\#(\mathcal{S}_i)}, \tag{11}$$

where $\mathcal{S}_i$ is the set of all estimated points with ground truth instance label i.

Results. The evaluation result is reported in Tbl. 3. Our approach achieves the highest AOS, which is 1 % higher than InSeg [58] and 3.7% higher than PanopticFusion [38]. The use of a confidence-based approach handles better the label consistency and thus improves the final AOS score. We provide an example of how our method improves temporal consistency under non-uniform distributed points in the supplementary material.

4.4. Ablation Study

We ablate our network with two components, *i.e.* geometric descriptor $\mathbf{g}_i$ and relative pose descriptor $R_{i\to j}$. The experiment setup is the same as in [64], which makes the ablation comparable to Tbl. 1. The result is reported in Tbl. 4. More ablation studies are in the supplementary material.

Our vanilla network without $\mathbf{g}_i$ and $R_{i\to j}$ outperform baselines in most of the metrics. With $\mathbf{g}_i$, there is a consistent improvement on all metrics except the Pred. in *mRecall*. Compared to VGfM [15] in Tbl. 1, VGfM [15] fails to improve the performance of IMP [66] with sparse input. Our gated geometric feature aggregation improves its baseline with sparse input, bringing a more consistent performance gain than VGfM [15]. The $R_{i\to j}$ improves the *mRecall* performance with *GT* and *Dense* inputs but decreases the model *Recall* performance. This behavior suggests that $R_{i\to j}$ helps handle class imbalance issues. The combination of both components achieves the best *Recall*. However, the model tends to focus on dominant classes, resulting in slightly worse performance in *mRecall*.

	Method		*Recall*(%)			*mRecall*	
	$\mathbf{g}_i$	$R_{i\to j}$	Rel.	Obj.	Pred.	Obj.	Pred.
GT			61.9	76.4	95.6	74.3	69.2
	✓		62.9	77.9	**95.9**	74.2	64.3
		✓	60.4	76.3	95.0	75.3	**73.2**
	✓	✓	**66.1**	**81.2**	95.6	**77.4**	71.5
Dense			30.2	54.0	88.5	44.9	33.2
	✓		33.9	56.4	89.7	45.7	33.8
		✓	28.7	52.7	88.2	**47.5**	**34.3**
	✓	✓	**34.1**	**58.1**	**89.9**	43.0	33.3
Sparse			9.6	28.6	90.0	25.6	17.7
	✓		9.8	28.5	90.0	**25.7**	**18.1**
		✓	9.6	28.1	90.2	23.3	16.9
	✓	✓	**9.9**	**29.5**	**90.4**	23.5	16.5

Table 4. Ablation study on the proposed network. We ablate the proposed gated geometric feature ($\mathbf{g}_i$) and the relative pose descriptor ($R_{i\to j}$) using the same experiment setup as in Tbl. 1.

	Front end	*Back end*		
	Sparse Mapping	2D Entity Est.	Label Fusion	Scene Graph Est.
Mean [ms]	14.7	124.6	14.2	52.5

Table 5. Runtime [ms] of the different components of our method.

4.5. Runtime

We report the runtime of our system on 3RScan [60] sequence 4acaebcc-6c10-2a2a-858b-29c7e4fb410d in Tbl. 5. The analysis is done with a machine equipped with an Intel Core i7-8700 3.2GHz CPU with 12 threads and a NVidia GeForce RTX 2080ti GPU.

5. Conclusion

We present a novel method that estimates 3D scene graphs from RGB images incrementally. Our method runs in real-time and does not rely on depth inputs, which could benefit other tasks, such as robotics and AR, that have hardware limitations and real-time demand. The experiment results indicate that our method outperforms others in three different input types. The provided ablation study demonstrates the effectiveness of our design. Our vanilla network, without any geometric input and relative pose descriptor, still outperforms other baselines. Our method provides a novel architecture for estimating scene graphs with only RGB input. The multiview feature is proven to be more powerful than existing 3D methods. Our method can be further improved in many directions. In particular, using semi-direct SLAM methods such as SVO [12] might improve the handling of untextured regions where feature-based methods often fail. In addition, the multiview image encoder could be replaced with a more powerful encoder to improve the scene graph estimation with a computational penalty.

References

[1] Iro Armeni, Zhi-Yang He, JunYoung Gwak, Amir R Zamir, Martin Fischer, Jitendra Malik, and Silvio Savarese. 3D SceneGgraph: A Structure for Unified Semantics, 3D Space, and Camera. In *ICCV*, 2019. 1, 2, 3

[2] Garrick Brazil, Gerard Pons-Moll, Xiaoming Liu, and Bernt Schiele. Kinematic 3d object detection in monocular video. In *European Conference on Computer Vision*, pages 135–152. Springer, 2020. 2

[3] Carlos Campos, Richard Elvira, Juan J Gómez Rodríguez, José MM Montiel, and Juan D Tardós. Orb-slam3: An accurate open-source library for visual, visual–inertial, and multimap slam. *IEEE Transactions on Robotics*, 37(6):1874–1890, 2021. 2, 3, 6

[4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 4

[5] Chia-Tche Chang, Bastien Gorissen, and Samuel Melchior. Fast oriented bounding box optimization on the rotation group $SO(3, \mathbb{R})$. *ACM Trans. Graph.*, 30(5):122:1–122:16, Oct. 2011. 4

[6] Xieyuanli Chen, Andres Milioto, Emanuele Palazzolo, Philippe Giguere, Jens Behley, and Cyrill Stachniss. Suma++: Efficient lidar-based semantic slam. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4530–4537. IEEE, 2019. 2

