

MSG-Loc: Multi-Label Likelihood-based Semantic Graph Matching for Object-Level Global Localization

Gihyeon Lee¹, Jungwoo Lee¹, Juwon Kim¹, Young-Sik Shin^{2†}, and Younggun Cho^{1†}

Abstract—Robots are often required to localize in environments with unknown object classes and semantic ambiguity. However, when performing global localization using semantic objects, high semantic ambiguity intensifies object misclassification and increases the likelihood of incorrect associations, which in turn can cause significant errors in the estimated pose. Thus, in this letter, we propose a multi-label likelihood-based semantic graph matching framework for object-level global localization. The key idea is to exploit multi-label graph representations, rather than single-label alternatives, to capture and leverage the inherent semantic context of object observations. Based on these representations, our approach enhances semantic correspondence across graphs by combining the likelihood of each node with the maximum likelihood of its neighbors via context-aware likelihood propagation. For rigorous validation, data association and pose estimation performance are evaluated under both closed-set and open-set detection configurations. In addition, we demonstrate the scalability of our approach to large-vocabulary object categories in both real-world indoor scenes and synthetic environments.

Index Terms—Semantic Scene Understanding, Localization, Graph Matching, Object-based SLAM.

I. INTRODUCTION

GLOBAL localization on a prior map, which estimates the pose of a robot from current sensor observations without any prior pose information, is an essential task in robotics [1]. This capability is important for simultaneous localization and mapping (SLAM) tasks such as relocalization after tracking failures and loop closures [2, 3]. While recent feature-based methods [4] improve robustness to viewpoint and appearance changes, their performance still degrades under severe

Manuscript received: June 30, 2025; Accepted: November 18, 2025. This letter was recommended for publication by Editor Sven Behnke upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (RS-2022-II220448), National Research Foundation of Korea (NRF) grant (RS-2025-02217000 and RS-2025-24803365), Industrial Strategic Technology Development Program (no.20018745) funded by the Ministry of Trade, Industry & Energy (MOTIE, Korea), and National Research Council of Science & Technology as part of the project titled “Development of Core Technologies for Robot General Purpose Task Artificial Intelligence Framework” (NK254G). (Corresponding Authors: Young-Sik Shin and Younggun Cho)

¹Gihyeon Lee, ¹Jungwoo Lee, ¹Juwon Kim, and ^{1†}Younggun Cho are with the Electrical and Computer Engineering, Inha University, Incheon, South Korea (e-mail: leekh951@inha.edu, pihsdneirf@inha.edu, marimo117@inha.edu, yg.cho@inha.ac.kr).

^{2†}Young-Sik Shin is with Depart. of AI Machinery, Korea Institute of Machinery and Materials, Daejeon, South Korea (e-mail: yshin86@kimm.re.kr).

Our supplementary materials and code are available at <https://sparolab.github.io/research/msg-loc/>.

Digital Object Identifier (DOI): see top of this page.

©2026 IEEE

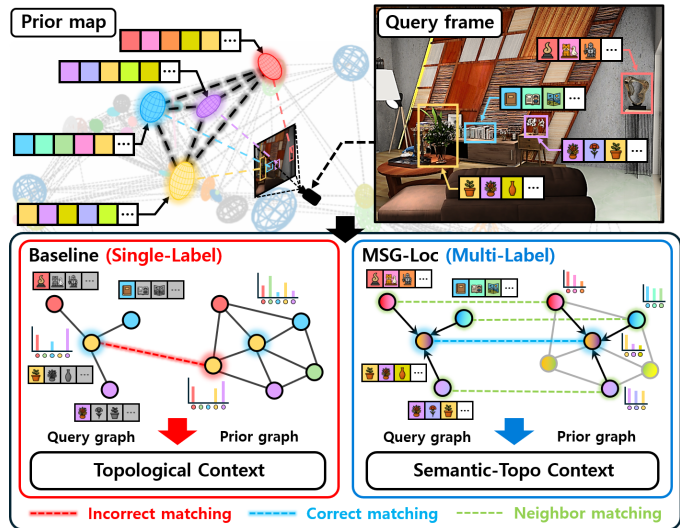


Fig. 1. **MSG-Loc**. Our method preserves multi-label hypotheses for both landmarks and observed objects, and estimates their likelihoods. It propagates context-aware likelihoods and calculates similarity across the semantic graph to mitigate viewpoint-dependent misclassifications. This enables reliable object-level global localization even in scenes of severe semantic ambiguity.

variations. In contrast, object-level approaches maintain more reliable global localization in such conditions.

However, most existing object-based methods rely on single-label predictions, neglecting semantic uncertainty and label inconsistency prevalent in practical scenarios [5]. In particular, the single-label assumption has limitations in open-set environments where unknown objects are present. In this context, prior studies [6–11] have introduced semantic graph-based matching methods leveraging relationships between adjacent objects, but struggle to capture semantic variations of individual objects induced by partial observations or viewpoint changes. For instance, the graph kernel-based approach [11] constructs kernel descriptors that account for object detection uncertainty in the prior graph. Nevertheless, since it assigns only a single label per object in query frames, it is limited in handling the semantic ambiguity inherent in open-set observations. Additionally, fixed descriptors derived from sparse object distributions are sensitive to semantic inconsistencies, often leading to matching errors.

Recently, vision-language model (VLM)-based methods [12, 13] have been explored to recognize out-of-distribution classes. However, these approaches have limited applicability in real-world scenarios as they assign to each landmark a single text label, which is typically generated manually or

via offline captioning. Therefore, it is necessary to develop flexible subgraph matching approaches capable of integrating node-level uncertainty and adapting to sparse and ambiguous observations in open-set environments.

In this letter, we propose *MSG-Loc*, a novel semantic graph matching framework based on multi-label likelihoods for global localization, as shown in Fig. 1. The key idea is to preserve multi-label hypotheses observed from multiple viewpoints during map construction and to exploit these hypotheses during querying, thereby mitigating the semantic ambiguity inherent to single-label predictions. To this end, we define a multi-label likelihood for each node-pair and leverage maximum-likelihood propagation from neighboring matches, which leads to more reliable candidate selection for global localization. The main contributions of this work are summarized as follows:

- **Semantic Disambiguation:** To the best of our knowledge, this is the first approach to capture semantic uncertainty and consistency by exploiting multi-label likelihoods to estimate inter-object correspondences.
- **Context-Aware Stability:** The proposed method propagates multi-label likelihoods via 1-hop neighbors in the semantic graph, enabling stable data association and pose estimation under sparse and ambiguous observations.
- **Scalability to Large Class Sets:** Our method achieves reliable performance over a wide vocabulary of object classes, as the multi-label likelihood is not constrained by the size or composition of the label set.
- **Cross-Paradigm Compatibility:** The proposed method shows effectiveness across multiple paradigms, including open-set and closed-set detection as well as zero-shot and supervised classification tasks.

II. RELATED WORK

A. Object-based SLAM

Unlike traditional feature-based methods [2], object-based approaches leverage high-level semantics to enhance localization accuracy and robustness. In semantic SLAM frameworks, accurate object representation and recognition form the basis for building semantically consistent maps.

