

AirObject: A Temporally Evolving Graph Embedding for Object Identification

Nikhil Varma Keetha^{1,2} Chen Wang¹ Yuheng Qiu¹ Kuan Xu³ Sebastian Scherer¹
¹Carnegie Mellon University ²IIT (ISM) Dhanbad, India ³Geek+ Corp
 {keethanikhil, xukuanhit}@gmail.com, chenwang@dr.com, {yuhengq, basti}@andrew.cmu.edu

Abstract

Object encoding and identification are vital for robotic tasks such as autonomous exploration, semantic scene understanding, and re-localization. Previous approaches have attempted to either track objects or generate descriptors for object identification. However, such systems are limited to a “fixed” partial object representation from a single viewpoint. In a robot exploration setup, there is a requirement for a temporally “evolving” global object representation built as the robot observes the object from multiple viewpoints. Furthermore, given the vast distribution of unknown novel objects in the real world, the object identification process must be class-agnostic. In this context, we propose a novel temporal 3D object encoding approach, dubbed AirObject, to obtain global keypoint graph-based embeddings of objects. Specifically, the global 3D object embeddings are generated using a temporal convolutional network across structural information of multiple frames obtained from a graph attention-based encoding method. We demonstrate that AirObject achieves the state-of-the-art performance for video object identification and is robust to severe occlusion, perceptual aliasing, viewpoint shift, deformation, and scale transform, outperforming the state-of-the-art single-frame and sequential descriptors. To the best of our knowledge, AirObject is one of the first temporal object encoding methods. Source code is available at <https://github.com/Nik-V9/AirObject>.

1. Introduction

Object encoding and identification are crucial for robotic tasks such as autonomous exploration, semantic scene understanding, and loop closure in simultaneous localization and mapping (SLAM). For example, object-based semantic SLAM and identification of revisited objects require robust and efficient object encodings [41, 45, 46]. Prior approaches proposed in the literature have attempted to track object detections [47], use keypoint features [10], and generate graph-based embeddings for object matching [50]. However, such systems are limited to a “fixed” object rep-

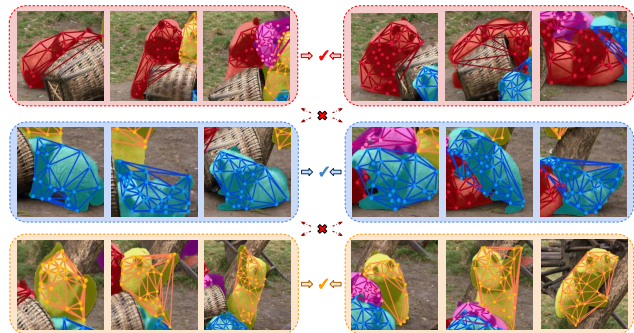


Figure 1. Temporally evolving topological graph representations of objects within a video sequence. We propose a method, AirObject, to match these temporally evolving representations and alleviate problems caused by perceptually-aliased occluded single frame representations.

resentation from a single viewpoint and are not robust to severe occlusion, viewpoint shift, perceptual aliasing, or scale transform. These single frame representations tend to lead to false correspondences amongst perceptually-aliased objects, especially when severely occluded. Hence, a robust object encoding method that aggregates the temporally “evolving” object structures is necessary since we often observe more information when the camera or object moves, as shown in Figure 1.

In this work, we propose a novel temporal encoding approach, dubbed AirObject, that encapsulates the evolving topological graph-based representations of objects. It is very simple and only contains three modules. Specifically, we use extracted deep learned keypoint features [9] across multiple frames to form sequences of object-wise topological graph neural networks (GNNs), which on embedding generate temporal 3D object descriptors. We next employ a graph attention-based sparse encoding method on these topological GNNs to generate content graph features and location graph features representing the structural information of the object. Then, these graph features are aggregated across multiple frames using a single-layer temporal convolutional network to generate a temporal 3D object descrip-

tor. These generated object descriptors, which accumulate knowledge across multiple evolving representations of the objects, are robust to severe occlusion, viewpoint changes, deformation, perceptual aliasing, and the scale transform.

In summary, we make the following contributions:

- To learn the geometric relationship of the keypoints, we construct topological object graphs for each frame using Delaunay triangulation.
- We introduce a simple yet effective temporal object encoding method to aggregate and embed multiple instances into an object descriptor.
- Extensive experiments show that AirObject consistently provides the state-of-the-art performance for video object identification on four large-scale datasets.

2. Related Work

In this section, we review single-frame and sequential methods based on handcrafted and deep learned features. Furthermore, we also review Visual Place Recognition (VPR) methods that can be extended to object identification. Object tracking methods based on networks such as Mask R-CNN [19] are not included because they are unsuitable for object re-identification.

2.1. Single-frame Representations

Handcrafted features such as SIFT [32] and SURF [4] have been widely used in classical approaches for loop closure, object matching, and VPR. One such classical approach, fast appearance-based mapping (FABMAP) [17], utilizes the trained visual vocabulary of SURF features, obtained by hierarchical k-means clustering, for identifying revisited objects through feature distribution matching. Further extending this idea, a binary descriptor ORB [36] was used in DBoW2 [12] to achieve better speed. Further building on the notion of vocabulary-based retrieval, several approaches [13, 15, 30, 39] have been proposed. However, these handcrafted features based methods are sensitive to environmental changes and lead to false matches when the local descriptors are not discriminative enough.