[7] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34, 2021. 4

[8] Brian Curless and Marc Levoy. A Volumetric Method for Building Complex Models from Range Images. In *Proceedings Conference on Computer Graphics and Interactive Techniques*, 1996. 6

[9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6

[10] Helisa Dhamo, Azade Farshad, Iro Laina, Nassir Navab, Gregory D Hager, Federico Tombari, and Christian Rupprecht. Semantic Image Manipulation Using Scene Graphs. In *CVPR*, 2020. 1

[11] Helisa Dhamo, Fabian Manhardt, Nassir Navab, and Federico Tombari. Graph-to-3d: End-to-end generation and manipulation of 3d scenes using scene graphs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16352–16361, 2021. 1

[12] Christian Forster, Zichao Zhang, Michael Gassner, Manuel Werlberger, and Davide Scaramuzza. Svo: Semidirect visual odometry for monocular and multicamera systems. *IEEE Transactions on Robotics*, 33(2):249–265, 2016. 8

[13] Sarthak Garg, Helisa Dhamo, Azade Farshad, Sabrina Musatian, Nassir Navab, and Federico Tombari. Unconditional scene graph generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16362–16371, 2021. 1

[14] Stefano Gasperini, Frithjof Winkelmann, Alvaro Marcos-Ramiro, Micheal Schmidt, Nassir Navab, Benjamin Busam, and Federico Tombari. Holistic segmentation. *arXiv preprint arXiv:2209.05407*, 2022. 4

[15] Paul Gay, James Stuart, and Alessio Del Bue. Visual Graphs from Motion (VGfM): Scene Understanding with Object Geometry Reasoning. In *ACCV*. Springer, 2018. 1, 2, 3, 4, 5, 6, 7, 8

[16] 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*, 2019. 2

[17] Markus Grotz, Peter Kaiser, Eren Erdal Aksoy, Fabian Paus, and Tamim Asfour. Graph-based visual semantic perception for humanoid robots. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, pages 869–875. IEEE, 2017. 1

[18] Jiuxiang Gu, Shafiq Joty, Jianfei Cai, Handong Zhao, Xu Yang, and Gang Wang. Unpaired image captioning via scene graph alignments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10323–10332, 2019. 1

[19] Jiuxiang Gu, Handong Zhao, Zhe Lin, Sheng Li, Jianfei Cai, and Mingyang Ling. Scene Graph Generation with External Knowledge and Image Reconstruction. In *CVPR*, 2019. 1

[20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 6

[21] Mehdi Hosseinzadeh, Yasir Latif, Trung Pham, Niko Suenderhauf, and Ian Reid. Structure aware slam using quadrics and planes. In *Asian Conference on Computer Vision*, pages 410–426. Springer, 2018. 2

[22] Hou-Ning Hu, Qi-Zhi Cai, Dequan Wang, Ji Lin, Min Sun, Philipp Krahenbuhl, Trevor Darrell, and Fisher Yu. Joint monocular 3d vehicle detection and tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5390–5399, 2019. 2

[23] N. Hughes, Y. Chang, and L. Carlone. Hydra: A real-time spatial perception system for 3D scene graph construction and optimization. In *Robotics: Science and Systems (RSS)*, 2022. 1, 2, 3

[24] Justin Johnson, Agrim Gupta, and Li Fei-Fei. Image Generation from Scene Graphs. In *CVPR*, 2018. 1

[25] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5830–5840, 2021. 4

[26] Andrej Karpathy and Li Fei-Fei. Deep Visual-Semantic Alignments for Generating Image Descriptions. In *CVPR*, 2015. 1

[27] Ue-Hwan Kim, Jin-Man Park, Taek-Jin Song, and Jong-Hwan Kim. 3-d scene graph: A sparse and semantic representation of physical environments for intelligent agents. *IEEE transactions on cybernetics*, 50(12):4921–4933, 2019. 1, 2, 3, 7

[28] Chi Li, Jin Bai, and Gregory D Hager. A unified framework for multi-view multi-class object pose estimation. In *Proceedings of the european conference on computer vision (ECCV)*, pages 254–269, 2018. 2

[29] Kejie Li, Daniel DeTone, Yu Fan Steven Chen, Minh Vo, Ian Reid, Hamid Rezatofighi, Chris Sweeney, Julian Straub, and Richard Newcombe. Odam: Object detection, association, and mapping using posed rgb video. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5998–6008, 2021. 2

[30] Kejie Li, Hamid Rezatofighi, and Ian Reid. Moltr: Multiple object localization, tracking and reconstruction from monocular rgb videos. *IEEE Robotics and Automation Letters*, 6(2):3341–3348, 2021. 2

[31] Yanwei Li, Hengshuang Zhao, Xiaojuan Qi, Liwei Wang, Zeming Li, Jian Sun, and Jiaya Jia. Fully convolutional networks for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 214–223, 2021. 4

[32] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 6

[33] Hengyue Liu, Ning Yan, Masood Mortazavi, and Bir Bhanu. Fully convolutional scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11546–11556, 2021. 1

[34] Cewu Lu, Ranjay Krishna, Michael Bernstein, and Li Fei-Fei. Visual relationship detection with language priors. In *European conference on computer vision*, pages 852–869. Springer, 2016. 7

[35] John McCormac, Ronald Clark, Michael Bloesch, Stefan Leutenegger, and Andrew Davison. Fusion++: Volumetric Object-Level SLAM. In *Int. Conf. on 3D Vis.*, 2018. 2, 4

[36] John McCormac, Ankur Handa, Andrew Davison, and Stefan Leutenegger. SemanticFusion: Dense 3D Semantic Mapping with Convolutional Neural Networks. In *Int. Conf. Robotics and Automation*, 2017. 2