SLAM++ [14] introduced an object-oriented SLAM approach that utilizes known object models. However, this approach lacks generalization to unknown objects. To address this limitation, CubeSLAM [15] proposed monocular object SLAM by estimating 3D cuboids directly from 2D bounding boxes. Yang and Scherer [16] further enhanced this method by integrating planar surfaces with cuboid objects, resulting in improved localization accuracy and richer environmental representation.

In recent years, dual quadrics have emerged as a general and compact representation of objects within SLAM frameworks. Since QuadricSLAM [17] formalized dual quadrics as compact 3D landmarks, subsequent works have adopted this representation for object-level SLAM owing to its analytic projection model and efficient parameterization. For example, EAO-SLAM [18] proposed a semi-dense object-based framework that leverages an ensemble strategy with statistical tests

to improve data association. QISO-SLAM [19] introduced a more precise observation model by incorporating instance segmentation-derived contour points. Recently, VOOM [20] proposed a hierarchical landmark representation combining high-level objects and low-level points, demonstrating superior localization performance over traditional object-based and feature-based methods.

B. Semantic Graph Matching for Global Localization

Despite these advancements, previous approaches focusing on dual quadric-based mapping and localization have not adequately addressed crucial aspects such as relocalization and loop detection, especially after tracking loss. To address these limitations, OA-SLAM [3] fused sparse points with ellipsoidal object models, enabling camera relocalization through object-wise matching from viewpoints where the point-only baseline [2] fails.

Beyond this relocalization method, graph-based approaches incorporating topological information have gained prominence for reliable localization in object-level maps. Guo et al. [10] introduced Semantic Histogram (SH), a histogram-of-label-paths descriptor computed on semantic graphs that captures higher-order context and enables faster graph matching than the random-walk approach [9]. Additionally, Zhou et al. [6] developed directional histogram descriptors tailored for object-based loop closure, demonstrating reliable performance even with extensive viewpoint differences. Wu et al. [8] expanded the object map proposed in EAO-SLAM [18] into a topological map framework, achieving multi-scene matching and semantic descriptor-driven relocalization. SemanticTopoLoop [7] proposed multi-level verification strategies, combining object data association and quadric-level topological maps for semantic loop closure, effectively reducing false positives.

However, these graph-based methods have mainly relied on topological context, overlooking the semantic ambiguity inherent to each object. To address this limitation, GOREloc [11] integrated the graph kernel-based matching strategy of Stumm et al. [21] with label statistics accumulated across keyframes. Nevertheless, the aforementioned methods still depend on single-label predictions per object in query frames, failing to fully capture semantic uncertainty at inference time.

In contrast to previous methods, we propose a multi-label likelihood-based graph matching framework that incorporates both the semantic context of candidate labels and the topological context of the graph to establish object correspondences. To estimate the multi-label likelihood, we employ the multi-label detection frequencies from the map and the normalized confidence distributions from the input frame.

III. MSG-LOC

A. Problem Definition

Our objective is to match a query graph \mathcal{G}^q to its corresponding subgraph in a prior graph \mathcal{G}^p for global localization, as illustrated in Fig. 2, where the superscripts q and p denote the query and prior domains, respectively.

For this purpose, the Prior Graph Management (PGM) and Query Graph Management (QGM) modules construct prior

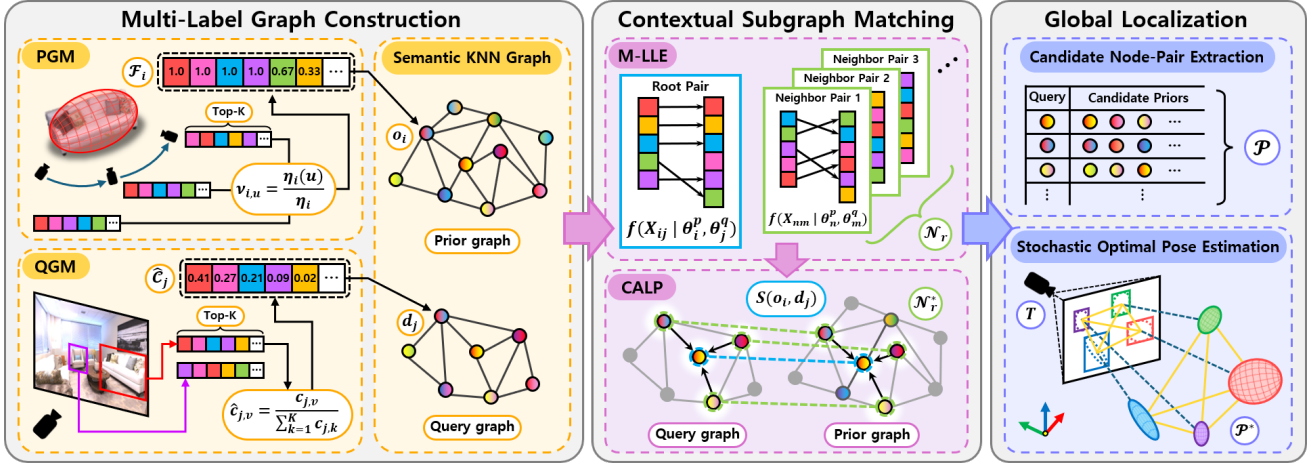


Fig. 2. **Overview of our method.** In the Multi-Label Graph Construction phase, we generate a k-nearest neighbor (KNN)-based semantic graph, embedding multi-label detection frequencies and normalized confidence scores via the Prior Graph Management (PGM) and Query Graph Management (QGM) modules. Multi-Label Likelihood Estimation (M-LLE) then computes semantic likelihoods from these attributes for both landmarks and observations. Context-Aware Likelihood Propagation (CALP) propagates the maximum likelihood from neighboring nodes to the root node, calculating its similarity score to ensure contextual consistency. Finally, candidate node-pairs are extracted, and the stochastic optimal pose estimation module determines the camera pose.

nodes \mathcal{O}^p and query nodes \mathcal{D}^q from synchronized RGB-D input, each characterized by multi-label attributes and defined in 3D Euclidean space. Each i^{th} prior object node $o_i \in \mathcal{O}^p$ contains dual quadric parameters $\mathbf{q}_i^p = [\mathbf{p}_i^p, \mathbf{r}_i, \mathbf{s}_i]$ and multi-label detection frequencies \mathcal{F}_i , where \mathbf{p}_i^p , \mathbf{r}_i , and \mathbf{s}_i denote the position, rotation, and scale vectors, respectively. For each input frame, the j^{th} query node $d_j \in \mathcal{D}^q$ comprises a bounding box B_j from the object detector, a normalized multi-label confidence $\hat{\mathcal{C}}_j$ derived from the classifier output \mathcal{C}_j , and a 3D position \mathbf{p}_j^q estimated using B_j and the depth image.