The recent success of convolutional neural networks (CNN) [28] in computer vision has led to the rise of deep learned features-based image retrieval. The methods employing deep learned features have shown tremendous improvements over handcrafted features. One such method, [5], uses a multi-scale feature encoding across two CNN architectures to generate CNN features that are viewpoint invariant, thereby providing considerable performance improvement. Another popular end-to-end deep learning-based approach, NetVLAD [1] generates descriptors inspired by the traditional vector of locally aggregated de-

scriptors (VLAD). Further exploring other input modalities such as RGB-D images and point cloud data, several approaches [40, 53, 54] have attempted to incorporate spatial/depth data into the RGB domain for object recognition.

Recently proposed deep learning method, SuperPoint [9] leverages a self-supervised framework to train interest point detectors and descriptor extractors. Further building on SuperPoint, SuperGlue [38] introduced a local feature matcher based on graph attention [43] where the interest points are nodes of a graph, and their associated descriptors are the node features. Both SuperPoint and SuperGlue have been widely adopted for the task of feature matching and hierarchical VPR [25, 37]. Similar to SuperGlue, Xu *et al.* embeds object-wise fully connected graph-based representations of SuperPoint features using a Sparse Object Encoder to generate object descriptors [50]. However, this approach doesn't consider the explicit geometry available from the SuperPoint interest points. It is limited to only a single viewpoint, making it susceptible to false matches under perceptual aliasing and occlusion. In this context, our proposed framework generates temporal object descriptors that aggregate structural knowledge across multiple object instances.

2.2. Sequential Representations

While single frame representations have been extensively used across the literature, temporal information for compact representations has received limited attention in robotics, especially for loop closure. However, there exists an extensive array of Spatio-temporal representation techniques in related fields of research [7, 16, 21, 49, 51]. Many of these approaches use LSTM [20], GRUs [6], Graph Convolutional Networks (GCNs) [27], and Temporal Convolutional Networks [3] to model the temporal relations. In the context of VPR, there have been a few approaches attempting to leverage Spatio-temporal information in terms of landmarks [22] and bio-inspired memory cells [34].

Recently, Facil *et al.* presented concatenation, fusion, and recurrence for learning sequential representations [11]. Before this work, the concatenation of descriptors within a sequence was explored [2], where binarization was used for efficient place recognition. Similarly, Neubert *et al.* employs recurrence for learning a bio-inspired topological map of the environment [33]. More recently, a temporal convolutional network, SeqNet, was proposed to learn sequential descriptors from single image descriptors for hierarchical VPR [14]. However, the learned temporal information for SeqNet depends on the underlying visual attributes from single image descriptors which don't provide explicit knowledge regarding structural relations. In this context, our approach learns temporal information across evolving topological graph representations, which provide spatial/structural information regarding the object.

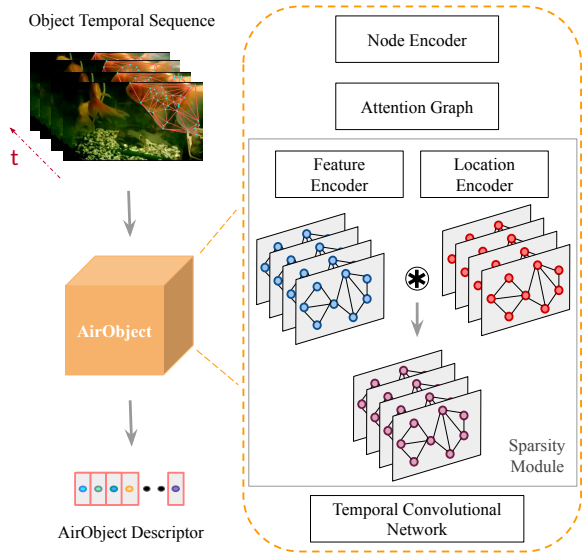


Figure 2. Schematic of our proposed approach. The SuperPoint based temporal topological graph representations of an object are input to the graph attention encoder to generate content graph features and location graph features. Then, these features are node-wise multiplied and temporally aggregated to generate the AirObject descriptors. Finally, we discuss the loss functions to supervise the encoders.

3. Proposed Approach

We propose a new architecture, AirObject, to encode objects from a video, which is shown in Figure 2. In this section, we will first present the topological graph representations of objects, and then describe the structure of the graph attention encoder and a temporal encoding approach to generate the AirObject descriptors. Finally, we discuss the loss functions to supervise the encoders.

3.1. Topological Graph Representations

Intuitively, a group of feature points on an object form a graphical representation where the feature points are nodes and their associated descriptors are the node features. Essentially, the graph’s nodes are the distinctive local features of the object, while the edges/structure of the graph represents the global structure of the object. We believe that embedding such a topological graph-based representation of an object containing both distinctive local features and global object structure will enable robust object identification similar to humans [42]. Hence based on this hypothesis, we formulate a procedure to generate topological graph representations of feature points corresponding to objects.