[37] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3d bounding box estimation using deep learning and geometry. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 7074–7082, 2017. 2

[38] Gaku Narita, Takashi Seno, Tomoya Ishikawa, and Yohsuke Kaji. PanopticFusion: Online Volumetric Semantic Mapping at the Level of Stuff and Things. In *IEEE Conf. Intelligent Robots and Syst.*, 2019. 2, 4, 7, 8

[39] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor Segmentation and Support Inference from RGBD Images. In *ECCV*, 2012. 7

[40] Lachlan Nicholson, Michael Milford, and Niko Sünderhauf. Quadricslam: Dual quadrics from object detections as landmarks in object-oriented slam. *IEEE Robotics and Automation Letters*, 4(1):1–8, 2018. 2

[41] Yinyu Nie, Xiaoguang Han, Shihui Guo, Yujian Zheng, Jian Chang, and Jian Jun Zhang. Total3dunderstanding: Joint layout, object pose and mesh reconstruction for indoor scenes from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 55–64, 2020. 2

[42] Ege Özsoy, Evin Pınar Örnek, Ulrich Eck, Tobias Czempiel, Federico Tombari, and Nassir Navab. 4d-or: Semantic scene graphs for or domain modeling. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2022. 1

[43] Ege Özsoy, Evin Pınar Örnek, Ulrich Eck, Federico Tombari, and Nassir Navab. Multimodal semantic scene graphs for holistic modeling of surgical procedures. In *Arxiv*, 2021. 1

[44] Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Real-time Progressive 3D Semantic Segmentation for Indoor Scene. In *IEEE Winter Conf. Appli. Comput. Vis.*, 2019. 2

[45] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *CVPR*, 2017. 5, 6, 7

[46] Lu Qi, Jason Kuen, Yi Wang, Jiuxiang Gu, Hengshuang Zhao, Zhe Lin, Philip Torr, and Jiaya Jia. Open-world entity segmentation. *arXiv preprint arXiv:2107.14228*, 2021. 4, 6

[47] Zachary Ravichandran, Lisa Peng, Nathan Hughes, J Daniel Griffith, and Luca Carlone. Hierarchical representations and explicit memory: Learning effective navigation policies on 3d scene graphs using graph neural networks. In *2022 International Conference on Robotics and Automation (ICRA)*, pages 9272–9279. IEEE, 2022. 1

[48] Antoni Rosinol, Arjun Gupta, Marcus Abate, Jingnan Shi, and Luca Carlone. 3D Dynamic Scene Graphs: Actionable Spatial Perception with Places, Objects, and Humans. *Comput. Res. Repository*, 2020. 2, 3

[49] Antoni Rosinol, Arjun Gupta, Marcus Abate, Jingnan Shi, and Luca Carlone. 3d dynamic scene graphs: Actionable spatial perception with places, objects, and humans. *arXiv preprint arXiv:2002.06289*, 2020. 2

[50] Antoni Rosinol, Andrew Violette, Marcus Abate, Nathan Hughes, Yun Chang, Jingnan Shi, Arjun Gupta, and Luca Carlone. Kimera: From slam to spatial perception with 3d dynamic scene graphs. *The International Journal of Robotics Research*, 40(12-14):1510–1546, 2021. 2, 3

[51] Martin Runz, Maud Buffier, and Lourdes Agapito. Maskfusion: Real-time recognition, tracking and reconstruction of multiple moving objects. In *2018 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*, pages 10–20. IEEE, 2018. 4

[52] Radu Bogdan Rusu and Steve Cousins. 3D is here: Point Cloud Library (PCL). In *IEEE International Conference on Robotics and Automation (ICRA)*, Shanghai, China, May 9-13 2011. 4

[53] Renato F. Salas-Moreno, Richard A. Newcombe, Hauke Strasdat, Paul H. J. Kelly, and Andrew J. Davison. SLAM++: Simultaneous Localisation and Mapping at the Level of Objects. In *CVPR*, 2013. 2

[54] Raluca Scona, Simona Nobili, Yvan R Petillot, and Maurice Fallon. Direct visual slam fusing proprioception for a humanoid robot. In *2017 IEEE/RSJ International Conference*

on Intelligent Robots and Systems (IROS), pages 1419–1426. IEEE, 2017. 1

[55] Philipp Seiwald, Shun-Cheng Wu, Felix Sygulla, Tobias FC Berninger, Nora-Sophie Staufenberg, Moritz F Sattler, Nicolas Neuburger, Daniel Rixen, and Federico Tombari. Lola v1. 1–an upgrade in hardware and software design for dynamic multi-contact locomotion. In *2020 IEEE-RAS 20th International Conference on Humanoid Robots (Humanoids)*, pages 9–16. IEEE, 2021. 1

[56] Olivier Stasse, Andrew J Davison, Ramzi Sellaouti, and Kazuhito Yokoi. Real-time 3d slam for humanoid robot considering pattern generator information. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 348–355. IEEE, 2006. 1

[57] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 945–953, 2015. 5

[58] Keisuke Tateno, Federico Tombari, and Nassir Navab. Real-Time and Scalable Incremental Segmentation on Dense SLAM. In *IEEE Conf. Intelligent Robots and Syst.*, 2015. 2, 4, 7, 8

[59] Keisuke Tateno, Federico Tombari, and Nassir Navab. Large scale and long standing simultaneous reconstruction and segmentation. *Computer Vision and Image Understanding*, 157:138–150, 2017. 2, 6, 8

[60] Johanna Wald, Armen Avetisyan, Nassir Navab, Federico Tombari, and Matthias Nießner. RIO: 3D Object Instance Re-Localization in Changing Indoor Environments. In *ICCV*, 2019. 2, 6, 7, 8

