

# Leveraging Geometric Priors for Unaligned Scene Change Detection

Ziling Liu<sup>1\*</sup>, Ziwei Chen<sup>1\*</sup>, Mingqi Gao<sup>1,3</sup>, Jinyu Yang<sup>4</sup>✉, Feng Zheng<sup>1,2</sup>✉

**Abstract**—Unaligned Scene Change Detection aims to detect scene changes between image pairs captured at different times without assuming viewpoint alignment. To handle viewpoint variations, current methods rely solely on 2D visual cues to establish cross-image correspondence to assist change detection. However, large viewpoint changes can alter visual observations, causing appearance-based matching to drift or fail. Additionally, supervision limited to 2D change masks from small-scale SCD datasets restricts the learning of generalizable multi-view knowledge, making it difficult to reliably identify visual overlaps and handle occlusions. This lack of explicit geometric reasoning represents a critical yet overlooked limitation. In this work, we introduce geometric priors for the first time to address the core challenges of unaligned SCD, for reliable identification of visual overlaps, robust correspondence establishment, and explicit occlusion detection. Building on these priors, we propose a training-free framework that integrates them with the powerful representations of a visual foundation model to enable reliable change detection under viewpoint misalignment. Through extensive evaluation on the PSCD, ChangeSim, and PASLCD datasets, we demonstrate that our approach achieves superior and robust performance. Our code will be released at <https://github.com/ZilingLiu/GeoSCD>.

## I. INTRODUCTION

Scene Change Detection (SCD) has broad applicability across multiple domains, including scene updating [1], [2], [3], object-level navigation in robotic systems [4], and environmental perception in autonomous driving [5], [6], [7]. While most existing SCD methods [7], [8], [9] are developed and evaluated under the assumption of viewpoint alignment, this constraint limits their applicability to real-world scenarios. Usually, images captured at different times exhibit viewpoint discrepancies caused by sensor placement, motion, or environmental factors, as commonly encountered in autonomous driving, drone, or mobile robot applications.

Several recent works have attempted to address unaligned viewpoints in SCD by establishing correspondences between image pairs, typically relying on 2D cues such as optical flow [10], feature correlation [6], or attention-based semantic alignment [11], and then integrating these correspondences with a change decoder to predict changes. Although these methods achieve reasonable performance on existing evaluation datasets [5], [6] with limited and homogeneous viewpoint variations, they struggle when applied to more complex scenes with diverse viewpoint changes. This is mainly because these approaches typically rely on supervision from

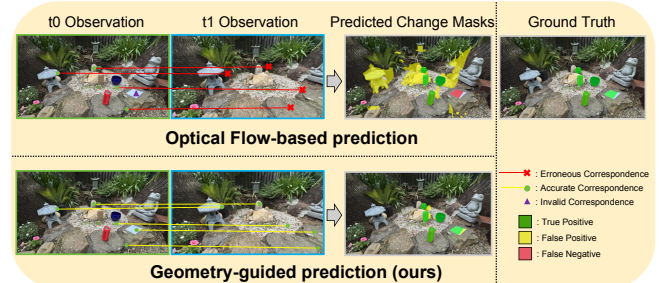


Fig. 1: Comparison between GFM-based geometric correspondences and optical flow-based correspondences. Invalid correspondence means that the matched point falls outside the image bounds. Leveraging geometric correspondences, our method generates more accurate change masks.

only 2D change masks and are trained on limited SCD data. As a result, they fail to learn generalizable multi-view knowledge, hindering their ability to maintain robust correspondence and performance across diverse viewpoint changes. Furthermore, relying solely on 2D information presents significant challenges when dealing with unaligned viewpoints. For example, these methods often struggle to robustly determine the visual overlap between images and exclude irrelevant information outside the overlapping regions. Additionally, viewpoint shifts inevitably cause occlusions, where regions visible in one image are obscured in the other. With only 2D cues, it is difficult to explicitly detect occluded areas and avoid misinterpreting them as changes. In contrast, 3D information provides a natural solution to these challenges. Just as humans can effortlessly establish associations and perceive occlusions when observing a scene from different viewpoints thanks to our innate geometric reasoning, robust unaligned SCD also requires explicit geometric understanding.

Recently, pretrained geometric foundation models (GFM) [12], [13], [14], [15] have shown strong generalization in recovering 3D geometry, such as poses, depth, and structure, directly from multi-view images across diverse scenes and viewpoint disparities. In real-world scenarios, when scene changes occur, the global geometric structure of the scene typically remains consistent. Under such conditions, GFM can robustly reconstruct per-pixel geometry and camera poses from images captured at different time points. This enables a unified 3D representation of the scene that is invariant to viewpoint changes, providing consistent holistic geometry before and after changes. By reprojecting pixels using estimated camera poses, we can

\*Equal contribution. ✉Corresponding author.

<sup>1</sup>Z. Liu, Z. Chen, M. Gao and F. Zheng are with Southern Univ. of Science and Technology (SUSTech), China. <sup>2</sup>F. Zheng is also with Spatialtemporal AI, China. <sup>3</sup>M. Gao is also with Univ. of Sheffield, UK. <sup>4</sup>J. Yang is with Harbin Institute of Technology, Shenzhen, China. (Emails: {liuzl2024, chenwz2023}@sustech.edu.cn; f.zheng@ieee.org; mingqi.gao@outlook.com; jinyu.yang96@outlook.com).

effectively identify overlapping regions and establish robust correspondences before and after changes. These projection-based geometric correspondences also naturally handle regions that have undergone changes, outperforming 2D feature-based methods that rely on appearance consistency. We demonstrate the advantage of geometric correspondence over visual correspondence for unaligned SCD in Fig. 1. Moreover, GFM reconstruction enables explicit detection of occluded regions, allowing us to effectively exclude them from genuine change areas.