Given an object, we extract a set of feature points corresponding to the object, where the position of each feature point is denoted as $p_i = (x, y), i \in [1, N]$ and the associated descriptor as $d_i \in \mathbb{R}^{D_p}$, where D_p is the descriptor




Figure 3. Topological Graph Representations of Objects. These representations are generated by using Delaunay Triangulation on object-wise grouped SuperPoint keypoints.

dimension. In practice, these object-wise grouped feature points could be obtained using point detector SuperPoint [9] along with ground-truth instance segmentations or masks from an off-the-shelf network like Mask R-CNN [19] or an open-world object detector [23]. Given these object-wise grouped feature points, our goal is to generate a topological graph representation that leverages the explicit geometry provided by the feature point positions.

We build a topological graph structure for the feature points corresponding to an object by using Delaunay triangulation [29] on the positions of the feature points as shown in Figure 3. Delaunay triangulation is a mathematical formulation, where given a set of discrete points, the objective is to provide a triangulation avoiding narrow and intersecting triangles such that no discrete point is present inside the circumcircle of any triangle of the triangulation. This particular property coupled with fast compute time makes it suited to generate a triangular mesh representation that intuitively captures the feature points’ local & global structure. We believe that this mesh representation will enable the graph attention encoder to reason better about the object’s structure, thereby making the final temporal object descriptor robust to deformation or occlusion.

3.2. Graph Attention Encoder

The topological graph representations of the object feature points are inputted to the Graph Attention-based sparse object encoder [50]. A particular desirable property of a sparse object encoder is that a single keypoint should only have a local effect on sparse locations of the object descriptor so that addition or removal of keypoints doesn’t significantly change the object descriptor. Furthermore, the object encoder should encode distinct keypoints to unique locations within the object descriptor. To facilitate this, the encoder comprises a Node Encoder, a two-layer Graph Attention (GAT) [43] module followed by a Sparsity module that contains two parallel heads, namely a feature and a location encoder whose outputs are element-wise multiplied.

The Node Encoder takes the positions and descriptors of

the topological graph-based object representation and encodes the graph’s node features x_i as the concatenation of the point descriptor and a transform of the point position:

$$x_i^{(1)} = [d_i \parallel \text{MLP}(p_i)], \quad x_i^{(1)} \in \mathbb{R}^{D_n}, \quad (1)$$

where $D_n = D_p + D_m$, \parallel denotes the concatenation operator, the Multi-layer Perceptron (MLP) module maps $\mathbb{R}^2 \mapsto \mathbb{R}^{D_m}$, and x_i^l denotes the output of the l -th layer of the GNN. In practice, we normalize the position p_i to $[-1, 1]$ by the object size, which considers the object center as the origin. The concatenation of the transformation of the position as opposed to a summation similar to SuperGlue [38] helps the encoder explicitly learn the object structure since the relative position information is not blended into the descriptor. Furthermore, this concatenation operation enables the Sparsity module of the encoder to learn sparse non-zero locations for the keypoints based on the object structure.

The graph’s node features x_i and the adjacency matrix from the topological graph structure are input to a two-layered GAT to enable structured attention-based message propagation between the features of the object¹. This helps the encoder reason about global feature interactions between the distinctive local keypoint-based features of the object. Afterward, the output of the GAT is passed on to the Sparsity module to encode the graph embeddings such that learned location features of the node decide the sparse location of the keypoint on the temporal object descriptor.

The Sparsity module of the encoder contains two parallel heads where each head has two stacked sparsity layers to learn the location feature x_i^L and the content feature x_i^C , whose inputs are x_i from the GAT (we leave out the layer index (l) for simplicity). The sparsity layers for the location node features and content node features are defined as:

$${}^{(l+1)}x_i^L = \text{ReLU}(W_L^{(l)} \cdot {}^{(l)}x_i^L), \quad (2a)$$

$${}^{(l+1)}x_i^C = \text{ReLU}(W_C^{(l)} \cdot {}^{(l)}x_i^C), \quad (2b)$$

where $W_L^{(l)}, W_C^{(l)} \in \mathbb{R}^{D_o \times D_n^l}$, $D_n^l < D_o$ are the learnable location and content weights, respectively, and D_o is the dimension of the temporal object descriptor. Then, these location and content features are node-wise multiplied to generate the structural graph features x_i^S .

3.3. Temporal Encoding

Given a sequence of topological graph object representations input to the graph attention encoder, we first obtain a sequence of structural graph features. Then, these structural graph features are node-wise stacked to form a tensor of size $(N_s \times D_o)$ where N_s is the number of structural graph feature nodes across a sequence, and D_o is the temporal object

¹Refer to [43] for more details on GNNs and GAT.

descriptor dimension. To perform temporal aggregation and encode these features into a single object descriptor, we use a Temporal Convolutional Network (TCN) comprising of a single layer 1-D convolution (with bias, without padding) followed by a Sequence Average Pooling (SAP) layer and an L2-Normalization layer.

In the convolution layer, we use a 1-D filter of length 1 that operates with stride 1 across the N_s dimension of the input tensor. We use a filter length of 1 to ensure that the encoding process is not limited to a fixed sequence length and is compatible with the varying total number of structural graph nodes. The dimension of the structural graph features, D_o , forms the input channels (feature maps) for the convolution layer while the output channels are also set to D_o . Thus, the convolutional kernel is of size $D_o \times 1 \times D_o$. The output of the 1-D convolution layer is further input to an SAP layer across the N_s dimension to obtain a descriptor of size $1 \times D_o$. To ensure compatibility with cosine similarity, we then use an L2-Normalization across the D_o dimension to obtain the final temporal AirObject descriptor A_k .

