

XDG: Accelerated Visual Disambiguation

Gonglin Chen^{1,2} Ben Southall³ Hanyuan Xiao^{1,2} Wenbin Teng^{1,2}
 Haolin Xiong^{1,2} Tianwen Fu^{1,2} Junyi Ouyang^{1,2} Kshitij Singh Minhas³
 Supun Samarasekera³ Rakesh Kumar³ Yajie Zhao^{1,2}

¹USC Institute for Creative Technologies ²University of Southern California ³SRI International




Figure 1. **XDG efficiently removes visually plausible false matches while preserving reconstruction quality.** Top left: an example of a doppelganger pair from two distinct parts of a scene (the louvre museum) from AerialMegaDepth [26]. The images share similar visual structure, but the circled details reveal inconsistent local geometry and appearance. The highlighted blue and green camera poses show where this pair is placed in each reconstruction. A vanilla COLMAP reconstruction is corrupted by multiple similar false matches, while XDG filters doppelganger edges and recovers a camera layout comparable to Doppelgangers++ (DG++) [31] and close to ground truth. Across pairwise and SfM benchmarks, XDG maintains comparable disambiguation performance compared to DG++ while significantly reduce the inference cost.

Abstract

Visual aliasing, also known as the doppelganger problem, remains a key challenge for structure-from-motion (SfM): visually similar but physically distinct surfaces can produce incorrect image matches and degrade reconstruction quality. Previous work mitigates this issue with geometry-aware foundation-model features, but places a heavy transformer classifier on top of the backbone, making large-scale disambiguation expensive. We introduce XDG, an efficient visual disambiguation model designed for scalable SfM. Our key observation is that a 3D foundation model already performs the cross-view geometric reasoning necessary for visual disambiguation, so doppelganger classi-

fication should adapt the backbone representation directly rather than relearn pair reasoning in a separate heavy decoder. XDG fine-tunes Depth Anything 3 with lightweight LoRA adapters and repurposes its camera tokens as compact pair-level classification tokens. A compact MLP head predicts whether a candidate image pair observes the same 3D surface. Extensive experiments show that XDG provides a favorable accuracy–efficiency tradeoff: it remains competitive with the state-of-the-art disambiguation method across pairwise and reconstruction benchmarks and delivers more than a 3× inference speedup. On individual LaMAR scenes containing thousands of images, XDG saves more than 10 hours of visual disambiguation processing. Code is available at <https://github.com/xtcpete/xdg>.

1. Introduction

Accurate 3D reconstruction from unordered image collections is fundamental to computer vision and increasingly critical for producing the geometric supervision required by emerging 3D foundation models. Modern structure-from-motion (SfM) systems rely on robust image retrieval, local feature matching, and geometric verification to build match graphs before reconstruction [22]. Despite these advances, SfM pipelines remain vulnerable to visual aliasing: distinct physical surfaces can share highly similar appearance, causing false image matches that pass local matching and contaminate the reconstruction. This failure mode is formalized as the doppelganger problem [2]. If such edges remain in the match graph, SfM can merge unrelated structures, register cameras incorrectly, or produce fragmented and distorted reconstructions, as shown in Fig. 1. A practical disambiguation system must therefore satisfy two competing requirements. It must be geometrically informed enough to distinguish true overlap from repeated appearance, while also being efficient enough to run on large SfM graphs containing thousands of candidate image pairs. Early work addressed this problem with a CNN-based binary classifier trained on labeled doppelganger pairs, where tentative correspondences from LoFTR [24] were used to estimate an affine warp before classification [2]. Doppelgangers++ (DG++) [31] improved generalization by leveraging MAST3R [11], a large 3D model, and then training transformer-based classifiers [25] on its features. DG++ already disambiguates many challenging cases effectively; its runtime, rather than its accuracy, becomes the practical bottleneck at scale. Its geometry-aware backbone and heavy transformer classifiers must be evaluated over every candidate edge, so processing scenes with thousands of images can require tens of hours.

In this work, we introduce XDG, an efficient geometry-aware doppelganger classifier designed for scalable SfM disambiguation. Our key idea is to adapt the 3D backbone directly and classify from its native representation with a lightweight classifier. XDG uses Depth Anything 3 (DA3) [13] as the geometry-aware backbone and inserts low-rank adaptation (LoRA) modules [9] into its attention and feed-forward projections. Instead of training a separate transformer decoder, XDG repurposes DA3’s stage-wise camera tokens as compact pair-level classification tokens. These tokens already participate in DA3’s alternating local and global feature interactions, so after LoRA fine-tuning, they efficiently summarize the relationship between views. A small MLP head then predicts whether the pair is a true match or a doppelganger edge.

This design differs from prior geometry-aware disambiguation in three ways. First, XDG adapts the foundation model itself to the doppelganger disambiguation task using parameter-efficient LoRA updates, while keeping the pre-

trained DA3 weights frozen. Second, XDG removes the heavy post-backbone transformer classifier and performs classification directly from compact camera tokens. Third, XDG evaluates both image orders for robustness to input order, but fuses the resulting tokens before classification, avoiding the two separate order-specific classification heads used in prior work [31]. Importantly, our ablation in Tab. 6 shows that backbone replacement alone is insufficient: using DA3 with a DG++-style dense-token transformer classifier provides only limited acceleration. The main efficiency gain instead comes from the proposed modules, which classify image pairs directly from the backbone’s native camera tokens, enforce order consistency through symmetric token aggregation, and use a lightweight MLP head for final prediction.