[61] Johanna Wald, Helisa Dhamo, Nassir Navab, and Federico Tombari. Learning 3D Semantic Scene Graphs from 3D Indoor Reconstructions. In *CVPR*, 2020. 1, 2, 3, 4, 6, 7

[62] Johanna Wald, Nassir Navab, and Federico Tombari. Learning 3d semantic scene graphs with instance embeddings. *International Journal of Computer Vision*, pages 1–22, 2022. 7

[63] Xin Wei, Yifei Gong, Fudong Wang, Xing Sun, and Jian Sun. Learning canonical view representation for 3d shape recognition with arbitrary views. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 407–416, 2021. 5

[64] Shun-Cheng Wu, Johanna Wald, Keisuke Tateno, Nassir Navab, and Federico Tombari. Scenegraphfusion: Incremental 3d scene graph prediction from rgb-d sequences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7515–7525, 2021. 1, 2, 3, 4, 5, 6, 7, 8

[65] Yanmin Wu, Yunzhou Zhang, Delong Zhu, Yonghui Feng, Sonya Coleman, and Dermot Kerr. Eao-slam: Monocular semi-dense object slam based on ensemble data association. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4966–4973. IEEE, 2020. 2

[66] Danfei Xu, Yuke Zhu, Christopher Choy, and Li Fei-Fei. Scene Graph Generation by Iterative Message Passing. In *CVPR*, 2017. 1, 2, 4, 5, 6, 7, 8

[67] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, Attend and Tell: Neural Image Caption Generation with Visual Attention. In *Int. Conf. Mach. Learn.*, 2015. 1

[68] Jianwei Yang, Jiasen Lu, Stefan Lee, Dhruv Batra, and Devi Parikh. Graph R-CNN for Scene Graph Generation. In *ECCV*, 2018. 7

[69] Shichao Yang and Sebastian Scherer. Cubeslam: Monocular 3-d object slam. *IEEE Transactions on Robotics*, 35(4):925–938, 2019. 2

[70] Cheng Zhang, Zhaopeng Cui, Yinda Zhang, Bing Zeng, Marc Pollefeys, and Shuaicheng Liu. Holistic 3d scene understanding from a single image with implicit representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8833–8842, 2021. 2

[71] Jiazhao Zhang, Chenyang Zhu, Lintao Zheng, and Kai Xu. Fusion-Aware Point Convolution for Online Semantic 3D Scene Segmentation. In *CVPR*, 2020. 2

[72] Yiwu Zhong, Jing Shi, Jianwei Yang, Chenliang Xu, and Yin Li. Learning to generate scene graph from natural language supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1823–1834, 2021. 1

Supplementary Material: Incremental 3D Semantic Scene Graph Prediction from RGB Sequences

Shun-Cheng Wu[1] Keisuke Tateno[2] Nassir Navab[1] Federico Tombari[1,2]
[1]Technische Universität München [2]Google

arXiv:2305.02743v2 [cs.CV] 6 May 2023

6. Method Comparison and Implementation

We provide an overview of all the methods evaluated in Sec. 4.2, including our adaptation to enable comparisons among them in our experiments. As a brief recall, we compare our method to two image-based methods, *i.e*. IMP [16] and VGfM [2], and two point-based methods, 3DSSG [12] and SGFN [14]. All methods shared a similar scene graph generation pipeline:

1. using a node and an edge encode to compute an initial node and edge embeddings.
2. using message passing to calculate messages for updating node and edge features.
3. updating node and edge feature with the messages from step 2.
4. integrating class prediction over time.

For all methods, we closely follow the original implementation from the respective papers. We refer interested readers to check out their papers for all detail. Here, we describe our adaptation to enable fair comparison in our experiments. We report the comparison of all five methods in Tab. 6.

Object Detection. IMP and VGfM rely on a regional proposal network to detect objects. We replace it with our entity detection methods described in Sec. 3.1.2.

Geometric Features in VGfM. For VGfM, they extract geometric features from ellipsoids. In our implementation, we replace the use of ellipsoids with oriented bounding boxes, which can provide equivalent information as needed by VGfM.

Fusion for IMP. For IMP, as mentioned in the main paper, we added the voting mechanism in [5] to fuse multiple predictions. For handling the incremental nature, our experiment in Table 1 follows the setup in [14]. We do a single global estimation of all methods and provide incremental estimations of our methods. Hence, no modification is needed for all the baseline methods.

Entity visibility graph and the neighbour graph for the *GT* setup Unlike in *Dense* and *Sparse*, where we run an incremental estimation system to obtain $\mathcal{G}_c$ and $\mathcal{G}_p$, the *GT* requires an additional procedure to obtain the entity visibility graph $\mathcal{G}_c = (\mathcal{V}, \mathcal{K}, \mathcal{E}_c)$ and the neighbour graph $\mathcal{G}_p = (\mathcal{V}, \mathcal{E}_p)$. The entities $\mathcal{V}$ are directly inherited from the ground truth annotation from the 3RScan dataset [11]. The proximity edges $\mathcal{E}_p$ is estimated with the same strategy as described in Sec. 3.1.4, using the ground truth bounding boxes of entities. Here we detail how we estimate $\mathcal{K}$ and $\mathcal{E}_c$ for the *GT* setup. For $\mathcal{K}$ and $\mathcal{E}_c$, we follow a similar approach as in VGfM [2], where we first find all relevant frames across all entities and then select keyframes with our keyframe selection strategy (Sec. 7). We use the ground truth instance mask from [11] to check if an entity appears in an image. However, an entity may be heavily occluded and cannot provide reasonable image features. To avoid this, we estimate the occupancy of an entity in an image as the ratio of the number of relevant pixels over all pixels within the bounding box of the entity. Furthermore, to prevent an entity is not aligned to the image coordinate, which will cause the occupancy value to be very low even if it is not occluded, we downscale each input mask by a factor of 8 with a maximum relevant selection, *i.e*. a downscaled pixel is considered relevant if one among the eight pixels in the original image is relevant. This gives the visibility of each entity on each input frame. We then apply the keyframe selection strategy to prevent duplicate views and to ensure good view coverage.