Based on these capabilities, we propose a novel two-stage framework that first leverages GFM to establish robust geometric understanding and then utilizes it to guide change mask prediction. This approach integrates the power of the visual foundation model SAM [16], drawing inspiration from zero-shot methods [17], [18] to achieve strong cross-dataset generalization. We evaluated our method against several state-of-the-art works on the PSCD, ChangeSim, and PASLCD datasets. The results demonstrate that our approach consistently maintains superior and robust performance across diverse viewpoint changes and scenes. Our contributions are as follows:

- We are the first to introduce geometric priors from GFM to accurately identify visual overlaps, establish robust correspondences, and explicitly detect occlusions, which are difficult to achieve with 2D visual cues only.
- We propose a novel training-free framework that integrates geometric priors with a visual foundation model, eliminating the need for large-scale annotated datasets and enhancing generalization capability.
- Through extensive evaluation on PSCD, ChangeSim, and PASLCD datasets, we demonstrate that our method achieves leading and robust performance under various viewpoint changes and scene types. Our ablation studies further validate the effectiveness of geometric priors for unaligned SCD.

## II. RELATED WORK

### A. Scene Change Detection

With the advance of deep feature learning, numerous SCD methods based on deep neural networks have emerged. In traditional aligned settings, mainstream approaches [8], [19], [9], [7] typically adopt Siamese CNN [20] or Transformer [21] architectures with multi-scale feature fusion or temporal attention to improve feature alignment and change modeling, achieving strong performance on aligned datasets. Meanwhile, with the rise of visual foundation models, some studies [17], [18] have also explored training-free SCD frameworks, such as leveraging tracking [22] or segmentation models [16] to perform zero-shot change detection, thereby avoiding the need for large-scale annotated datasets. For the more challenging case of unaligned viewpoints, several works have made initial attempts. For example, DOF-CDNet [10] employs optical flow to align cross-temporal image pairs, but it is prone to errors under large viewpoint differences or in the presence of object motion. CSCDNet [6] introduces a correlation layer in the

decoder to enhance robustness under large viewpoint shifts. Despite this, the method remains limited by its reliance on feature correlation, which lacks explicit geometric reasoning and hinders generalization in complex 3D scenes. Furthermore, its local window-based correlation strategy proves inadequate for handling significant viewpoint changes. In contrast, RSCD [11] utilizes frozen DINOv2 [23] features and bidirectional cross-attention to establish correspondences. Nevertheless, without geometric grounding in its backbone and due to its dependence on semantic alignment, the approach struggles to generalize across diverse multi-view data and is highly susceptible to geometric variations in the scene, such as spatial layout shifts and viewpoint disparities.

### B. Geometric Foundation Model

Geometric foundation models (GFMs) are attention-based models that directly infer core 3D properties including camera parameters, depth, and per-pixel point maps from images, thereby providing explicit multi-view geometric priors [12]. Recent advances [14], [13], [15], [24] have shifted the focus from robust two-view point-map regression [12], [25] toward scalable multi-view and multi-attribute prediction. Beyond reshaping traditional reconstruction pipelines, GFMs have demonstrated strong impact across diverse downstream tasks. For example, they enable realistic scene and camera control in video generation [26], replace traditional pose optimization with direct point-map regression in SLAM [27], and enhance representation learning for 3D scene understanding [28]. These applications highlight the effectiveness of GFMs' geometric priors for processing multi-view data and boosting downstream performance. Motivated by these advances, we are the first to adopt GFMs in unaligned SCD to address the key challenges of viewpoint misalignment

## III. METHOD

### A. Overview

As shown in Fig. 2, our framework comprises two core modules: the Geometric Priors Generation module and the Geometry-guided Change Mask Prediction module. Given an image pair  $\{I_1, I_2\}$  captured at two different times with unaligned viewpoints, where  $I_1$  serves as the query image and  $I_2$  serves as the reference image, the Geometric Priors Generation Module leverages the GFM to recover depth maps and camera parameters. In this work, we adopt VGGT [13] due to its demonstrated accuracy and efficiency. Using these reconstruction results, the module establishes pixel-wise correspondences and identifies the visual overlap between the image pair. Furthermore, it explicitly detects occluded regions to allow subsequent filtering. The Geometry-guided Change Mask Prediction module then integrates these geometric cues with the Visual Foundation Model SAM [16], inspired by zero-shot SCD approaches [29]. Guided by the geometric correspondences, initial change proposals are generated within overlap regions by measuring the similarity between the SAM encoder intermediate features of matched pixels. These proposals are then refined using the occlusion masks to produce coarse change masks. Subsequently, the coarse masks are