Using these graph formulations, candidate node-pairs \mathcal{P} are extracted from \mathcal{G}^p and \mathcal{G}^q by integrating topological structure and multi-label semantic features through the Multi-Label Likelihood Estimation (M-LLE) and Context-Aware Likelihood Propagation (CALP) modules. Finally, the camera pose T is estimated by iteratively determining the optimal node correspondences \mathcal{P}^* .

B. Graph Construction via Multi-Label Prediction

By utilizing multi-label rather than single-label statistics, the prior graph forms the basis for stable and consistent object descriptions, even in the presence of transient detection errors and viewpoint variations. Likewise, by adopting multi-label confidence distributions, the query graph captures latent semantic labels beyond what single-label approaches represent at the object level.

1) *Prior Graph Management*: Across multiple keyframes, the prior graph \mathcal{G}^p preserves semantic consistency of individual objects by assigning accumulated multi-label detection counts to the node set \mathcal{O}^p . Each prior object node $o_i = \{\mathbf{q}_i^p, \mathcal{F}_i\}$ is represented by \mathbf{q}_i^p , the optimized 3D object. Specifically, the set $\mathcal{F}_i = \{(l_{i,u}^p, \nu_{i,u})\}_{u=1}^U$ is constructed by aggregating top- K predicted labels across keyframes, where U denotes the number of unique labels of o_i . The per-label detection frequency $\nu_{i,u}$ of the u^{th} unique label $l_{i,u}^p$ is defined relative to all detections of o_i , as follows:

$$\nu_{i,u} = \frac{\eta_i(u)}{\eta_i}, \quad (1)$$

where η_i is the total number of detections of o_i , and $\eta_i(u)$ counts detections where $l_{i,u}^p$ appears among the predicted labels. After node generation, the prior graph $\mathcal{G}^p = (\mathcal{O}^p, \mathcal{E}^p)$ is constructed by merging, for each keyframe, all k-nearest neighbor (KNN) edges \mathcal{E}^p formed based on \mathbf{p}_i^p .

2) *Query Graph Management*: For each input frame, the query graph encapsulates semantic uncertainty by assigning normalized top- K multi-label confidence distributions to each detected object node in the set \mathcal{D}^q . Each query detection node $d_j = \{B_j, \mathbf{p}_j^q, \hat{\mathcal{C}}_j\}$ is represented as a bounding box B_j with the position vector \mathbf{p}_j^q . Notably, the set $\hat{\mathcal{C}}_j = \{(l_{j,v}^q, \hat{c}_{j,v})\}_{v=1}^K$ contains the v^{th} top label $l_{j,v}^q$ and its corresponding normalized confidence score $\hat{c}_{j,v}$, where K is the number of top-ranked labels retained for each detection. The normalization process is defined as follows:

$$\hat{c}_{j,v} = \frac{c_{j,v}}{\sum_{k=1}^K c_{j,k}}, \quad (2)$$

where $c_{j,v}$ denotes the raw confidence score for $l_{j,v}^q$, obtained from $(l_{j,v}^q, c_{j,v}) \in \mathcal{C}_j$. After generating the nodes, the query graph $\mathcal{G}^q = (\mathcal{D}^q, \mathcal{E}^q)$ is constructed by forming KNN edges \mathcal{E}^q based on \mathbf{p}_j^q .

C. Multi-Label Likelihood Estimation

Ambiguous detections often assign non-negligible confidence scores to several conceptually similar labels for a single object, such as ‘cup’, ‘mug’, and ‘glass’. Constraining multi-label candidates to a single deterministic label overlooks this semantic uncertainty and often results in mismatches between visually similar objects. So, we focus on leveraging the semantic context among candidate labels as a key feature for object matching, as illustrated in Fig. 3.

Under the assumption of pairwise independence for tractability, we formulate the multi-label likelihood in terms of empirical evidence $(l_{i,u}^p, \nu_{i,u})$ and perceptual cues $(l_{j,v}^q, \hat{c}_{j,v})$ as follows:

$$f(X_{ij} | \theta_i^p, \theta_j^q) \equiv \sum_{u=1}^U \sum_{v=1}^K \mathbb{I}(l_{i,u}^p = l_{j,v}^q) \cdot \nu_{i,u} \cdot \hat{c}_{j,v}, \quad (3)$$

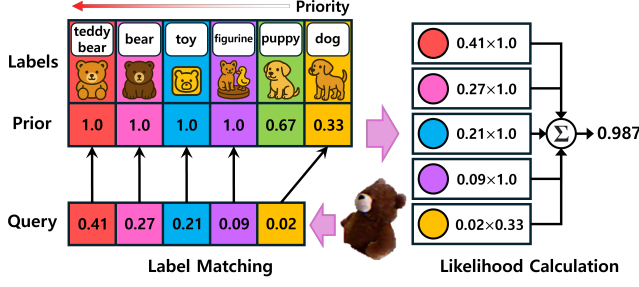


Fig. 3. **Multi-label likelihood estimation algorithm.** This figure illustrates how M-LLE computes semantic likelihoods between landmark and observation nodes by aggregating detection frequencies and normalized confidence scores. The probabilistic estimation of likelihoods enables data association that is agnostic to the object classification paradigm within our framework.

where θ_i^p , θ_j^q , and X_{ij} denote parameters associated with \mathcal{F}_i , $\hat{\mathcal{C}}_j$, and a random variable representing the correspondence between o_i and d_j , respectively. The indicator function $\mathbb{I}(\cdot)$ returns 1 if the given condition is true, and 0 otherwise. That is, the indicator function restricts the summation to only those label pairs that are identical, so that the likelihood reflects the degree of alignment between the empirical and perceptual label distributions.

Importantly, since M-LLE does not rely on a fixed set of class categories, it is not constrained by the size or composition of the label vocabulary. As a result, it is designed to be detector-agnostic and remains applicable to both open-set and closed-set configurations.

D. Context-Aware Likelihood Propagation

While the M-LLE evaluates node-pairs independently for efficiency, it may overlook contextual correlations across neighboring objects. Thus, it is necessary to enable correct pair selection by jointly considering semantic alignment and topological consistency across the entire graph via CALP.

First, we define the set of neighboring node-pairs $\mathcal{N}_r = \{(o_n, d_m)\}$ for each root node-pair (o_i, d_j) , where o_n and d_m denote the n^{th} and m^{th} 1-hop neighbors of o_i and d_j in the prior and query graphs, respectively. To reflect the similarity of relative distances to the root node in both graphs, the corresponding weight w_{nm} is defined as:

$$w_{nm} = \frac{1}{1 + |\delta_n^p - \delta_m^q|}, \quad (4)$$

where $\delta_n^p = \|\mathbf{p}_i^p - \mathbf{p}_n^p\|_2$ and $\delta_m^q = \|\mathbf{p}_j^q - \mathbf{p}_m^q\|_2$ are the Euclidean distances from the root node to each neighbor node in the prior and query graphs, respectively. The vectors \mathbf{p}_n^p and \mathbf{p}_m^q represent the positions of nodes o_n and d_m .