3.4. Loss Functions

The Graph Attention encoder is supervised by a sparse location loss, dense feature loss [50]. The objective of the sparse location loss is to ensure that location feature x_i^L is a sparse vector. The sparse location loss L_s is defined as the l_1 -norm of x_i^L .

$$L_s = \sum_{i=1}^N \|\phi(x_i^L)\|_1, \quad (3)$$

where $\phi(x) = x/\|x\|_2$ is a l_2 -normalization to prevent the location features from being zero. Given that the sparse location loss ensures that keypoints are encoded into sparse locations on the object descriptor, the objective of the dense feature loss is to ensure that distinctive keypoints are encoded to unique sparse locations on the object descriptor. Hence, dense feature loss L_d is defined as the negative l_1 -norm of the location features.

$$L_d = \max \left(0, \delta - \phi \left(\left\| \sum_{i=1}^N (x_i^L) \right\|_1 \right) \right), \quad (4)$$

where $\delta > 0$ is a positive constant. Intuitively, the combined optimization of both sparse location loss and dense feature loss enables the object encoder to encode graph representations such that the similar keypoints are encoded to similar locations, while distinctive keypoints cover different locations retaining the density of the object descriptor.

Finally, the Graph Attention Encoder and Temporal Encoder are supervised for object identification using a triplet style matching loss. The objective of the matching loss L_m is to maximize the cosine similarity of positive object pairs

Table 1. Quantitative Results: Performance comparison on four datasets.

Methods	YT-VIS			UVO			OVIS			TAO-VOS		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
Single Frame Descriptors:												
2D Baseline	68.93	80.93	74.45	73.18	81.85	77.27	29.45	68.32	41.15	46.80	72.49	56.87
NetVLAD	77.52	52.41	62.54	86.19	72.15	78.55	49.41	29.45	36.90	78.44	35.57	48.95
SeqNet ($s_l = 1$)	71.59	66.38	68.89	69.44	81.03	74.79	39.80	44.87	42.18	66.94	49.79	57.10
<i>Ours</i> : AirObject ($s_l = 1$)	79.47	73.49	76.36	82.99	82.51	82.75	38.58	55.40	45.49	66.46	60.96	63.59
Sequential Descriptors:												
3D Baseline	65.14	86.48	74.31	73.78	77.00	75.35	25.58	80.29	38.80	40.60	75.84	52.89
SeqNet	76.73	86.48	81.31	96.57	72.63	82.91	70.19	42.09	52.62	71.10	71.54	71.32
<i>Ours</i> : AirObject	85.09	82.36	83.70	94.31	83.79	88.74	69.86	42.47	52.82	72.81	71.21	72.00

and minimize the cosine similarity of negative object pairs.

$$L_m = \sum_{\{p,q\} \in P^+} (1 - C(A_p, A_q)) + \sum_{\{p,q\} \in P^-} \max(0, C(A_p, A_q) - \lambda), \quad (5)$$

where $\lambda = 0.2$, C is the cosine similarity, and P^+ and P^- are positive and negative object pairs, respectively.

4. Experimental Results

4.1. Datasets

For video object identification, we require video object sequences where objects are associated across multiple frames. Hence, to train and evaluate our proposed approach, we used four video instance segmentation datasets: YouTube Video Instance Segmentation (YT-VIS) [52], Unidentified Video Objects (UVO) [48], Occluded Video Instance Segmentation (OVIS) [35], and Tracking Any Object with Video Object Segmentation (TAO-VOS) [8, 44]. All these datasets contain a large object vocabulary and various challenging scenarios, including perceptually-aliased occluded objects, as described below:

1) *YT-VIS*: This large-scale dataset contains over 3,000 high-resolution youtube videos annotated with $\approx 5,000$ unique video instances across a 40-category object label set including common objects such as animals, vehicles, and persons. Furthermore, it contains a large amount of perceptually-aliased objects in various environmental backdrops, making it challenging. We split the sequences into a training split of 2,485 videos and a test split of 500 videos.

2) *UVO*: This dataset contains class-agnostic video instance segmentations of open-world objects. In particular, all objects present within the Kinetics400 [24] dataset videos are densely annotated such that there are about 13 objects per video. The open-world class-agnostic nature of the objects coupled with crowded scenes and complex background motions make this a challenging dataset to test the

robustness of our proposed object encoding method. We use the Dense-Annotation training and validation sets containing 393 videos with ground-truth instances for evaluation.

3) *OVIS*: This large-scale dataset was designed with the philosophy of perceiving occluded objects in videos. So for the same, this dataset contains long videos with severely occluded objects annotated with high-quality instance masks across 25 semantic categories. The severe occlusions, long video duration, and crowded scenes make this dataset very challenging for object identification. We use the training set containing 607 videos for evaluation.

4) *TAO-VOS*: This dataset is a subset of the Tracking Any Object (TAO) dataset [8] with masks for video object segmentation. TAO is a benchmark federated object tracking dataset comprising videos from 7 datasets captured in diverse environments. In particular, the large object vocabulary ranging from outdoor vehicles to indoor household objects makes this dataset challenging. We use both the training and validation sets containing 626 videos with ground-truth instances for evaluation.