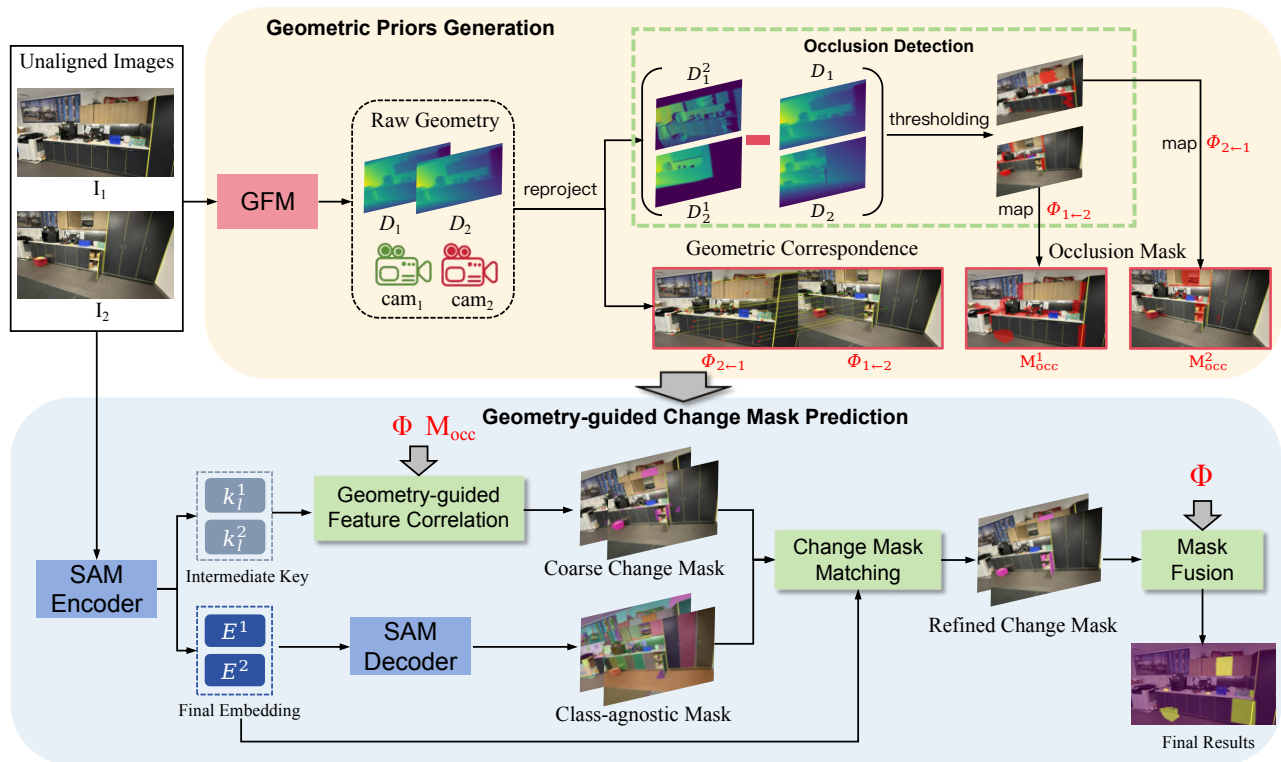


Fig. 2: Framework overview. We first derive geometric correspondences, visual overlaps, and occlusion masks from GFM reconstructions. These cues then guide SAM to generate coarse change masks via feature correlation, which are further refined and fused to produce the final results.

matched with class-agnostic segmentation masks from SAM based on mask intersection and semantic consistency, yielding refined change masks for  $I_1$  and  $I_2$ . Finally, the two masks are fused using the geometric correspondences to produce the final result. Additionally, we introduce a preprocessing step that mitigates the impact of significant illumination changes before feeding images into the GFM.

### B. Geometric Priors Generation

a) *Geometric Correspondence Generation*: We adopt the Geometric Foundation Model (GFM)  $\mathcal{G}$  to reconstruct geometric information from the input image pair  $\{I_1, I_2\}$ :

$$\mathcal{K}, \mathcal{T}, \mathcal{D} = \mathcal{G}(I_1, I_2), \quad (1)$$

where  $\mathcal{K} = \{K_1, K_2\}$  denotes the camera intrinsics,  $\mathcal{T} = \{T_1, T_2\}$  with  $T_i \in \text{SE}(3)$  denotes the camera extrinsics, and  $\mathcal{D} = \{D_1, D_2\}$  denotes the predicted depth maps. Let  $\Omega_i = [0, W_i] \times [0, H_i]$  denote the pixel domain of image  $I_i$ , with  $(W_i, H_i)$  being its width and height. For each pixel  $p_1 \in \Omega_1$ , we establish its correspondence  $\Phi_{2 \leftarrow 1}(p_1) \in \Omega_2$  by reprojecting  $p_1$  using the estimated depth and camera parameters:

$$x_2^1 = R_{2 \leftarrow 1} D_1(p_1) K_1^{-1} \tilde{p}_1 + t_{2 \leftarrow 1}, \quad (2)$$

$$p_2^1 = \Phi_{2 \leftarrow 1}(p_1) = \pi(K_2 x_2^1), \quad (3)$$

where  $\tilde{p}_1$  is the homogeneous representation of  $p_1$ ,  $R_{2 \leftarrow 1}$  and  $t_{2 \leftarrow 1}$  are the rotation and the translation in the transformation  $T_{j \leftarrow i} = T_j T_i^{-1} = \begin{bmatrix} R_{j \leftarrow i} & t_{j \leftarrow i} \\ 0 & 1 \end{bmatrix}$ ,  $\pi$  is perspective projection

and  $p_2^1$  is the corresponding pixel of  $p_1$  in  $I_2$ . Similarly, the inverse correspondence  $\Phi_{1 \leftarrow 2}$  can be derived. We then determine the overlapping region in  $I_1$  as the set of pixels whose reprojections fall within the valid image domain:

$$\Omega_{1 \rightarrow 2} = \{p_1 \in \Omega_1 \mid p_2^1 \in \Omega_2\}. \quad (4)$$

Symmetrically, we can obtain the overlapping  $\Omega_{2 \rightarrow 1} \subset \Omega_2$ .

b) *Occlusion Detection*: In addition to pixel-level correspondences, geometric information enables the identification of occluded regions across views. Specifically, for each pixel  $p_1 \in \Omega_{1 \rightarrow 2}$  within overlap region in  $I_1$ , we compute its reprojected depth  $D_2^1(p_1)$  at the corresponding location  $p_2^1$  in  $I_2$  using the 3D point  $x_2^1$  obtained from Eq 2:

$$D_2^1(p_1) = e_3^\top x_2^1, e_3 = [0, 0, 1]^\top. \quad (5)$$

The occlusion state of  $p_1$  is then determined by comparing the reprojected depth with the estimated depth  $D_2$  at the corresponding pixel in  $I_2$ :

$$M_{\text{occ}}^1(p_1) = \begin{cases} 1, & \text{if } D_2^1(p_1) - D_2(p_2^1) > \tau, \\ 0, & \text{otherwise,} \end{cases} \quad (6)$$

where  $\tau$  is a depth consistency threshold. To ensure robustness across varying scene scales and depth uncertainties,  $\tau$  is adaptively computed. It combines a geometric constraint proportional to the median scene depth and a statistical constraint based on the median absolute deviation (MAD) of



Fig. 3: Visualization of detected occlusions. Red regions indicate occluded areas.

the depth differences:

$$\tau = \max(\alpha \cdot \text{med}(D_2), \text{med}(\Delta D) + \kappa \cdot \text{MAD}(\Delta D)), \quad (7)$$

where  $\Delta D$  is the depth difference in Eq 6,  $\alpha = 0.03$ , and  $\kappa = 2.5$ . Pixels that satisfy the above condition are visible in  $I_1$  but occluded in  $I_2$ . Similarly, we can obtain  $M_{\text{occ}}^2$ . We visualize a pair of predicted occlusion masks in Fig. 3.