7. Keyframe selection strategy.

Selecting keyframes is crucial when the estimation quality is solely based on multiview images. Having diverse view coverage of objects usually results in better feature representation of objects [4, 13, 17]. Unlike the keyframes selection in ORBSLM [3], which focuses more on the pose estimation quality than the view coverage, we select keyframes mainly based on the pose difference and, in addition, the quality of detected objects. A frame is selected as a keyframe only 1) it has at least one valid object detected

Method	Node Type	Edge Type	Message Passing Type	Message Update Method	Fusion
IMP	Image ROI	Image ROI Union	Prime-Dual	GRU	Voting [5]
VGfM	Image ROI	Image ROI Union	Prime-Dual + Geo. description	GRU	Temporal Gate
3DSSG	Points	Points Union	Triplet	concatenation	N/A
SGFN	Points	geometric description	FAN	concatenation	running mean
Ours	MV Image ROIs	geometric description	FAN + Gated Points	GRU	running mean

Table 6. A summary of the modules used for different methods. The text colored in cyan involves our modification to make all methods comparable.

and 2) its pose is dissimilar to other existing keyframes. We measure the validity of bounding boxes by checking if their minimum width and height are larger than 200 pixels, and the pose difference threshold is set to 5 degrees in rotation and 0.3 meters in translation. This keyframe selection method is used for all input, *i.e. GT*, *Dense*, and *Sparse*, cases in our experiments.

8. Data Distribution

We provide the class distribution on objects and predicates in Fig. 5. It can be seen that the setup in [12] has severe long-tail data distribution. After mapping to 20 NYUv2 labels [8], the distribution is relatively well distributed but still unbalanced. The unbalanced distribution indicates that the *mRecall* metric reflects better the model performance.

9. Multi-view Feature Encoding

As mentioned in Sec. 3.2.1, the multi-view feature of a node is computed by aggregating image features of the node from multiple images. This is essentially the task of 3D shape recognition with arbitrary views. We compare the use of simple mean aggregation, as in MVCNN [9], with the state-of-the-art method, CVR [13] on ScanNet [1] dataset. Since the main comparison is on the multi-view feature aggregation, we use the same backbone, ResNet18 [3], for both methods. All networks are trained from scratch using the splits from ScanNet, following the same training approach described in the respective papers. We report the mean intersection-over-union (*mIoU*), precision (*mPrec*), and recall (*mRecall*).

The result is shown on Tbl. 7. It can be seen that the use of simple mean operation outperforms canonical transformation in CVR [13]. Although the number reported in CVR [13], it outperforms MVCNN in ModelNet40 [15], ScanObjectNN [10] and RGBD [6]. We investigated the difference between the three datasets mentioned above and ScanNet and found that the images from ScanNet [1] contain more background objects while the others have non-cluttered backgrounds. We demonstrate some example images in Fig. 6. The performance inconsistency may indicate

	mIoU(%)	*mPrec*(%)	*mRecall*(%)
CVR [13]	32.3	45.2	54.4
MVCNN [9]	**39.2**	**50.5**	**61.4**

Table 7. Object classification result on ScanNet dataset. A simple averaging of overall image features (MVCNN) outperforms the sophisticated multi-view image encoding method (CVR).

that using the mean operator makes the model less sensitive to background things.

10. Comparing confidence and IoU-based methods

We provide an example of the difference between using the maximum IoU and our maximum mean confidence to find the most probable correspondence with a sparse point map. In figure 7, given three consecutive frames at $t = n$, $t = n+1$ and $t = n+2$ for the label association and fusion, their entity maps and the association result are shown on the second and third columns (separated by white space). The second column is the label fusion result with IoU [7], and the third column is ours. When using IoU (second column), the table label at $t = n$ is assigned to the floor at $t = n + 2$. This is due to the map points created at $t = n + 1$ on the floor (carpet) having larger IoU than the table. With our approach (third column), the table label at $t = n$ remains at the table at $t = n + 2$. This shows that our method provides more consistent label association. Note that the label colors are different between IoU and ours since the segment color are randomly in each run.

11. Ablation Study

We provide two additional ablation studies: (1) the use of edge descriptor (Tab. 8) and (2) the effect of the sigmoid gate on the geometric feature (Tab. 9). In Tab. 8, it can be seen that the use of our relative pose descriptor $R_{i\to j}$ leads to better overall performance in *GT* and *Sparse* settings while having slightly worse performance in *Dense* setting. In Tab. 9, using the gated geometric feature consistently achieves better performance in all three setups.