Then, the best neighboring node-pair set \mathcal{N}_r^* is selected from \mathcal{N}_r by maximizing the weighted neighbor likelihood for each neighboring query node as follows:

$$\mathcal{N}_r^* = \bigcup_{m=1}^M \arg \max_{(o_n, d_m) \in \mathcal{N}_r} w_{nm} f(X_{nm} | \theta_n^p, \theta_m^q), \quad (5)$$

where M denotes the number of neighboring nodes connected to the root node in the query graph. The weighted likelihoods

Algorithm 1 Candidate Extraction and Pose Estimation

Require: Query graph \mathcal{G}^q , prior graph \mathcal{G}^p , candidate pair number τ , iteration number N_{iter}

Ensure: Estimated camera pose T , correspondence set \mathcal{P}^*

- 1: # M-LLE and CALP are described in Sec. III-C and III-D
- 2: $\mathcal{P} \leftarrow \text{GetCandidates}(\text{CALP}(\text{M-LLE}(\mathcal{G}^p, \mathcal{G}^q)), \tau)$ using Eq. (9)
- 3: $(\mathcal{W}, T, \mathcal{P}^*) \leftarrow (0, \text{None}, \text{None})$
- 4: **for** iter = 1 **to** N_{iter} **do**
- 5: $\mathcal{P}_{\text{samp}} \leftarrow \text{RandomSample}(\mathcal{P}, 3)$
- 6: **if not** $\text{IsValidSample}(\mathcal{P}_{\text{samp}})$ **then**
- 7: **continue**
- 8: **end if**
- 9: $\mathcal{T}_{\text{tmp}} \leftarrow \text{P3P}(\mathcal{P}_{\text{samp}})$
- 10: **for** $T_{\text{tmp}} \in \mathcal{T}_{\text{tmp}}$ **do**
- 11: $\mathcal{W}_{\text{tmp}}, \mathcal{P}_{\text{tmp}}^* \leftarrow \text{CalculateWAS}(T_{\text{tmp}}, \mathcal{P})$ using Eq. (9)
- 12: **if** $\mathcal{W}_{\text{tmp}} > \mathcal{W}$ **then**
- 13: $(\mathcal{W}, T, \mathcal{P}^*) \leftarrow (\mathcal{W}_{\text{tmp}}, T_{\text{tmp}}, \mathcal{P}_{\text{tmp}}^*)$
- 14: **end if**
- 15: **end for**
- 16: **end for**
- 17: **return** T, \mathcal{P}^*

of the neighbors identified by \mathcal{N}_r^* are propagated to the root node-pair, and the similarity score is calculated as follows:

$$S(o_i, d_j) = f(X_{ij} | \theta_i^p, \theta_j^q) + \frac{1}{|\mathcal{N}_r^*|} \sum_{(o_n, d_m) \in \mathcal{N}_r^*} w_{nm} f(X_{nm} | \theta_n^p, \theta_m^q). \quad (6)$$

This process suppresses the influence of local outliers and improves matching uniformity between the sparse query graph and the dense prior graph. That is, it alleviates the limitations of the subgraph extraction method proposed in previous work [11], which is sensitive to structural discrepancies.

E. Global Localization

1) *Candidate Node-pair Extraction:* In large-scale prior maps, the extraction of candidate nodes reduces complexity and mitigates outlier correspondences, enhancing the stability of alignment-based pose estimation. We extract the top- τ prior nodes for each query node using the propagated similarity scores from CALP, and generate the candidate node-pair set \mathcal{P} as follows:

$$\mathcal{P} = \{(o_i, d_j) \mid \text{rank}(S(o_i, d_j)) \leq \tau\}, \quad (7)$$

where $\text{rank}(\cdot)$ selects the top τ prior nodes o_i corresponding to d_j based on $S(o_i, d_j)$, thereby forming a reliable subgraph.

2) *Stochastic Optimal Pose Estimation:* We adopt a stochastic optimal iterative approach inspired by [11], as outlined in Algorithm 1. At each of the N_{iter} iterations, three unique node-pairs $\mathcal{P}_{\text{samp}}$ are randomly sampled from \mathcal{P} . Both subgraphs induced by this sampled set $\mathcal{P}_{\text{samp}}$ must preserve identical edge connectivity, and previously used sampled sets are never reused. Given $\mathcal{P}_{\text{samp}}$ that satisfies the required conditions, provisional camera poses \mathcal{T}_{tmp} are then estimated by solving the Perspective-3-Point (P3P) problem. Subsequently, for each $(o_i, d_j) \in \mathcal{P}$, the dual quadric \mathbf{q}_i^p of o_i is transformed by each $T_{\text{tmp}} \in \mathcal{T}_{\text{tmp}}$ to generate a projected prior bounding

box B_i . Both B_i and the corresponding query bounding box B_j are modeled as Gaussian distributions, denoted as $D_i = G(\mu_i, \Sigma_i)$ and $D_j = G(\mu_j, \Sigma_j)$, respectively. The degree of alignment between D_i and D_j is measured by the normalized Wasserstein distance proposed by [22] as follows:

$$W_n(D_i, D_j) = \exp\left(-\frac{\sqrt{W_2^2(D_i, D_j)}}{C}\right), \quad (8)$$

where $W_2^2(D_i, D_j)$ is the 2nd order Wasserstein distance between two Gaussians, and C is a scale factor. The best correspondences $\mathcal{P}^* = \{(o_i, d_j)^*\}$ are selected from \mathcal{P} by maximizing $W_n(D_i, D_j)$ for each query node. The Wasserstein Alignment Score (WAS) is then calculated as follows:

$$\mathcal{W} = \frac{1}{|\mathcal{P}^*|} \sum_{(o_i, d_j)^* \in \mathcal{P}^*} W_n(D_i, D_j). \quad (9)$$

After all iterations, the algorithm returns the camera pose T and \mathcal{P}^* that achieve the maximum \mathcal{W} as the final solution.

IV. EXPERIMENTS

A. Setup

1) *Datasets*: In this section, we evaluate the proposed method on real-world TUM RGB-D [23] and synthetic ICL-LM [24] public benchmarks, as well as a custom resort-interior dataset captured in the real world. From TUM RGB-D, we use the ‘Fr2_desk’ and ‘Fr2_person’ sequences, which share similar layouts, with ‘Fr2_person’ including dynamic factors. In the ICL-LM dataset, we select the ‘Walk’, ‘Head’, ‘VR’, ‘DJI’, and ‘Ground’ sequences. These sequences involve complex object layouts with ambiguous classes, captured by diverse platforms in a virtual indoor setting. The custom resort dataset includes rooms where many objects are partially occluded by walls. It is divided into ‘Resort_1’ and ‘Resort_2’, which feature stable and dynamic camera movements, respectively. Object-level maps are built from ‘Fr2_desk’, ‘Walk’, and ‘Resort_1’, with data association and pose estimation evaluated on all sequences.