### C. Geometry-guided Change Mask Prediction

a) *Geometry-guided feature correlation*: We feed the image pair  $I_1$  and  $I_2$  into the SAM encoder and extract multi-head key features from its  $l$ -th transformer layer, denoted as  $k_l^1$  and  $k_l^2$ , with  $l = 17$  following [17]. In addition, we also obtain the final-layer embedding of the encoder  $E^1$  and  $E^2$ . All extracted features are then resized to image resolution for subsequent pixel-level operations. For each pixel  $p_1 \in \Omega_{1 \rightarrow 2}$  within overlap, we compute the cosine similarity between  $k_l^1(p_1)$  and its geometrically corresponding feature  $k_l^2(p_2^1)$  across all attention heads. The similarities from different heads are averaged to produce the similarity map  $S_l^1$ . An initial change proposal  $P_{\text{init}}^1$  is derived by applying the Adaptive Threshold Function [17] to  $S_l^1$ . To filter out occluded regions, we refine  $P_{\text{init}}^1$  using:

$$P_{\text{ref}}^1 = P_{\text{init}}^1 \odot (1 - M_{\text{occ}}^1), \quad (8)$$

where  $M_{\text{occ}}^1$  denote the occlusion masks. The refined proposal  $P_{\text{ref}}^1$  also serves as the coarse change mask.

b) *Change Mask Matching & Mask Fusion*: We then obtain the fine change mask in  $I_1$  via Geometric-Semantic Mask Matching [17], which selects class-agnostic SAM masks  $\{\text{mask}_i\}_{\text{SAM}}$  that sufficiently overlap with  $P_{\text{ref}}^1$  and exhibit semantic inconsistency across the two images:

$$M_{\text{change}}^1 = \text{GSM\_Match}(P_{\text{ref}}^1, \{\text{mask}_i\}_{\text{SAM}}). \quad (9)$$

The same procedure is applied symmetrically to  $I_2$ , obtaining  $M_{\text{change}}^2$ . Using the established pixel correspondences  $\Phi_{1 \leftarrow 2}$  between  $I_1$  and  $I_2$ , we map  $M_{\text{change}}^2$  onto the query view and fuse it with  $M_{\text{change}}^1$  to obtain the final change mask:

$$M_{\text{final}} = M_{\text{change}}^1 \cup \Phi_{1 \leftarrow 2}(M_{\text{change}}^2). \quad (10)$$

TABLE I: Properties of datasets after processing for unaligned SCD evaluation. † indicates the dataset is processed into unaligned settings.

	PSCD <sup>†</sup> [6]	ChangeSim <sup>†</sup> [32]	PASLCD <sup>†</sup> [29]
Image pairs	2079	4887	1500
Data source	Real	Synthetic	Real
Resolution	224×224	640×480	504×284
Environment	Outdoor	Indoor	Both
Camera intervals	stride-1/2	stride-5/10/15	stride-5/10/15
Camera translation	✗	✓	✓
Camera rotation	✓	✓	✓
Rotation Range	{10°, 20°}	(0°, 90°]	(0°, 90°]
Aligned pair	✓	✓	✗

### D. Handling Illumination Variations

Significant illumination or color differences between input views can cause GFM reconstruction failures, as shown in Fig. 5(a). To address this, we apply preprocessing only before feeding images into GFM, when the brightness or color distribution difference between the input image pair exceeds a predefined threshold. We experiment with two approaches: Color Transfer [30] and Retinex [31]. Specifically, Retinex extracts the reflectance components of the image pair to suppress illumination variations, while Color Transfer attempts to reduce the RGB distribution gap between the images. Although Retinex produces images that differ significantly in appearance from the originals, it enforces stronger consistency between the input pair. Our experiments (Fig. 5) demonstrate that GFM benefit from Retinex preprocessing, achieving more reliable reconstruction and improved performance under significant illumination changes.

## IV. EXPERIMENTS

### A. Datasets

Previous unaligned SCD evaluations [11] often relied on augmented versions of VL-CMU-CD [5] and PSCD [6]. However, VL-CMU-CD suffers from limited viewpoint variations, noisy ground truth [19], and augmentation-induced artifacts. While PSCD offers viewpoint changes via panoramic images, its standard 15-view cropping [11] only introduces fixed camera rotations without translational motion, and is restricted to street-level scenarios. To ensure a more comprehensive and realistic evaluation, we instead adopt ChangeSim [32] and PASLCD [29], which provide diverse environments and complex 6-DoF camera movements. Table I summarizes the properties of all evaluated datasets.

ChangeSim is a simulated dataset of indoor warehouse scenes, comprising 10 scenes (6 for training and 4 for testing) with 80 sequences, rendered along simulated UAV trajectories with diverse change types and pixel-level annotations. We construct three unaligned subsets by sampling neighboring view intervals of 5, 10, and 15 from the coarsely aligned sequences each representing a different degree of viewpoint difference. For each sampled pair, we compute the visual overlap regions using the provided depth and camera poses and mask out the ground-truth labels outside the overlap.