4.2. Implementation Details

The AirObject configurations are $D_p = 256$, $D_m = 16$, and $D_o = 2048$. To test the generalizability, we train only on the YT-VIS train split and evaluate on all four datasets. Firstly, the Graph Attention Encoder coupled with a single-layer perceptron, pretrained on the COCO dataset [31], is finetuned on the YT-VIS training split for single-frame object matching using the sparse location loss, dense feature loss, and matching loss. For finetuning, we employ a batch size of 16 and a learning rate of $1e-4$ with Adam optimizer [26]. Then, we freeze the Graph Attention Encoder and train the Temporal Encoder on the YT-VIS training split using the object matching loss. During training, we use a batch size of 16 containing object sequences of length $s_l \leq 4$ and a learning rate of $1e-4$ with Adam optimizer.

4.3. Evaluation Criteria

For testing the object identification performance, we consider the first half of the video object sequence as the




Figure 4. Precision-Recall Curves for Video Object Identification on four datasets. The area under the curve is shown in the [brackets].




Figure 5. Precision-Recall Comparison with Local Feature Matching (SuperGlue) for Video Object Identification on two datasets. The area under the curve is shown in the [brackets].

query and the second half as the reference and evaluate by matching the query and reference object sequences within a video. To determine a match between an object pair, we compute the cosine similarity between the descriptors and define a matching threshold ρ . Based on the True Positives and False Positives, we calculate the Precision, Recall, and F1-Score accordingly. Furthermore, by varying the threshold ρ values, we generate Precision-Recall curves and calculate the area under the curves (AUCs).

4.4. Comparison to State-of-the-art Methods

We compare our proposed approach against several benchmark baselines: 2D Baseline, 3D Baseline, NetVLAD [1], and SeqNet [14]. For the 2D Baseline, similar to [50], we use the Graph Attention Encoder coupled with a single layer perceptron to perform single-frame object matching. While, for the 3D Baseline, we consider the simple case of averaging the 2D Baseline’s single-frame object descriptors to obtain temporal object descriptors.

For the NetVLAD baseline, we obtain object descriptors of size $D_{nv} = 8192$ using NetVLAD (32 clusters) on the object-wise grouped features obtained using Superpoint and ground-truth instances. We pre-train NetVLAD on COCO and finetune on the YT-VIS training split. For SeqNet, we use the NetVLAD object descriptors as the backbone and a temporal filter length of 1 to support varying object sequence lengths. We train SeqNet on the YT-VIS training split with an output descriptor dimension of $D_s = 4096$.

Table 1 and Figure 4 contain quantitative comparisons of AirObject and the baseline methods. It can be observed that AirObject outperforms all the methods both in F-1 and AUC consistently across all four datasets. In particular, AirObject outperforms the best performing single-frame and sequential descriptor methods, *i.e.*, 2D Baseline, NetVLAD, and SeqNet on average by 11.88 %, 17.58 %, 2.28 % for F-1 and by 15.05 %, 15.36 %, 0.86 % for AUC respectively. Furthermore, the performance gap between single-frame and temporal methods is consistently large for both metrics, indicating the importance of temporal information. This difference is particularly pronounced for the OVIS and TAO-VOS datasets showing that temporal information encapsulating the evolving object structure helps for mitigating perceptual aliasing and severe occlusions. To further test the applicability of the proposed approach to single-frame object matching, we evaluate SeqNet and AirObject using a sequence length (s_l) of 1. From Table 1, we can observe that AirObject also provides the best performance for single-frame object identification.

As another baseline, we compare the performance of our method against SuperGlue [38]. While initially proposed for pose estimation and homography, SuperGlue has achieved state-of-the-art performance for feature matching [18, 25, 37]. So, we extend SuperGlue to object identification using the inlier to outlier ratio as the matching score. From Figure 5, we can observe that SuperGlue provides high recall and low precision matching. We believe that high perceptual aliasing coupled with no background context leads to low precision for SuperGlue. Furthermore, given the large array of objects present within a single video, the compute time required for local feature matching makes SuperGlue infeasible for object identification. Thus further solidifying the necessity for our proposed robust object encoding and identification method, AirObject.

In Figure 6, we show a set of example object video sequences, illustrating the matches retrieved with our method compared to SeqNet, along with the topological graph-based object representations. We believe that SeqNet’s dependence on the underlying single-frame descriptor for visual attributes and structural information leads to false matches due to the lack of explicit spatial/structural infor-



Figure 6. **Qualitative Results.** In these examples, the proposed AirObject successfully retrieves the matching temporal object sequence, while SeqNet produces incorrect matches. The retrieved objects with our approach are particularly challenging, with a combination of severe occlusion, deformation, perceptual aliasing, and clutter. Please note that SeqNet does not use the graph representations illustrated, and they are only provided for visualization of the object and its structure.

Table 2. Ablation Study: Topological Graph Representations

Methods	PR-AUC (%)			
	YT-VIS	UVO	OVIS	TAO-VOS
2D Baseline (w/o Topological)	77.87	84.96	40.82	66.03
2D Baseline (w Topological)	81.00	87.25	40.11	61.60
AirObject (w/o Topological)	86.48	93.82	63.03	79.49
AirObject (w Topological)	91.20	95.80	63.07	80.09

mation. In contrast, AirObject leverages both explicit spatial information from the topological graph representations and temporal information from the object’s evolving topology, thereby leading to more robust object identification.