2) *Method Configurations*: We adopt SH [10] and GOREloc [11] as baselines for comparison. Both methods rely on single-label, graph-based descriptors. Experiments were conducted under the following configurations:

- **Closed-set**: We utilize a YOLOv8 [25] model trained on COCO [26] and LVIS [27] datasets to detect and classify objects within a predefined set of categories. This configuration is denoted as Y.
- **Open-set**: We employ Grounding DINO [28] for zero-shot object detection, and integrate foundation models such as Open-Vocabulary SAM (OVSAM) [29] and Tokenize Anything via Prompting (TAP) [30] for zero-shot classification. These configurations are denoted as G+O and G+T, respectively.

3) *Implementation Details*: To construct prior maps, we adopt an object-based SLAM framework [17] that leverages 2D bounding-box detections. We evaluate all methods on COCO (80 classes) and LVIS (1203 classes) categories. SH is evaluated on COCO only due to computational constraints.

Detection confidence thresholds are set to 0.1 for YOLOv8 and 0.14 for Grounding DINO, with non-maximum suppression applied at an overlap ratio above 0.6. Hyperparameters are determined empirically based on preliminary experiments: we set $K = 5$ in Eq. (2), $\tau = 3$ in Eq. (7), and $C = 100$ in Eq. (8) for all experiments unless otherwise specified.

B. Evaluation Metrics

1) *Precision, Recall, and F1-score*: We evaluate multi-object data association by projecting associated objects onto the image plane based on the ground-truth pose. Precision, Recall, and F1-score are defined as:

$$P = \frac{TP}{TP + FP}, \quad R = \frac{TP}{TP + FN}, \quad F_1 = \frac{2 \cdot P \cdot R}{P + R}, \quad (10)$$

where TP, FP, and FN denote true positives, false positives, and false negatives, respectively.

2) *Multiple Object Tracking Accuracy*: To evaluate the temporal consistency of data association across frames, we employ multiple object tracking accuracy (MOTA):

$$\text{MOTA} = 1 - \frac{\sum_t (\text{FN}_t + \text{FP}_t + \text{IDS}_t)}{\sum_t \text{GT}_t}, \quad (11)$$

where FN_t , FP_t , IDS_t , and GT_t denote false negatives, false positives, ID switches, and ground-truth detections at frame t . The association cost is defined in Eq. (8).

3) *Translation Error and Success Rate*: For pose estimation, we use translation error (TE):

$$\text{TE} = \|\hat{t} - t\|_2, \quad (12)$$

where \hat{t} and t are the estimated and ground-truth positions, respectively. The mean TE, calculated over all success frames, is denoted as $\overline{\text{TE}}$. The success rate (SR) is the fraction of poses with TE below a threshold ($@\text{TE}$): SR_{succ} for success frames, SR_{all} for all frames.

4) *Shannon Entropy*: To quantify frame-wise semantic uncertainty, we compute Shannon entropy for each object:

$$H(\hat{\mathcal{C}}_j) = -\sum_{v=1}^K \hat{c}_{j,v} \log \hat{c}_{j,v}, \quad (13)$$

where $\hat{c}_{j,v}$ is the normalized confidence score corresponding to the label $l_{j,v}^q$ in $\hat{\mathcal{C}}_j$, as detailed in Sec. III-B.

C. Data Association

1) *Semantic Disambiguation*: As shown in Table I, MSG-Loc consistently outperforms SH and GOREloc in both F_1 and MOTA. Averaged over all sequences under both Y and G+T configurations, MSG-Loc achieves relative improvements of 10.6% in F_1 and 33.4% in MOTA compared to the best baseline. Interestingly, while MSG-Loc maintains stable performance across all configurations, the baselines tend to degrade under the open-set configuration. As illustrated in Fig. 4, the open-set configuration distributes confidence across multiple labels, reflecting broader semantic uncertainty. Consequently, single-label approaches exhibit increased noise due to label inconsistency. In contrast, the proposed method

TABLE I
PERFORMANCE COMPARISON ON COCO CATEGORIES

Sequence	Method	$F_1 \uparrow$	MOTA \uparrow	SR _{succ} [%] \uparrow			\overline{TE} [m] \downarrow
				@0.5 m	@1 m	@2 m	
Fr2_desk	SH (Y)	0.905	0.786	75.08	80.30	86.13	0.674
	GOREloc (Y)	0.882	0.739	69.34	75.73	81.65	0.845
	MSG-Loc (Y)	0.942	0.873	88.85	91.96	95.64	0.330
	SH (G+T)	0.898	0.774	77.34	81.14	87.00	0.655
	GOREloc (G+T)	0.899	0.771	79.01	83.61	88.27	0.581
	MSG-Loc (G+O)	0.942	0.868	88.25	91.94	96.27	0.336
Walk	MSG-Loc (G+T)	0.957	0.909	94.37	95.99	97.46	0.231
	SH (Y)	0.847	0.633	55.11	61.97	65.85	3.039
	GOREloc (Y)	0.861	0.668	51.95	61.98	69.53	2.623
	MSG-Loc (Y)	0.957	0.882	80.16	87.98	90.43	0.922
	SH (G+T)	0.801	0.541	40.57	51.58	58.12	4.610
	GOREloc (G+T)	0.831	0.595	45.64	54.05	60.48	4.254
Resort_1	MSG-Loc (G+O)	0.938	0.829	69.92	83.14	88.91	1.240
	MSG-Loc (G+T)	0.940	0.860	80.62	93.86	95.85	0.595
	SH (Y)	0.909	0.786	65.32	74.19	81.45	1.001
	GOREloc (Y)	0.932	0.833	61.65	81.20	88.72	0.835
	MSG-Loc (Y)	0.959	0.907	74.38	89.16	95.57	0.514
	SH (G+T)	0.751	0.472	33.33	46.22	52.94	2.482
Resort_2	GOREloc (G+T)	0.731	0.432	32.85	44.19	52.03	2.636
	MSG-Loc (G+O)	0.899	0.780	78.85	88.93	93.28	0.578
	MSG-Loc (G+T)	0.945	0.880	76.83	84.75	88.91	0.789

Bold: Best performance.




Fig. 4. **Frame-wise semantic uncertainty analysis.** Shannon entropy is calculated for each object using the normalized confidence scores of its top-5 predicted LVIS [27] labels. The closed-set detector (YOLOv8 [25]) consistently produces low entropy, as it tends to allocate most of its confidence to a single predicted label. Conversely, the open-set detector with a zero-shot classifier (TAP [30]) produces a confidence distribution spanning multiple labels, resulting in higher entropy.

suppresses such noise by leveraging both semantic and topological contexts, suggesting that the multi-label likelihood can effectively handle semantic uncertainty regardless of the underlying classification paradigm.

2) *Category Scalability:* Under COCO-to-LVIS scaling, Table II indicates that MSG-Loc improves MOTA by 9.3% on average across all sequences when transitioning from closed-set (Y) to open-set (G+T). In contrast, GOREloc shows reduced association performance in most sequences under the same transition. As category diversity increases, single-label approaches tend to be more susceptible to misclassifications. The proposed multi-label approach preserves a broader set of semantic candidates, reducing misclassification and improving inter-object discriminability. Additionally, MSG-Loc (G+O) generally achieves lower performance than MSG-Loc (G+T) across both COCO and LVIS categories, suggesting that multi-label predictions generated by OVSAM may be less stable than those produced by TAP.