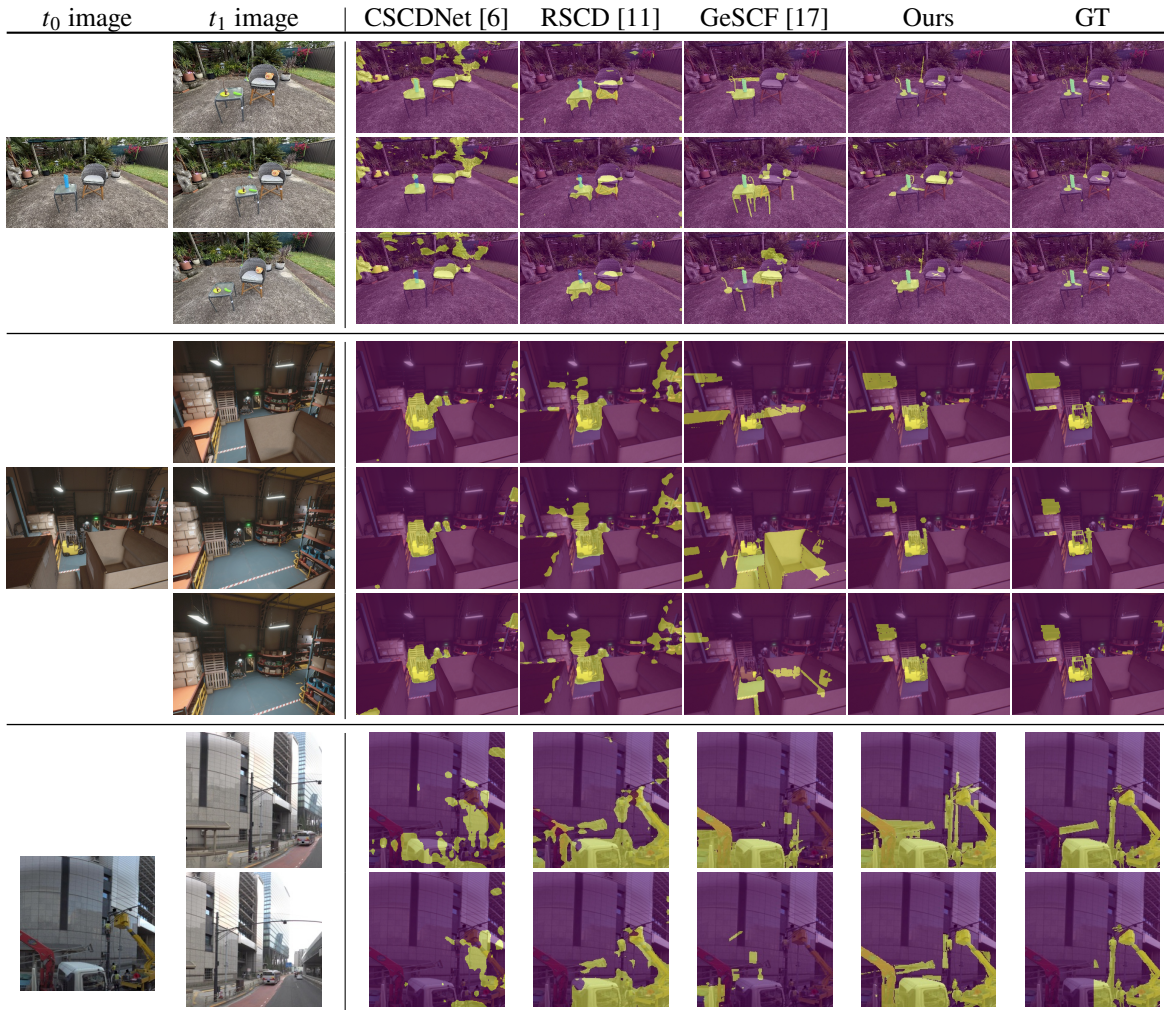


Fig. 4: Visualization of our results compared with baseline methods. From top to bottom, the three groups correspond to PASLCD<sup>†</sup>, ChangeSim<sup>†</sup> (stride-5/10/15), and PSCD<sup>†</sup> (stride-1/2). Our method produces cleaner masks and handles viewpoint changes more robustly.

PASLCD is a recently proposed challenging 3D change detection dataset comprising 10 real-world scenes, including 5 indoor and 5 outdoor. The dataset covers diverse change types, with approximately 70% structural changes and 30% surface-level modifications. For each scene, multi-view images are provided with SfM-estimated [33] camera poses, and three subsets are constructed by pairing each test image with its 5, 10, and 15 nearest reference images. As PASLCD does not provide per-view depth maps, we first reconstruct each scene using 3DGS [34] based on the camera poses and multi-view images. Depth maps for image pairs are then rendered from 3DGS, from which we compute the overlap regions and process the ground truth masks consistent with the processing applied to ChangeSim.

### B. Metrics

Following prior work, we use F1-score, the harmonic mean of precision and recall, as the main evaluation metric, and report mIoU in the ablation study.

### C. Baselines

In our experiments, we compare our proposed framework against three representative training-based SCD methods, CSCDNet [6], C-3PO [19], and RSCD [11], as well as two training-free approaches, ZSCD [18] and GeSCF [17]. In addition, for the ablation study, we construct a baseline by integrating the SOTA optical flow method SEA-RAFT [35] with our Geometry-guided Change Mask Prediction module, replacing geometric correspondences with optical flow and without handling occlusions. Among them, CSCDNet [6] and RSCD [11] explicitly incorporate mechanisms to handle viewpoint misalignment. We retrain all three training-based baselines on ChangeSim [32] and PSCD [6]. To further examine the generalization of training-based methods to varying viewpoint differences, we adopt the viewpoint augmentation strategy introduced in RSCD [11] and design fine-grained experimental groups. Specifically, for each method, we evaluate two variants: one without augmentation and one with a hybrid multi-viewpoint augmentation strategy. On

TABLE II: Quantitative comparison of training-based and training-free scene change detection (SCD) methods on the ChangeSim, PSCD, and PASLCD datasets under varying disparity conditions. The best score in each column is highlighted in **bold**, and the second-best score is indicated with underline. We report the F1-score, higher values indicate better performance.