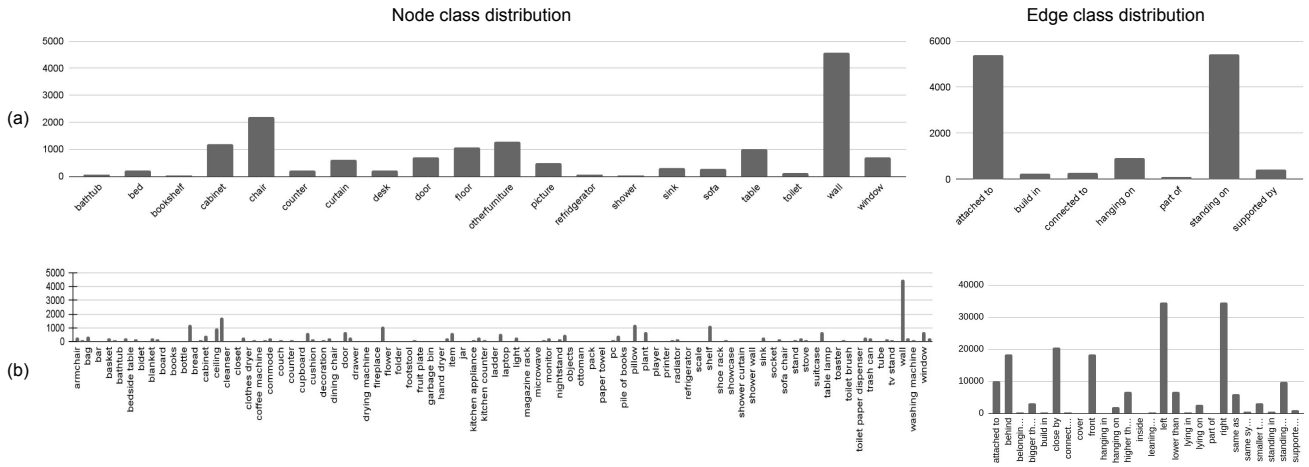

Figure 5. We provide the class distribution for two scene graph experiments in Tbl. 1 and Tbl. 2. The row (a) is the distribution for the experiment setup in [14] and in Tbl. 1. (b) is the distribution for the experiment setup in [12] and in Tbl. 2.

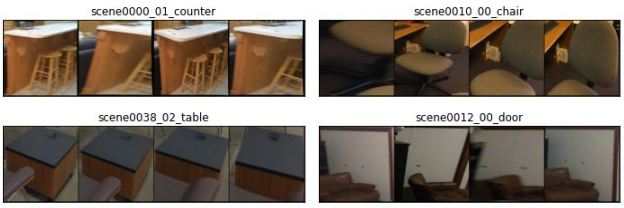

Figure 6. We provide examples of multi-view images of the same object in the ScanNet dataset. For display purposes, we only select four views of the same object. It can be seen that the multi-view observation of an object in a standard indoor dataset includes some non-target objects in the field of view. Those objects are considered noise which affects the multi-view image encoding.

	Edge Type	*Recall*(%) Rel.	Obj.	Pred.	*mRecall* Obj.	Pred.
GT	*Points*	62.0	77.9	**95.9**	68.4	69.0
	$R_{i\to j}$	**66.1**	**81.2**	95.6	**77.4**	**71.5**
Dense	*Points*	**35.1**	**57.5**	89.6	**47.9**	33.2
	$R_{i\to j}$	34.1	58.1	**89.9**	43.0	**33.3**
Sparse	*Points*	10.0	28.7	**90.6**	21.1	16.3
	$R_{i\to j}$	**9.9**	**29.5**	90.4	**23.5**	**16.5**

Table 8. Ablation study on the use of input type for computing edge feature. The experiment setup is the same as in Tbl. 1. Here *Points* means taking the point cloud union as in [12], and $R_{i\to j}$ is the relative pose descriptor described in Sec. 3.2.2.

11.1. Comparison in edge descriptor

compare non-learned and learned descriptors.

	Gate	*Recall*(%) Rel.	Obj.	Pred.	*mRecall* Obj.	Pred.
GT		61.5	77.1	95.3	77.1	70.9
	✓	**66.1**	**81.2**	**95.6**	**77.4**	**71.5**
Dense		32.9	55.5	89.1	41.0	31.4
	✓	**34.1**	**58.1**	**89.9**	**43.0**	**33.3**
Sparse		8.6	26.9	**90.5**	24.4	15.6
	✓	**9.9**	**29.5**	90.4	**23.5**	**16.5**

Table 9. Ablation study on the proposed gated geometric feature. We ablate the use of a sigmoid function of a gate for the input geometric feature using the same experimental setup as in Tbl. 1.

11.2. Gate on the geometric feature

whether to use the gate in the geometric feature or to use different ways of selecting keyframes. We provide an additional ablation study on the effect of using a sigmoid gate when including the geometric feature in our message passing payer.

12. Additional Results

12.1. Per-class prediction result

In Tbl. 10, We provide the per entity class recall for the experiment reported in Tbl. 1. Our method has dominant performance on most of the classes regardless of the input segmentation types.

12.2. Without consider *None* estimation

Our evaluations follow the line of work [12, 14] which consider *None* relationship is crucial, unlike other work [2,