Fig. 5(a) qualitatively illustrates the data association performance of the proposed method. In cluttered indoor scenarios, MSG-Loc (G+T) successfully matches multiple ambiguous objects, whereas the baseline occasionally confuses similar ones, resulting in matching errors.

TABLE II
PERFORMANCE COMPARISON ON LVIS CATEGORIES

Sequence	Method	$F_1 \uparrow$	MOTA \uparrow	SR _{succ} [%] \uparrow			\overline{TE} [m] \downarrow
				@0.5 m	@1 m	@2 m	
Fr2_desk	GOREloc (Y)	0.951	0.895	86.36	89.99	93.69	0.429
	MSG-Loc (Y)	0.970	0.936	96.23	97.40	98.19	0.211
	GOREloc (G+T)	0.906	0.790	74.13	80.53	87.49	0.673
	MSG-Loc (G+O)	0.918	0.822	91.34	96.56	98.84	0.245
	MSG-Loc (G+T)	0.976	0.947	97.40	98.28	99.58	0.159
	Fr2_person	GOREloc (Y)	0.894	0.776	57.13	73.90	81.71
MSG-Loc (Y)		0.875	0.730	58.84	74.27	78.56	1.044
GOREloc (G+T)		0.754	0.448	38.73	45.06	54.12	2.304
MSG-Loc (G+O)		0.850	0.631	62.16	68.73	73.81	1.079
MSG-Loc (G+T)		0.901	0.773	76.07	79.00	82.54	0.777
Walk		GOREloc (Y)	0.886	0.751	63.47	77.70	86.78
	MSG-Loc (Y)	0.926	0.838	78.10	90.84	95.33	0.576
	GOREloc (G+T)	0.914	0.809	64.51	74.35	80.07	2.385
	MSG-Loc (G+O)	0.951	0.885	72.34	86.04	94.82	0.716
	MSG-Loc (G+T)	0.961	0.906	84.82	94.24	98.25	0.426
	Head	GOREloc (Y)	0.789	0.575	29.01	44.17	58.02
MSG-Loc (Y)		0.895	0.774	42.58	59.32	73.13	2.403
GOREloc (G+T)		0.775	0.548	20.84	33.63	43.96	6.537
MSG-Loc (G+O)		0.872	0.702	27.98	46.49	63.10	3.895
MSG-Loc (G+T)		0.913	0.791	44.64	63.63	76.57	2.493
Resort_1		GOREloc (Y)	0.828	0.624	61.83	77.75	89.23
	MSG-Loc (Y)	0.883	0.759	80.43	93.99	96.32	0.432
	GOREloc (G+T)	0.893	0.744	57.58	69.70	78.45	1.190
	MSG-Loc (G+O)	0.952	0.895	81.58	90.50	95.45	0.488
	MSG-Loc (G+T)	0.953	0.900	88.01	94.51	95.73	0.398
	Resort_2	GOREloc (Y)	0.820	0.620	58.27	81.77	90.04
MSG-Loc (Y)		0.815	0.604	62.03	77.82	88.91	0.882
GOREloc (G+T)		0.634	0.262	19.93	34.78	42.03	3.273
MSG-Loc (G+O)		0.832	0.648	54.46	65.10	73.76	1.526
MSG-Loc (G+T)		0.867	0.724	54.28	61.95	71.98	1.614

D. Pose Estimation

1) *Pose Accuracy:* As shown in Tables I and II, MSG-Loc demonstrates higher SR_{succ} and lower \overline{TE} than the baselines in most cases, including complex scenes. In the ‘Fr2_desk’, ‘Walk’, and ‘Resort_1’ sequences, our method achieves a success rate above 70% at the 0.5 m threshold and \overline{TE} below 1 m across both categories. The baselines construct their descriptors from all edge-connected neighbors using a single-label representation, which can introduce noise from misclassifications during matching. In contrast, the proposed method selectively integrates only the most probable neighboring node-pairs, indirectly guiding the selection of objects with minimal geometric error relative to the observations.

2) *Open-set Stability of Localization:* GOREloc shows a large performance gap between the closed-set and open-set configurations across most sequences. These fluctuations stem from the graph kernel-based approach, which amplifies mismatches under label noise. This effect becomes more evident when partial observations are frequent and object detections are sparse. In contrast, MSG-Loc maintains comparatively stable pose estimation performance across diverse sequences. Notably, MSG-Loc (G+T) achieves additional improvements in all metrics when scaling from COCO to LVIS.

To qualitatively demonstrate these advantages, Fig. 5(b) visualizes pose estimation results across three complete sequences. Gray lines denote frames that failed due to sparse detections or invalid correspondence graphs, while dense scatter points indicate a high rate of successful and precise localization. This highlights the robustness of MSG-Loc under various conditions, including category expansion.

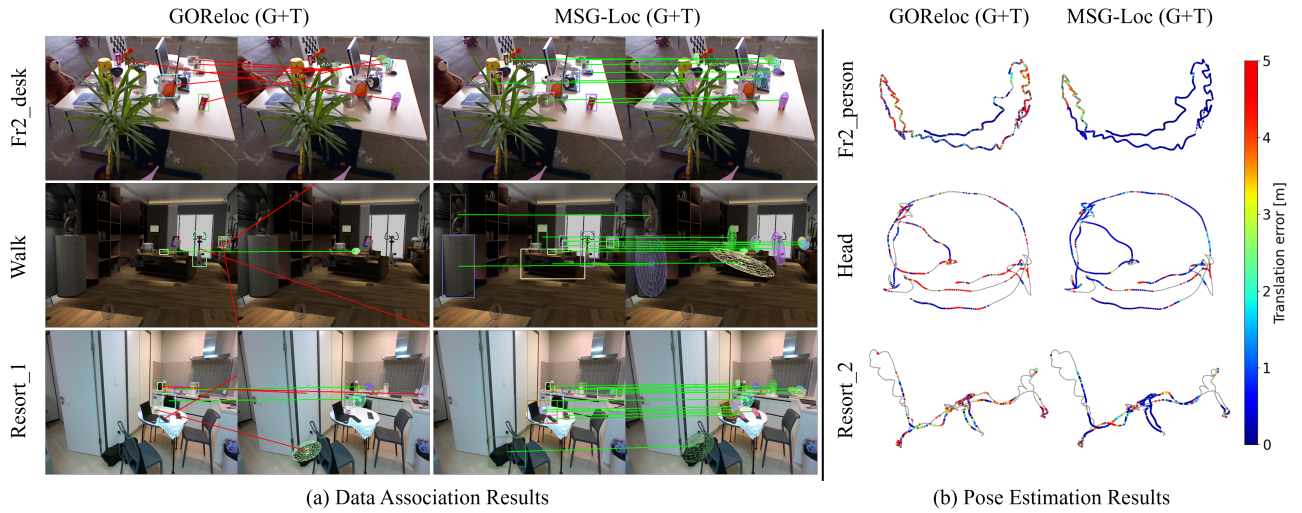


Fig. 5. **The qualitative results of data association and pose estimation.** (a) For each pair, the query frame image is presented on the left, and the corresponding landmark rendering is on the right. Green lines represent correct associations, whereas red lines indicate incorrect ones. (b) The color bar on the right denotes the magnitude of translation error. Scatter points represent successful pose estimates and are colored according to their translation error as indicated by the color bar, whereas continuous gray lines denote failures.