Training Dataset	Method	Aug	ChangeSim <sup>†</sup> [32]					PSCD <sup>†</sup> [6]				PASLCD <sup>†</sup> [29]			
			Aligned	stride-5	stride-10	stride-15	Avg	Aligned	stride-1	stride-2	Avg	stride-5	stride-10	stride-15	Avg
ChangeSim	CSCDNet [6]	✗	0.368	0.334	0.305	0.280	0.322	0.285	0.233	0.193	0.237	0.122	0.120	0.121	0.121
	CSCDNet [6]	✓	0.354	0.328	0.305	0.285	0.318	0.226	0.196	0.162	0.195	0.103	0.102	0.103	0.103
	C-3PO [19]	✗	0.403	0.351	0.303	0.268	0.331	0.206	0.137	0.121	0.155	<u>0.131</u>	0.120	0.128	0.126
	C-3PO [19]	✓	0.384	0.338	0.299	0.265	0.322	0.227	0.134	0.102	0.154	0.097	0.089	0.090	0.092
	RSCD [11]	✗	0.364	0.315	0.280	0.255	0.303	0.347	0.277	0.209	0.278	0.085	0.083	0.081	0.083
	RSCD [11]	✓	0.392	0.355	0.328	0.305	0.345	0.312	0.276	0.219	0.269	0.073	0.072	0.071	0.072
PSCD	CSCDNet [6]	✗	0.274	0.238	0.211	0.189	0.228	<u>0.533</u>	0.395	0.296	0.408	0.114	0.115	0.116	0.115
	CSCDNet [6]	✓	0.243	0.212	0.188	0.171	0.204	<u>0.503</u>	<u>0.443</u>	<b>0.411</b>	<b>0.452</b>	0.099	0.097	0.100	0.099
	C-3PO [19]	✗	0.303	0.257	0.220	0.192	0.243	0.486	0.337	0.232	0.352	0.076	0.064	0.071	0.070
	C-3PO [19]	✓	0.329	0.267	0.219	0.183	0.249	0.458	0.372	0.298	0.376	0.101	0.102	0.097	0.100
	RSCD [11]	✗	0.277	0.235	0.208	0.187	0.227	<b>0.544</b>	0.356	0.241	0.380	0.129	0.128	0.120	0.126
	RSCD [11]	✓	0.240	0.206	0.183	0.166	0.199	0.490	0.419	<u>0.353</u>	0.421	<u>0.131</u>	<u>0.129</u>	<u>0.131</u>	<u>0.130</u>
Zero-Shot	ZSCD [18]	-	0.288	0.292	0.281	0.266	0.282	0.298	0.254	0.219	0.257	0.120	0.102	0.094	0.105
	GeSCF [17]	-	<u>0.513</u>	<u>0.451</u>	<u>0.396</u>	<u>0.346</u>	<u>0.426</u>	0.330	0.269	0.166	0.255	0.094	0.088	0.088	0.090
	Ours	-	<b>0.569</b>	<b>0.557</b>	<b>0.517</b>	<b>0.481</b>	<b>0.531</b>	<u>0.533</u>	<b>0.454</b>	0.350	<u>0.446</u>	<b>0.351</b>	<b>0.312</b>	<b>0.291</b>	<b>0.318</b>

ChangeSim, aligned pairs are sampled with probability 0.4 and unaligned pairs at three viewpoint intervals (stride-5/10/15) with probability 0.2 each. On PSCD, aligned pairs are sampled with probability 0.6 and unaligned pairs at two viewpoint intervals (stride-1/2) with probability 0.2 each. In addition, on ChangeSim, we introduce a single-viewpoint augmentation setting, where training uses only one viewpoint interval of unaligned pairs, to assess model robustness under unseen viewpoint conditions. All experiments are conducted on NVIDIA V100 GPUs with a batch size of 4, using the default settings of the respective methods. Due to space limitations, the detailed hyperparameter settings of our method for different datasets are available in our open-source repository.

## V. RESULTS

### A. Viewpoint Robustness and Cross-Dataset Generalization

Table II presents the F1-scores of baselines and proposed method across different datasets, while also analyzing the effect of different viewpoint augmentation on training-based models. For the augmentation group, we adopt the aforementioned hybrid augmentation strategy. Furthermore, we evaluate the cross-dataset generalization of these methods. The key observations are summarized as follows:

**1) Robustness to Viewpoint Variations:** Our framework demonstrates strong performance across different viewpoint intervals and datasets, highlighting its robustness to viewpoint variations. On ChangeSim, our method achieves an average F1-score of 0.531, substantially outperforming all baselines and surpassing the second-best method (0.426) by 24.6%. On PASLCD, it consistently achieves the best results across all viewpoint variation levels, with an average score of 0.318 compared to the second-best method (0.130). On PSCD, our method achieves competitive performance across viewpoint separations, obtaining the second-highest average score. This is partly due to PSCD’s cropped panorama pairs with fixed overlaps, which favor some training-based methods that implicitly exploit such regularities. In contrast, our framework

relies on general geometric reasoning rather than dataset-specific patterns, ensuring robust performance across datasets.

**2) Different-Viewpoint Augmentation Does Not Always Help:** Table II shows that different-viewpoint augmentation does not consistently improve training-based methods. RSCD benefits from augmentation in the unaligned setting due to its full-image cross-attention, which provides stronger cross-view matching. This also explains why RSCD improves even on the aligned ChangeSim setting, where residual misalignments remain. In contrast, C-3PO and CSCDNet exhibit mixed behavior: augmentation improves their performance on PSCD but reduces it on ChangeSim. Overall, across nearly all methods, augmentation tends to cause performance drops under the aligned setting.

**3) Strong Cross-dataset Generalization of Our Method:** The results in Table II show that training-based methods suffer substantial performance drops under zero-shot cross-dataset evaluation, and both training-based and training-free baselines struggle on the challenging real-world PASLCD dataset. In contrast, our method maintains consistently strong performance, demonstrating robust cross-dataset generalization.

### B. Qualitative Comparison

In Fig. 4, we present qualitative comparisons across different datasets and viewpoint differences. For ChangeSim and PSCD, training-based models are trained on the corresponding dataset with viewpoint augmentation, while for PASLCD, we use the model trained on PSCD with viewpoint augmentation. The results show that training-based methods produce noisy masks even on in-domain data. For instance, RSCD exhibits a considerable number of false positives on ChangeSim (group 2, 4th column). Their performance further degrades on unseen datasets like PASLCD, where CSCDNet misclassifies complex background regions as changes, and RSCD produces coarse masks despite roughly detecting the changed regions. Meanwhile, the training-free method GeSCF lacks the ability to handle viewpoint variations, resulting in substantial performance degradation under larger viewpoint differences (group

TABLE III: F1-score of training-based models under different training strategies on ChangeSim. Each group represents models trained either without augmentation or with a single-viewpoint interval or with a hybrid augmentation strategy.