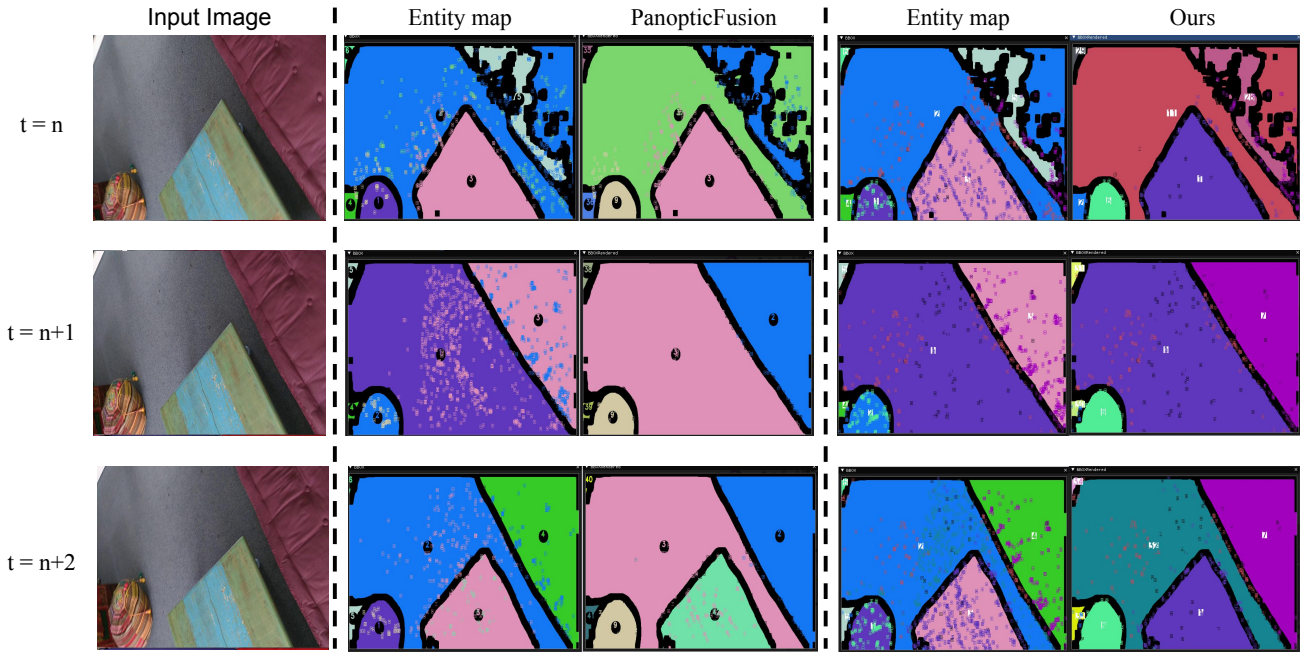

Figure 7. An example of comparing the use of IoU and our label association approach with the sparse input points. Our method handles non-uniformly distributed data better than the IoU-based method.

		bath.	bed	bkshf	cab.	chair	cntr.	curt.	desk	door	floor	ofurn	pic.	refri.	show.	sink	sofa	table	toil.	wall	wind.	avg
GT	IMP	0.000	0.667	0.143	0.562	0.688	0.677	0.712	0.292	0.541	0.957	0.262	0.727	0.000	0.143	0.617	0.579	0.731	0.889	0.856	0.560	0.530
	VGfM	0.750	0.833	**0.286**	0.423	0.854	0.677	0.712	0.042	0.658	**0.969**	0.319	0.693	0.000	0.286	0.633	0.592	0.641	0.852	0.847	0.429	0.575
	3DSSG	0.500	0.333	0.000	0.477	0.838	0.516	0.432	0.417	0.622	0.957	0.288	0.205	0.000	0.143	0.717	0.605	0.563	0.630	0.609	0.357	0.460
	SGFN	1.000	0.833	0.143	0.385	0.696	**0.839**	0.577	0.417	0.631	0.963	0.372	0.830	0.111	0.143	0.783	0.434	0.647	0.593	0.718	0.429	0.577
	Ours	**1.000**	**1.000**	0.000	**0.619**	**0.927**	0.645	**0.847**	**0.583**	**0.838**	**0.969**	**0.539**	**0.943**	**0.667**	**1.000**	**0.900**	**0.671**	**0.808**	**1.000**	**0.877**	**0.655**	**0.774**
Dense	IMP	0.000	0.667	0.000	0.381	0.453	0.000	0.477	0.000	0.081	0.951	0.199	0.023	0.000	0.000	0.200	0.474	0.485	**0.667**	0.770	**0.179**	0.300
	VGfM	0.000	0.667	0.000	0.346	0.494	0.000	0.486	0.042	0.198	0.957	0.141	0.011	0.000	0.000	0.233	**0.579**	0.569	0.630	**0.780**	**0.179**	0.316
	3DSSG	0.250	0.667	0.000	0.200	0.510	**0.258**	**0.505**	0.000	**0.477**	0.914	0.147	0.034	**0.222**	**0.143**	0.250	0.474	0.425	0.259	0.519	0.131	0.319
	SGFN	**0.750**	0.333	0.000	**0.508**	0.636	0.194	0.405	0.083	0.387	**0.969**	0.230	**0.114**	0.111	0.000	**0.383**	0.553	**0.623**	0.519	0.730	0.131	0.383
	Ours	**0.750**	**1.000**	0.000	0.504	**0.656**	0.194	0.459	**0.125**	0.342	**0.969**	**0.251**	0.057	0.000	**0.143**	**0.383**	**0.579**	0.599	**0.667**	0.761	0.155	**0.430**
Sparse	IMP	0.000	**0.333**	0.000	0.235	0.146	**0.129**	0.252	**0.167**	0.099	0.798	0.068	0.023	0.000	**0.286**	0.183	**0.329**	**0.281**	0.259	0.324	**0.214**	0.206
	VGfM	0.000	**0.333**	0.000	**0.273**	0.150	0.097	0.243	0.042	0.081	**0.810**	0.047	0.011	0.000	0.000	0.150	0.276	0.263	0.222	0.325	0.202	0.176
	3DSSG	0.000	0.000	0.000	0.085	0.045	0.000	0.108	0.000	0.063	0.558	0.005	0.011	0.000	0.000	0.017	0.000	0.186	0.000	0.048	0.060	0.059
	SGFN	0.000	0.000	0.000	0.058	0.081	0.032	0.072	0.042	0.063	0.712	0.010	0.034	0.000	0.000	0.083	0.092	0.174	0.000	0.097	0.107	0.083
	Ours	**0.500**	**0.333**	**0.286**	**0.273**	**0.227**	**0.129**	**0.270**	0.042	**0.135**	**0.810**	**0.110**	**0.068**	0.000	0.000	**0.217**	0.211	0.263	**0.296**	**0.342**	0.179	**0.235**

Table 10. The per-class *Recall* of all methods in 3RScan dataset [11] with 20 node classes.