TABLE III
COMPARISON OF POSE ESTIMATION WITH FEATURE-BASED METHOD

Sequence	Method	SR _{all} [%]↑			
		@0.5 m	@1 m	@2 m	@5 m
Fr2_person	ORB-SLAM2	8.68	8.68	8.68	8.68
	SH w/ O	36.59	41.87	48.66	75.02
	GOREloc w/ O	41.78	48.46	57.14	80.30
	MSG-Loc w/ O	59.48	61.54	66.56	87.44
Head	ORB-SLAM2	47.82	48.97	48.97	48.97
	SH w/ O	49.05	50.65	53.40	58.59
	GOREloc w/ O	49.43	51.72	55.23	61.12
	MSG-Loc w/ O	57.45	63.33	67.91	72.04
VR	ORB-SLAM2	26.98	27.12	27.12	27.12
	SH w/ O	27.46	28.56	30.87	39.56
	GOREloc w/ O	27.75	29.43	32.26	41.48
	MSG-Loc w/ O	30.10	34.04	37.30	44.46
DJI	ORB-SLAM2	5.90	5.90	5.90	5.90
	SH w/ O	6.80	9.65	16.79	36.98
	GOREloc w/ O	8.20	11.24	16.44	32.08
	MSG-Loc w/ O	12.54	17.19	22.89	33.73
Ground	ORB-SLAM2	2.56	2.75	2.75	2.75
	SH w/ O	3.69	4.31	5.06	10.87
	GOREloc w/ O	2.94	3.62	4.88	10.56
	MSG-Loc w/ O	2.62	3.50	5.75	10.87
Resort_2	ORB-SLAM2	8.48	15.18	16.22	16.22
	SH w/ O	9.09	16.34	18.41	23.66
	GOREloc w/ O	8.96	16.04	18.23	23.78
	MSG-Loc w/ O	12.99	20.30	22.93	29.15

*w/ O' denotes feature augmentation with ORB-SLAM2.

3) *Scene Invariance*: Table III presents a comparison of ORB-SLAM2 with the feature-augmented variants of SH, GOREloc, and MSG-Loc under the open-set configuration with COCO categories. The integrated approaches consistently outperform the purely feature-based ORB-SLAM2 across all evaluated sequences. On the more ambiguous ICL-LM sequences, the baselines tend to remain comparable to ORB-SLAM2 at the strict 0.5 m threshold, in contrast to MSG-Loc. Notably, our method achieves successful pose estimation within 0.5 m TE for more than 50% of frames on the 'Fr2_person' and 'Head' sequences. These results indicate that multi-label reasoning improves inlier discriminability across varying scene conditions and viewpoints, thereby increasing the overall pose-estimation success rate over all frames.

TABLE IV
ABLATION STUDY ON TOP- τ UNDER OPEN-SET CONFIGURATION

Top- τ	Method	Fr2_person: SR _{succ} [%]↑			Head: SR _{succ} [%]↑		
		@0.5 m	@1 m	@2 m	@0.5 m	@1 m	@2 m
$\tau = 1$	SH	34.07	41.50	50.98	6.46	11.73	18.03
	GOREloc	31.71	38.78	51.02	6.43	11.04	19.77
	M-LLE	49.87	56.33	66.66	32.04	42.96	54.79
	M-LLE + CALP	53.45	58.00	64.65	35.97	49.11	60.80
$\tau = 3$	SH	40.99	48.04	56.85	7.44	12.36	19.79
	GOREloc	47.12	55.58	66.12	8.32	14.65	23.42
	M-LLE	62.51	68.72	74.67	30.82	43.50	53.77
	M-LLE + CALP	67.05	69.38	75.07	36.07	49.36	60.02

TABLE V
ABLATION STUDY ON TOP- K UNDER OPEN-SET CONFIGURATION

Method	Top- K	Fr2_person: SR _{succ} [%]↑			Head: SR _{succ} [%]↑		
		@0.5 m	@1 m	@2 m	@0.5 m	@1 m	@2 m
M-LLE	$K = 1$	30.30	36.66	47.41	13.33	21.05	25.88
	$K = 3$	72.23	77.30	81.73	36.90	52.86	64.40
	$K = 5$	73.95	77.59	81.43	38.45	59.88	74.32
M-LLE + CALP	$K = 1$	54.03	59.61	65.38	32.78	44.65	51.56
	$K = 3$	75.60	78.85	82.40	43.91	60.22	70.90
	$K = 5$	76.07	79.00	82.54	44.64	63.63	76.57

E. Ablation Study

1) *Analysis of the Correspondence Candidate Set*: Table IV presents an ablation study on Top- τ under the COCO open-set configuration, comparing the baselines with the proposed modules. As τ increases from 1 to 3, SR_{succ} generally improves across most methods. While M-LLE shows slight regressions in a few cases, its combination with CALP yields consistent improvements across both sequences and most thresholds. These results indicate that expanding the candidate set necessitates topology-based consistency constraints to effectively suppress noisy candidates and ensure valid inference.

2) *Analysis of Semantic Context*: Table V reports a Top- K ablation of the proposed modules under the LVIS open-set configuration. Compared to $K = 1$, M-LLE achieves a substantial performance gain at $K = 3$, improving SR_{succ} by an average of 35.1 percentage points across both se-

TABLE VI
RUNTIME COMPARISON WITH BASELINES

	TUM RGB-D			ICL-LM		
	MSG-Loc	GOReloc	SH	MSG-Loc	GOReloc	SH
Preprocessing	4.17 ms	4.23 ms	19.58 ms	4.38 ms	4.49 ms	27.14 ms
Graph Matching	8.28 ms	0.50 ms	168.5 ms	4.85 ms	1.10 ms	454.5 ms
Pose Estimation	14.95 ms	4.11 ms	3.73 ms	7.30 ms	3.24 ms	3.27 ms
Total	27.40 ms	8.84 ms	191.8 ms	16.53 ms	8.83 ms	484.9 ms

quences and thresholds. This suggests that retaining more semantic candidates enables better exploitation of semantic context. Under the same setting, M-LLE + CALP demonstrates consistent improvements as K increases from 1 to 5 by leveraging per-object multi-label predictions and graph-based context to enhance candidate reliability and suppress noise. Consequently, integrating M-LLE and CALP improves the stability and generalization of global localization.