Augmentation	Method	Aligned	stride-5	stride-10	stride-15
w/o Aug	RSCD [11]	0.364	0.315	0.280	0.255
	C-3PO [19]	<u>0.403</u>	0.351	0.303	0.268
	CSCDNet [6]	0.368	0.334	0.305	0.280
stride-5	RSCD [11]	0.398	<u>0.356</u>	0.324	0.297
	C-3PO [19]	0.396	0.343	0.301	0.265
	CSCDNet [6]	0.367	0.340	0.316	0.293
stride-10	RSCD [11]	0.388	0.345	0.315	0.291
	C-3PO [19]	0.374	0.328	0.291	0.260
	CSCDNet [6]	0.333	0.313	0.294	0.275
stride-15	RSCD [11]	0.383	0.339	0.311	0.288
	C-3PO [19]	0.380	0.331	0.291	0.260
	CSCDNet [6]	0.324	0.302	0.286	0.2713
Hybrid	RSCD [11]	0.392	0.355	<u>0.328</u>	<u>0.305</u>
	C-3PO [19]	0.384	0.338	0.299	0.265
	CSCDNet [6]	0.354	0.328	0.305	0.285
w/o Aug	Ours	<b>0.569</b>	<b>0.557</b>	<b>0.517</b>	<b>0.481</b>

2, 5th column) and even near-complete failure in some cases (group 3, 5th column). In contrast, our method effectively handles viewpoint changes, maintaining stable performance and producing more accurate masks even under significant viewpoint shifts. Our method can also precisely detect small or challenging changed objects, such as the desktop items and background broom in group 1. Overall, our approach consistently achieves higher accuracy and stronger viewpoint robustness across datasets and varying viewpoint conditions.

### C. Viewpoint Generalization

Beyond evaluating the effect of different viewpoint augmentation strategies on training-based methods, we further examine the single-viewpoint augmentation setting on ChangeSim to assess their generalization ability to unseen viewpoint differences. As shown in Table III, training-based models exhibit very limited viewpoint generalization: when trained with only stride-10 or stride-15 pairs, their performance drops significantly on aligned cases and on smaller viewpoint intervals (stride-5). This indicates that these models fail to learn transferable geometric knowledge from limited single-viewpoint samples. More importantly, even under the corresponding large-viewpoint conditions, single-viewpoint augmentation does not yield a clear advantage over small-viewpoint augmentation. Overall, these results highlight the inherent limitations of training-based approaches in handling multi-view variations. In contrast, our method consistently achieves stable and superior performance across different viewpoint levels without requiring any augmentation.

### D. Ablation Study

We conducted an ablation study on the main modules of our method using the PSCD and PASLCD datasets, evaluating three settings. The first setting (OF + FC) establishes dense




Fig. 5: Effect of preprocessing on GFM reconstruction and change detection. Retinex improves robustness under illumination changes.

pixel correspondences using the optical flow method SEARAF [35], which are then directly used to guide feature correlation and mask fusion. The second setting (GM + FC) replaces optical flow with the proposed geometric correspondences, while filtering out pixels outside the overlapping region. The third setting (GM + FC + Occ) further incorporates occlusion handling. For evaluation, PSCD [6] metrics are reported as the average over stride-1 and stride-2, while PASLCD [29] metrics are averaged over stride-5, stride-10, and stride-15.

TABLE IV: Ablation study on the PSCD and PASLCD datasets.

Method	PSCD <sup>†</sup> [6]		PASLCD <sup>†</sup> [29]	
	F1-Score	mIoU	F1-Score	mIoU
OF + FC	0.244	0.155	0.094	0.052
GM + FC	0.391	0.274	0.237	0.148
GM + FC + Occ	<b>0.402</b>	<b>0.282</b>	<b>0.321</b>	<b>0.218</b>

As shown in Table IV, geometric correspondences significantly outperform optical flow, proving that geometric priors enable more robust overlap identification and correspondence for unaligned SCD. Incorporating occlusion handling further enhances performance, particularly on real-world multi-view datasets PASLCD, where the F1-score improves by 0.084 (35%), underscoring the importance of occlusion reasoning in practical scenarios.

## VI. CONCLUSION

In this work, we tackled the challenging problem of scene change detection under unaligned viewpoints by incorporating geometric priors into SCD for the first time. Leveraging GFM, we construct robust geometric correspondences within visual overlap and detect occlusions, which are then integrated with SAM to form a training-free, geometry-aware framework that delivers reliable performance across diverse viewpoint differences and scenarios. Superior results of our approach demonstrate the critical role of geometric priors in enabling more robust and generalizable SCD. Moreover, by relaxing the conventional assumption of viewpoint alignment,

our approach further broadens the applicability of SCD to more flexible and realistic scenarios. Despite its advantages, our method has limitations. It currently cannot explicitly distinguish between independent moving objects and structural changes, and the performance may be affected by large pose or depth errors from GFM in extreme cases like severe occlusions or textureless regions. Additionally, our framework currently involves a high computational cost, requiring 18GB of GPU memory at a 512×512 resolution. Future work will focus on improving inference efficiency.

## VII. ACKNOWLEDGMENTS

This work was supported by the National Key Research and Development Program of China (2024YFE0203100).