4.5. Ablation Studies

4.5.1 Topological Graph Representations

To analyze the effectiveness of the topological graph representations within AirObject, we train the 2D Baseline and AirObject with a fully connected feature points graph as input. From Table 2, we can observe that AirObject and 2D Baseline with topological graph representations outperform those trained with a fully connected graph. This indicates that the topological graph representations help the Graph Attention Encoder reason better about the object geometry, leading to better performance across the four datasets.

Table 3. Ablation Study: Sequence Length

Methods	PR-AUC (%)			
	YT-VIS	UVO	OVIS	TAO-VOS
SeqNet ($s_l = 1$)	77.65	84.10	40.47	63.57
SeqNet ($s_l \leq 2$)	81.21	87.95	45.50	70.48
SeqNet ($s_l \leq 4$)	84.22	90.58	49.69	75.26
SeqNet ($s_l \leq 8$)	86.46	92.53	53.42	78.35
SeqNet ($s_l \leq 16$)	88.70	93.97	56.74	79.15
SeqNet	89.42	95.68	62.60	79.02
AirObject ($s_l = 1$)	85.41	90.69	44.18	69.91
AirObject ($s_l \leq 2$)	86.56	92.13	46.40	74.16
AirObject ($s_l \leq 4$)	88.19	93.10	48.61	77.36
AirObject ($s_l \leq 8$)	89.75	94.03	51.21	79.98
AirObject ($s_l \leq 16$)	91.39	94.68	53.85	81.43
AirObject	91.20	95.80	63.07	80.09

4.5.2 Sequence Length

To verify the effectiveness of temporal information for object identification, we analyze the performance of SeqNet and AirObject for varying object sequence lengths (s_l). In line with our intuition, from Table 3, we can observe that the performance improves as the amount of temporal encoded information increases. Furthermore, the increase in performance tends to decrease slightly as the sequence length is doubled. Further in line with this trend, for OVIS, the long average video duration (more temporal information) leads to a higher performance gap between AirObject using $s_l \leq 16$ and AirObject using half the object sequence

Table 4. Ablation Study: Unique Multi-Frame Graph Features

Methods	PR-AUC (%)			
	YT-VIS	UVO	OVIS	TAO-VOS
SeqNet	89.42	95.68	62.60	79.02
AirObject (Unique Features)	90.38	94.75	61.88	79.16
AirObject (All Features)	91.20	95.80	63.07	80.09

Table 5. Ablation Study: Temporal Encoder Architecture

Methods	PR-AUC (%)			
	YT-VIS	UVO	OVIS	TAO-VOS
AirObject (Single-layer Perceptron)	84.90	92.62	51.81	68.64
AirObject (GAT)	85.20	88.88	49.55	70.09
AirObject (Concat. SLP & SeqNet)	86.00	90.65	50.55	68.80
AirObject (Attention-based TCN)	88.77	93.38	59.22	76.61
AirObject (TCN)	91.20	95.80	63.07	80.09

length. Thus, this portrays the efficacy of temporal encoding for object identification.

4.5.3 Unique Multi-Frame Graph Features

To analyze the effect of recurring object visual features across multiple frames on the temporal descriptor, we perform an experiment encoding only the unique object visual features across the video. In particular, we use the location graph features x^C to identify and encode the unique structural graph features x^S across multiple frames. In Table 4, we present results for AirObject using unique multi-frame features and AirObject using all multi-frame features. It can be observed that using unique features across multiple frames leads to a performance drop. We deduce that the recurrence of visual features on an object weighs the temporal object descriptor distribution, making it more unique to that particular object, thereby leading to better performance for AirObject using all multi-frame graph features.

4.5.4 Temporal Encoder Architecture

Given the simple nature of our temporal encoder, we further test various model ablations to verify the effectiveness of our design. In particular, we test our TCN, an Attention-based TCN, a Single-layer Perceptron (SLP), and a Graph Attention network (GAT). Inspired by Spatio-temporal work in related fields [51], we also test a method for normalized concatenation of descriptors from AirObject (SLP) and SeqNet (with 2D Baseline backbone). From Table 5, we can observe that the various model ablations, including fully connected attention across multiple temporal states, do not lead to better temporal encoding. Amongst all the methods, using a single-layer TCN works the best, thus proving the simple yet effective nature of our proposed temporal object encoding method, AirObject.



Figure 7. Representative frames of two temporal sequences from the OVIS dataset where AirObject fails to identify the objects.

4.6. Limitations

Although our proposed method, AirObject, achieves the state-of-the-art performance for object identification across challenging scenarios, its success is partially based on the segmentation head’s success. Furthermore, we notice that it struggles to identify objects when two distinct parts of an object are observed across the temporal sequences. For example, in Figure 7, we show representative frames of two temporal sequences. Here, intuitively one would match the bikes across the two temporal sequences using either the color of the rider’s coat or the relative spatial position. However, since AirObject only uses a self-attention limited to the object, it fails to identify the bikes. We believe this is a limitation of our approach. In this context, an interesting avenue for future work would be to explore encoding temporal relative spatial information using cross-attention across adjacent objects within the scene and the background.