F. Runtime Analysis

All experiments were conducted on a desktop with an Intel i9-13900K CPU. The runtime comparison is summarized in Table VI. Despite maintaining high semantic discriminability, MSG-Loc achieves real-time performance, operating above 36.4 Hz across datasets. GOReloc attains the fastest runtime of about 8.8 ms through a lightweight descriptor-matching design. In contrast, SH shows notable computational overhead, running below 5.3 Hz due to the heavy cost of histogram-based matching. This comparison indicates that the proposed method achieves an effective balance between runtime efficiency and semantic correspondence accuracy for online operation.

V. CONCLUSION

In this work, we introduce *MSG-Loc*, a novel semantic graph matching framework that leverages multi-label likelihood for object-level global localization. The proposed method effectively bridges closed-set and open-set paradigms for object detection and classification. Specifically, M-LLE constructs likelihoods based on the multi-label hypotheses assigned to each observed object and each map landmark, resolving semantic ambiguity without relying on deterministic single-label predictions. Furthermore, CALP aggregates local maximum likelihoods via the graph structure, enabling consistent data association even with sparse or ambiguous observations. Experimental results demonstrate the scalability of MSG-Loc, maintaining robust performance even as the number of recognizable categories significantly increases.

In future work, we aim to extend this framework toward an open-world paradigm by incorporating vision-language embeddings and captions for interpretable object localization.

REFERENCES

- C. Cadena, L. Carlone, H. Carrillo, Y. Latif, D. Scaramuzza, J. Neira, I. Reid, and J. J. Leonard, "Past, present, and future of simultaneous localization and mapping: Toward the robust-perception age," *IEEE Transactions on robotics*, vol. 32, no. 6, pp. 1309–1332, 2016.
- R. Mur-Artal and J. D. Tardós, "Orb-slam2: An open-source slam system for monocular, stereo, and rgb-d cameras," *IEEE Transactions on Robotics*, vol. 33, no. 5, pp. 1255–1262, 2017.
- M. Zins, G. Simon, and M.-O. Berger, "Oa-slam: Leveraging objects for camera relocation in visual slam," in *2022 IEEE international symposium on mixed and augmented reality (ISMAR)*. IEEE, 2022, pp. 720–728.
- J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, "LoFTR: Detector-free local feature matching with transformers," *CVPR*, 2021.
- M. Michalkiewicz, S. Bai, M. Baktashmotlagh, V. Jampani, and G. Balakrishnan, "Not all views are created equal: Analyzing viewpoint instabilities in vision foundation models," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2025, pp. 9113–9123.
- B. Zhou, Y. Meng, and F. Kai, "Object-based loop closure with directional histogram descriptor," in *2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2022, pp. 1346–1351.
- Z. Cao, Q. Zhang, J. Guang, S. Wu, Z. Hu, and J. Liu, "Semantictopoloop: Semantic loop closure with 3d topological graph based on quadric-level object map," *IEEE Robotics and Automation Letters*, 2024.
- Y. Wu, Y. Zhang, D. Zhu, Z. Deng, W. Sun, X. Chen, and J. Zhang, "An object slam framework for association, mapping, and high-level tasks," *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 2912–2932, 2023.
- A. Gawel, C. D. Don, R. Siegwart, J. Nieto, and C. Cadena, "X-view: Graph-based semantic multi-view localization," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1687–1694, 2018.
- X. Guo, J. Hu, J. Chen, F. Deng, and T. L. Lam, "Semantic histogram based graph matching for real-time multi-robot global localization in large scale environment," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 8349–8356, 2021.
- Y. Wang, C. Jiang, and X. Chen, "Goreloc: Graph-based object-level relocation for visual slam," *IEEE Robotics and Automation Letters*, vol. 9, no. 10, pp. 8234–8241, 2024.
- S. Matsuzaki, T. Sugino, K. Tanaka, Z. Sha, S. Nakaoka, S. Yoshizawa, and K. Shintani, "Clip-loc: Multi-modal landmark association for global localization in object-based maps," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 13 673–13 679.
- S. Matsuzaki, K. Tanaka, and K. Shintani, "Clip-clique: Graph-based correspondence matching augmented by vision language models for object-based global localization," *IEEE Robotics and Automation Letters*, 2024.
- R. F. Salas-Moreno, R. A. Newcombe, H. Strasdat, P. H. Kelly, and A. J. Davison, "Slam++: Simultaneous localisation and mapping at the level of objects," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2013, pp. 1352–1359.
- S. Yang and S. Scherer, "Cubeslam: Monocular 3-d object slam," *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 925–938, 2019.
- S. Yang and S. Scherer, "Monocular object and plane slam in structured environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3145–3152, 2019.
- L. Nicholson, M. Milford, and N. Sünderhauf, "Quadricslam: Dual quadrics from object detections as landmarks in object-oriented slam," *IEEE Robotics and Automation Letters*, vol. 4, no. 1, pp. 1–8, 2019.
- Y. Wu, Y. Zhang, D. Zhu, Y. Feng, S. Coleman, and D. Kerr, "Eao-slam: Monocular semi-dense object slam based on ensemble data association," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 4966–4973.
- Y. Wang, B. Xu, W. Fan, and C. Xiang, "Qiso-slam: Object-oriented slam using dual quadrics as landmarks based on instance segmentation," *IEEE Robotics and Automation Letters*, vol. 8, no. 4, pp. 2253–2260, 2023.
- Y. Wang, C. Jiang, and X. Chen, "Voom: Robust visual object odometry and mapping using hierarchical landmarks," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 10 298–10 304.
- E. Stumm, C. Mei, S. Lacroix, J. Nieto, M. Hutter, and R. Siegwart, "Robust visual place recognition with graph kernels," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2016, pp. 4535–4544.
- C. Xu, J. Wang, W. Yang, H. Yu, L. Yu, and G.-S. Xia, "Detecting tiny objects in aerial images: A normalized wasserstein distance and a new benchmark," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 190, pp. 79–93, 2022.
- J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 573–580.
- S. Saeedi, E. D. C. Carvalho, W. Li, D. Tzoumanikas, S. Leutenegger, P. H. J. Kelly, and A. J. Davison, "Characterizing visual localization and mapping datasets," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 6699–6705.
- G. Jocher, A. Chaurasia, and J. Qiu, "Ultralytics yolov8," 2023. [Online]. Available: <https://github.com/ultralytics/ultralytics>
- T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*. Springer, 2014, pp. 740–755.
- A. Gupta, P. Dollar, and R. Girshick, "Lvis: A dataset for large vocabulary instance segmentation," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 5356–5364.
- S. Liu *et al.*, "Grounding dino: Marrying dino with grounded pre-training for open-set object detection," in *European conference on computer vision*. Springer, 2024, pp. 38–55.
- H. Yuan, X. Li, C. Zhou, Y. Li, K. Chen, and C. C. Loy, "Open-vocabulary sam: Segment and recognize twenty-thousand classes interactively," in *European Conference on Computer Vision*. Springer, 2024, pp. 419–437.
- T. Pan, L. Tang, X. Wang, and S. Shan, "Tokenize anything via prompting," in *European Conference on Computer Vision*. Springer, 2024, pp. 330–348.