## REFERENCES

- [1] J. Ackermann, J. Kulhanek, S. Cai, X. Haofei, M. Pollefeys, G. Wetzstein, L. Guibas, and S. Peng, “Cl-splats: Continual learning of gaussian splatting with local optimization,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025.
- [2] Z. Lu, J. Ye, and J. Leonard, “3dgs-cd: 3d gaussian splatting-based change detection for physical object rearrangement,” *IEEE Robotics and Automation Letters*, 2025.
- [3] L. Cheng, Z. Qi, Z. Zhou, C. Lu, and G. Xiong, “Lt-gaussian: Long-term map update using 3d gaussian splatting for autonomous driving,” in *2025 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2025, pp. 1427–1433.
- [4] B. Fu, J. Li, B. Zhang, R. Wang, and X. Chen, “Gs-lts: 3d gaussian splatting-based adaptive modeling for long-term service robots,” *arXiv preprint arXiv:2503.17733*, 2025.
- [5] P. F. Alcantarilla, S. Stent, G. Ros, R. Arroyo, and R. Gherardi, “Street-view change detection with deconvolutional networks,” *Autonomous Robots*, vol. 42, no. 7, pp. 1301–1322, 2018.
- [6] K. Sakurada, M. Shibuya, and W. Wang, “Weakly supervised silhouette-based semantic scene change detection,” in *2020 IEEE International conference on robotics and automation (ICRA)*. IEEE, 2020, pp. 6861–6867.
- [7] S. Chen, K. Yang, and R. Stiefelhagen, “Dr-tanet: Dynamic receptive temporal attention network for street scene change detection,” in *2021 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2021, pp. 502–509.
- [8] A. Varghese, J. Gubbi, A. Ramaswamy, and P. Balamuralidhar, “Changenet: A deep learning architecture for visual change detection,” in *Proceedings of the European conference on computer vision (ECCV) workshops*, 2018, pp. 0–0.
- [9] Z. Wang, Y. Zhang, L. Luo, and N. Wang, “Transcd: Scene change detection via transformer-based architecture,” *Optics Express*, vol. 29, no. 25, pp. 41 409–41 427, 2021.
- [10] K. Sakurada, W. Wang, N. Kawaguchi, and R. Nakamura, “Dense optical flow based change detection network robust to difference of camera viewpoints,” *arXiv preprint arXiv:1712.02941*, 2017.
- [11] C.-J. Lin, S. Garg, T.-J. Chin, and F. Dayoub, “Robust scene change detection using visual foundation models and cross-attention mechanisms,” *arXiv preprint arXiv:2409.16850*, 2024.
- [12] S. Wang, V. Leroy, Y. Cabon, B. Chidlovskii, and J. Revaud, “Dust3r: Geometric 3d vision made easy,” in *CVPR*, 2024.
- [13] J. Wang, M. Chen, N. Karaev, A. Vedaldi, C. Rupprecht, and D. Novotny, “Vggt: Visual geometry grounded transformer,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025.
- [14] J. Yang, A. Sax, K. J. Liang, M. Henaff, H. Tang, A. Cao, J. Chai, F. Meier, and M. Feiszli, “Fast3r: Towards 3d reconstruction of 1000+ images in one forward pass,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2025.
- [15] Q. Wang, Y. Zhang, A. Holynski, A. A. Efros, and A. Kanazawa, “Continuous 3d perception model with persistent state,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 10 510–10 522.
- [16] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo et al., “Segment anything,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 4015–4026.
- [17] J.-W. Kim and U.-H. Kim, “Towards generalizable scene change detection,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 24 463–24 473.
- [18] K. Cho, D. Y. Kim, and E. Kim, “Zero-shot scene change detection,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 3, 2025, pp. 2509–2517.
- [19] G.-H. Wang, B.-B. Gao, and C. Wang, “How to reduce change detection to semantic segmentation,” *Pattern Recognition*, 2023.
- [20] L. Leal-Taixé, C. Canton-Ferrer, and K. Schindler, “Learning by tracking: Siamese cnn for robust target association,” in *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 2016, pp. 33–40.
- [21] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [22] H. K. Cheng, S. W. Oh, B. Price, A. Schwing, and J.-Y. Lee, “Tracking anything with decoupled video segmentation,” in *ICCV*, 2023.
- [23] M. Oquab, T. Darcet, T. Moutakanni, H. V. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. Haziza, F. Massa, A. El-Nouby, R. Howes, P.-Y. Huang, H. Xu, V. Sharma, S.-W. Li, W. Galuba, M. Rabbat, M. Assran, N. Ballas, G. Synnaeve, I. Misra, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, “Dinov2: Learning robust visual features without supervision,” 2023.
- [24] S. Zhang, J. Wang, Y. Xu, N. Xue, C. Rupprecht, X. Zhou, Y. Shen, and G. Wetzstein, “Flare: Feed-forward geometry, appearance and camera estimation from uncalibrated sparse views,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 21 936–21 947.
- [25] V. Leroy, Y. Cabon, and J. Revaud, “Grounding image matching in 3d with mast3r,” in *European Conference on Computer Vision*. Springer, 2024, pp. 71–91.
- [26] W. Yu, J. Xing, L. Yuan, W. Hu, X. Li, Z. Huang, X. Gao, T.-T. Wong, Y. Shan, and Y. Tian, “Viewcrafter: Taming video diffusion models for high-fidelity novel view synthesis,” *arXiv preprint arXiv:2409.02048*, 2024.
- [27] Y. Liu, S. Dong, S. Wang, Y. Yin, Y. Yang, Q. Fan, and B. Chen, “Slam3r: Real-time dense scene reconstruction from monocular rgb videos,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 16 651–16 662.
- [28] D. Wu, F. Liu, Y.-H. Hung, and Y. Duan, “Spatial-mlm: Boosting mlm capabilities in visual-based spatial intelligence,” *arXiv preprint arXiv:2505.23747*, 2025.
- [29] C. J. Galappaththige, J. Lai, L. Windrim, D. Dansereau, N. Sunderhauf, and D. Miller, “Multi-view pose-agnostic change localization with zero labels,” in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 11 600–11 610.
- [30] E. Reinhard, M. Adhikhmin, B. Gooch, and P. Shirley, “Color transfer between images,” *IEEE Computer graphics and applications*, vol. 21, no. 5, pp. 34–41, 2002.
- [31] E. H. Land and J. J. McCann, “Lightness and retinex theory,” *Journal of the Optical society of America*, vol. 61, no. 1, pp. 1–11, 1971.
- [32] J.-M. Park, J.-H. Jang, S.-M. Yoo, S.-K. Lee, U.-H. Kim, and J.-H. Kim, “Changesim: Towards end-to-end online scene change detection in industrial indoor environments,” in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021, pp. 8578–8585.
- [33] J. L. Schönberger and J.-M. Frahm, “Structure-from-motion revisited,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [34] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3d gaussian splatting for real-time radiance field rendering,” *ACM Transactions on Graphics*, vol. 42, no. 4, July 2023. [Online]. Available: <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
- [35] Y. Wang, L. Lipson, and J. Deng, “Sea-raft: Simple, efficient, accurate raft for optical flow,” in *European Conference on Computer Vision*. Springer, 2024, pp. 36–54.