5. Conclusion

Discriminatively identifying objects is a critical and challenging problem for robotics tasks involving autonomous exploration and semantic localization & mapping. In this paper, we present a novel temporal object encoding method, AirObject, to generate global object descriptors. The proposed temporal encoding method accumulates structural knowledge across evolving topological graph representations of an object. Our experiments show that the proposed method leads to state-of-the-art performance in object identification on four challenging datasets. We also showcase that our AirObject descriptors are robust to severe occlusion, viewpoint shift, deformation, and scale transform. While we present the efficacy of AirObject for object identification, we envision AirObject to play a key role in endowing robots with general class-agnostic semantic knowledge for real-world robotic applications.

Acknowledgements

This work was supported by ONR Grant N0014-19-1-2266 and ARL DCIST CRA award W911NF-17-2-0181.

References

- [1] Relja Arandjelovic, Petr Gronat, Akihiko Torii, Tomas Padla, and Josef Sivic. Netvlad: Cnn architecture for weakly supervised place recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5297–5307, 2016. 2, 6
- [2] Roberto Arroyo, Pablo F Alcantarilla, Luis M Bergasa, and Eduardo Romera. Towards life-long visual localization using an efficient matching of binary sequences from images. In *2015 IEEE international conference on robotics and automation (ICRA)*, pages 6328–6335. IEEE, 2015. 2
- [3] Shaojie Bai, J Zico Kolter, and Vladlen Koltun. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. *arXiv preprint arXiv:1803.01271*, 2018. 2
- [4] Herbert Bay, Andreas Ess, Tinne Tuytelaars, and Luc Van Gool. Speeded-up robust features (surf). *Computer vision and image understanding*, 110(3):346–359, 2008. 2
- [5] Zetao Chen, Adam Jacobson, Niko Sünderhauf, Ben Uppcroft, Lingqiao Liu, Chunhua Shen, Ian Reid, and Michael Milford. Deep learning features at scale for visual place recognition. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3223–3230. IEEE, 2017. 2
- [6] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv preprint arXiv:1412.3555*, 2014. 2
- [7] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5868–5877, 2017. 2
- [8] Achal Dave, Tarasha Khurana, Pavel Tokmakov, Cordelia Schmid, and Deva Ramanan. Tao: A large-scale benchmark for tracking any object. In *European conference on computer vision*, pages 436–454. Springer, 2020. 5
- [9] Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superpoint: Self-supervised interest point detection and description. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 224–236, 2018. 1, 2, 3
- [10] M-P Dubuisson and Anil K Jain. A modified hausdorff distance for object matching. In *Proceedings of 12th international conference on pattern recognition*, volume 1, pages 566–568. IEEE, 1994. 1
- [11] Jose M Facil, Daniel Olid, Luis Montesano, and Javier Civera. Condition-invariant multi-view place recognition. *arXiv preprint arXiv:1902.09516*, 2019. 2
- [12] Dorian Gálvez-López and Juan D Tardos. Bags of binary words for fast place recognition in image sequences. *IEEE Transactions on Robotics*, 28(5):1188–1197, 2012. 2
- [13] Emilio Garcia-Fidalgo and Alberto Ortiz. ibow-lcd: An appearance-based loop-closure detection approach using incremental bags of binary words. *IEEE Robotics and Automation Letters*, 3(4):3051–3057, 2018. 2
- [14] Sourav Garg and Michael Milford. Seqnet: Learning descriptors for sequence-based hierarchical place recognition. *IEEE Robotics and Automation Letters*, 6(3):4305–4312, 2021. 2, 6
- [15] Mathias Gehrig, Elena Stumm, Timo Hinzmann, and Roland Siegwart. Visual place recognition with probabilistic voting. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3192–3199. IEEE, 2017. 2
- [16] Rohit Girdhar, Deva Ramanan, Abhinav Gupta, Josef Sivic, and Bryan Russell. Actionvlad: Learning spatio-temporal aggregation for action classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 971–980, 2017. 2
- [17] Arren Glover, William Maddern, Michael Warren, Stephanie Reid, Michael Milford, and Gordon Wyeth. Openfabmap: An open source toolbox for appearance-based loop closure detection. In *2012 IEEE International Conference on Robotics and Automation*, pages 4730–4735. IEEE, 2012. 2
- [18] Stephen Hausler, Sourav Garg, Ming Xu, Michael Milford, and Tobias Fischer. Patch-netvlad: Multi-scale fusion of locally-global descriptors for place recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14141–14152, 2021. 6
- [19] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 2, 3
- [20] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 2
- [21] Ahmad Jalal, Yeon-Ho Kim, Yong-Joong Kim, Shaharyar Kamal, and Daijin Kim. Robust human activity recognition from depth video using spatiotemporal multi-fused features. *Pattern recognition*, 61:295–308, 2017. 2
- [22] Edward Johns and Guang-Zhong Yang. Place recognition and online learning in dynamic scenes with spatio-temporal landmarks. In *BMVC*, pages 1–12. Citeseer, 2011. 2
- [23] 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. 3
- [24] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. 5
- [25] Nikhil Varma Keetha, Michael J Milford, and Sourav Garg. A hierarchical dual model of environment-and place-specific utility for visual place recognition. *IEEE Robotics and Automation Letters*, 2021. 2, 6
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 5
- [27] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. 2