16] which only consider edges with annotated non-*None* relationships. Both approaches have their advantages. Considering *None* estimation prevents excessive relationship estimation while also preventing potential relationship discovery, *e.g.* should exist but was not annotated. For further comparison and the interest of potential readers, we provide the evaluation result without considering the *None* relationship in Tab. 11 and Tab. 12, with the experiment setup as reported in Tbl. 1 and Tbl. 2.

References

[1] Angela Dai, Angel Xuan Chang, Manolis Savva, Maciej Halber, Tom Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. In *CVPR*, 2017. 2

[2] Paul Gay, James Stuart, and Alessio Del Bue. Visual Graphs from Motion (VGfM): Scene Understanding with Object Geometry Reasoning. In *ACCV*. Springer, 2018. 1, 3, 5

[3] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2

[4] Chih-Hui Ho, Bo Liu, Tz-Ying Wu, and Nuno Vasconcelos. Exploit clues from views: Self-supervised and regularized learning for multiview object recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9090–9100, 2020. 1

[5] Ue-Hwan Kim, Jin-Man Park, Taek-Jin Song, and Jong-Hwan Kim. 3-d scene graph: A sparse and semantic representation of physical environments for intelligent agents. *IEEE transactions on cybernetics*, 50(12):4921–4933, 2019. 1, 2

	Method	*Recall*(%)			*mRecall*(%)	
		Rel.	Obj.	Pred.	Obj.	Pred.
GT	IMP [16]	43.5	70.1	54.2	53.0	38.1
	VGfM [2]	49.0	69.4	61.9	57.5	44.6
	3DSSG [12]	40.4	58.0	69.3	46.8	58.7
	SGFN [14]	47.2	63.8	71.4	57.7	65.5
	Ours	**64.8**	**81.2**	**76.8**	**77.4**	**71.5**
Dense	IMP [16]	18.3	51.8	19.3	30.0	23.0
	VGfM [2]	20.8	53.3	22.1	31.6	24.4
	3DSSG [12]	15.1	41.4	26.1	31.9	26.6
	SGFN [14]	24.4	56.7	27.2	38.3	30.5
	Ours	**25.5**	**58.1**	**27.3**	**43.0**	**33.3**
Sparse	IMP [16]	3.5	27.5	4.0	20.6	14.0
	VGfM [2]	3.7	26.9	4.2	17.6	15.4
	3DSSG [12]	1.0	9.7	**9.2**	5.9	15.1
	SGFN [14]	2.3	12.6	7.0	8.3	14.4
	Ours	6.9	29.5	8.1	23.5	**16.5**
	Ours (i)	**7.1**	**30.2**	8.4	**24.5**	15.9

Table 11. We report top-1 recall of method without considering *None* relationship. The experiment setup is the same as in Tbl 1.

Method	*Recall*(%)			*mRecall*(%)	
	Rel.	Obj.	Pred.	Obj.	Pred.
IMP [16]	3.3	35.9	9.0	18.7	4.9
VGfM [2]	3.4	37.9	14.7	17.9	6.5
3DSSG [12]	7.3	29.6	**68.8**	11.7	**25.5**
SGFN [14]	4.5	29.4	45.8	11.8	13.5
Ours	**17.9**	**56.7**	50.4	**27.2**	23.9

Table 12. We report top-1 recall of method without considering *None* relationship. The experiment setup is the same as in Tbl 2.

[6] Kevin Lai, Liefeng Bo, Xiaofeng Ren, and Dieter Fox. A large-scale hierarchical multi-view rgb-d object dataset. In *2011 IEEE international conference on robotics and automation*, pages 1817–1824. IEEE, 2011. 2

[7] Gaku Narita, Takashi Seno, Tomoya Ishikawa, and Yohsuke Kaji. PanopticFusion: Online Volumetric Semantic Mapping at the Level of Stuff and Things. In *IEEE Conf. Intelligent Robots and Syst.*, 2019. 2

[8] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor Segmentation and Support Inference from RGBD Images. In *ECCV*, 2012. 2

[9] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 945–953, 2015. 2

[10] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1588–1597, 2019. 2

[11] Johanna Wald, Armen Avetisyan, Nassir Navab, Federico Tombari, and Matthias Nießner. RIO: 3D Object Instance Re-Localization in Changing Indoor Environments. In *ICCV*, 2019. 1, 4

[12] Johanna Wald, Helisa Dhamo, Nassir Navab, and Federico Tombari. Learning 3D Semantic Scene Graphs from 3D Indoor Reconstructions. In *CVPR*, 2020. 1, 2, 3, 5

[13] Xin Wei, Yifei Gong, Fudong Wang, Xing Sun, and Jian Sun. Learning canonical view representation for 3d shape recognition with arbitrary views. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 407–416, 2021. 1, 2

[14] Shun-Cheng Wu, Johanna Wald, Keisuke Tateno, Nassir Navab, and Federico Tombari. Scenegraphfusion: Incremental 3d scene graph prediction from rgb-d sequences. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7515–7525, 2021. 1, 3, 5

[15] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. 2

[16] Danfei Xu, Yuke Zhu, Christopher Choy, and Li Fei-Fei. Scene Graph Generation by Iterative Message Passing. In *CVPR*, 2017. 1, 3, 5

[17] Ze Yang and Liwei Wang. Learning relationships for multi-view 3d object recognition. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7505–7514, 2019. 1