- [28] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. 2
- [29] Der-Tsai Lee and Bruce J Schachter. Two algorithms for constructing a delaunay triangulation. *International Journal of Computer & Information Sciences*, 9(3):219–242, 1980. 3
- [30] Stefan Leutenegger, Margarita Chli, and Roland Y Siegwart. Brisk: Binary robust invariant scalable keypoints. In *2011 International conference on computer vision*, pages 2548–2555. Ieee, 2011. 2
- [31] 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. 5
- [32] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 2
- [33] Peer Neubert, Stefan Schubert, and Peter Protzel. A neurologically inspired sequence processing model for mobile robot place recognition. *IEEE Robotics and Automation Letters*, 4(4):3200–3207, 2019. 2
- [34] Vu Anh Nguyen, Janusz A Starzyk, and Wooi-Boon Goh. A spatio-temporal long-term memory approach for visual place recognition in mobile robotic navigation. *Robotics and Autonomous Systems*, 61(12):1744–1758, 2013. 2
- [35] Jiyang Qi, Yan Gao, Yao Hu, Xinggang Wang, Xiaoyu Liu, Xiang Bai, Serge Belongie, Alan Yuille, Philip Torr, and Song Bai. Occluded video instance segmentation. *arXiv preprint arXiv:2102.01558*, 2021. 5
- [36] Ethan Rublee, Vincent Rabaud, Kurt Konolige, and Gary Bradski. Orb: An efficient alternative to sift or surf. In *2011 International conference on computer vision*, pages 2564–2571. IEEE, 2011. 2
- [37] Paul-Edouard Sarlin, Cesar Cadena, Roland Siegwart, and Marcin Dymczyk. From coarse to fine: Robust hierarchical localization at large scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12716–12725, 2019. 2, 6
- [38] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4938–4947, 2020. 2, 4, 6
- [39] Dominik Schlegel and Giorgio Grisetti. Hbst: A hamming distance embedding binary search tree for feature-based visual place recognition. *IEEE Robotics and Automation Letters*, 3(4):3741–3748, 2018. 2
- [40] Max Schwarz, Hannes Schulz, and Sven Behnke. Rgb-d object recognition and pose estimation based on pre-trained convolutional neural network features. In *2015 IEEE international conference on robotics and automation (ICRA)*, pages 1329–1335. IEEE, 2015. 2
- [41] Akash Sharma, Wei Dong, and Michael Kaess. Compositional and scalable object slam. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11626–11632. IEEE, 2021. 1
- [42] Michael J Tarr and William G Hayward. The concurrent encoding of viewpoint-invariant and viewpoint-dependent information in visual object recognition. *Visual Cognition*, 25(1-3):100–121, 2017. 3
- [43] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017. 2, 3, 4
- [44] Paul Voigtlaender, Lishu Luo, Chun Yuan, Yong Jiang, and Bastian Leibe. Reducing the annotation effort for video object segmentation datasets. In *WACV*, 2021. 5
- [45] Chen Wang, Yuheng Qiu, Wenshan Wang, Yafei Hu, Seungchan Kim, and Sebastian Scherer. Unsupervised online learning for robotic interestingness with visual memory. *IEEE Transactions on Robotics*, 2021. 1
- [46] Chen Wang, Wenshan Wang, Yuheng Qiu, Yafei Hu, and Sebastian Scherer. Visual memorability for robotic interestingness via unsupervised online learning. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 52–68. Springer, 2020. 1
- [47] Chieh-Chih Wang, Charles Thorpe, Sebastian Thrun, Martial Hebert, and Hugh Durrant-Whyte. Simultaneous localization, mapping and moving object tracking. *The International Journal of Robotics Research*, 26(9):889–916, 2007. 1
- [48] Weiyao Wang, Matt Feiszli, Heng Wang, and Du Tran. Unidentified video objects: A benchmark for dense, open-world segmentation. *arXiv preprint arXiv:2104.04691*, 2021. 5
- [49] Lin Wu, Yang Wang, Ling Shao, and Meng Wang. 3-d personvlad: Learning deep global representations for video-based person reidentification. *IEEE transactions on neural networks and learning systems*, 30(11):3347–3359, 2019. 2
- [50] Kuan Xu, Chen Wang, Chao Chen, Wei Wu, and Sebastian Scherer. Aircode: A robust object encoding method. *IEEE Robotics and Automation Letters*, 2022. 1, 2, 3, 4, 6
- [51] Jinrui Yang, Wei-Shi Zheng, Qize Yang, Ying-Cong Chen, and Qi Tian. Spatial-temporal graph convolutional network for video-based person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3289–3299, 2020. 2, 8
- [52] Linjie Yang, Yuchen Fan, and Ning Xu. Video instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5188–5197, 2019. 5
- [53] Hasan FM Zaki, Faisal Shafait, and Ajmal Mian. Convolutional hypercube pyramid for accurate rgb-d object category and instance recognition. In *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1685–1692. IEEE, 2016. 2
- [54] Hasan FM Zaki, Faisal Shafait, and Ajmal Mian. Viewpoint invariant semantic object and scene categorization with rgb-d sensors. *Autonomous Robots*, 43(4):1005–1022, 2019. 2