We evaluate XDG on both pairwise visual disambiguation and downstream reconstruction. On the DG [2] and VisymScenes [31] pairwise benchmarks, XDG achieves accuracy comparable to the state-of-the-art method DG++ [31] while reducing per-pair inference time by more than $3\times$. Across multiple reconstruction benchmarks, including the unseen WRIVA [1] and LaMAR [21] datasets, XDG is competitive overall and significantly reduces total processing time. On the LaMAR scenes with 7.5K–9.3K images, XDG saves 10.85–13.91 hours of visual-disambiguation processing per scene compared with DG++. In summary, our contributions are:

1. We propose XDG, an efficient geometry-aware classifier for removing doppelganger edges from SfM match graphs.
2. We show that doppelganger classification does not require a heavy transformer classifier on top of a frozen 3D foundation model; instead, parameter-efficient LoRA adaptation with a simple MLP head can directly adapt the backbone to the task.
3. We demonstrate that XDG achieves competitive pairwise disambiguation and downstream SfM reconstruction performance while substantially reducing disambiguation cost, a key advantage for large-scale scene reconstruction.

2. Related Work

Local Feature Matching. Classical structure-from-motion (SfM) pipelines such as COLMAP [22] rely on local feature matching to establish image correspondences. Traditionally, hand-crafted descriptors such as SIFT [17] have been widely used for matching keypoints across image pairs. More recently, learning-based approaches have substantially improved matching robustness and coverage, leading to stronger reconstruction performance [3–5, 7, 11, 15, 19, 20, 24, 33]. However, these methods are primarily optimized to maximize the number of repeatable correspondences between visually similar regions, often using su-




Figure 2. **XDG architecture.** Given a candidate pair (I_A, I_B) , XDG processes both image orders using a shared DA3-Base backbone with LoRA adapters. Camera tokens from the reversed pass are realigned with the canonical image order and fused stage-wise. The fused tokens are projected, normalized, averaged across stages and views, and then classified by a lightweight MLP as either a true match or a doppelganger edge.

pervision derived from image overlap or geometric consistency. As a result, their training data may contain visually aliased or doppelganger patterns, while their objectives provide limited incentive to reject such false matches. Moreover, local matchers typically operate at the patch or key-point level, without explicitly modeling global image context or higher-level 3D scene consistency. Consequently, although highly effective for standard correspondence estimation, they remain vulnerable to producing spurious matches between distinct yet visually similar surfaces, particularly in doppelganger image pairs.

3D Learning. Recent advances in learning-based 3D reconstruction have produced a wave of methods that replace traditional keypoint matching and iterative optimization with feed-forward neural networks. DUST3R [28] introduced this paradigm by regressing dense point maps from pairs of uncalibrated images, enabling direct recovery of scene geometry and camera poses. Building on this direction, recent methods train large feed-forward transformers to reconstruct 3D geometry from one or many views in a single forward pass: VGGT [27] jointly predicts camera parameters, depth maps, point maps, and feature tracks; π^3 [29] improves robustness with a permutation-equivariant architecture that removes dependence on a fixed reference view; and Depth Anything 3 [13] adopts a unified design for consistent any-view geometry reconstruction. Despite their strong performance, purely feed-forward methods still struggle to match the precision and global consistency required by large-scale SfM pipelines, particularly when image collections contain repeated or visually aliased structures [18]. Their computational and memory costs also grow rapidly with the number of views, making them difficult to use directly as scalable reconstruction systems. More recently, GLUEMAP [18] addresses these limitations by

combining feed-forward local reconstruction with classical global motion averaging and bundle adjustment, while using DG++ [31] as a two-view disambiguation stage to remove harmful doppelganger edges. XDG complements this hybrid reconstruction framework by replacing the costly DG++ stage with a faster classifier that preserves competitive reconstruction accuracy in most settings.

Disambiguation in SfM and Image Matching. Visual disambiguation addresses a failure mode that is complementary to feature matching: instead of only finding correspondences between images, it must decide whether visually plausible correspondences actually arise from the same 3D structure. Classical SfM systems reduce false matches through geometric verification and robust optimization [22], but repeated structures and near-duplicate appearances can still survive these checks and produce incorrect image edges. Earlier ambiguity-handling methods reasoned over the structure of the image graph, loop constraints, duplicate scene components, missing correspondences, or geodesic context [8, 10, 30, 32, 34, 35]; however, these approaches rely largely on hand-designed graph or correspondence cues rather than directly learning from image content. Doppelgangers [2] first formulated this problem explicitly as doppelganger detection and trained a CNN-based binary classifier to reject ambiguous image pairs before reconstruction, improving SfM robustness but remaining limited by an appearance-centric representation. Doppelgangers++ [31] improves visual disambiguation by expanding the training data and using 3D-aware MAST3R [11] features with transformer-based classifiers, achieving stronger performance and integrating with both standard SfM and MAST3R-SfM [6] pipelines. However, its reliance on a heavy 3D foundation model plus duplicated transformer classifiers increases computational cost,

especially when large image collections require evaluating many candidate edges. Our work follows the same direction but takes a different efficiency route: we fine-tune DA3 with LoRA and classify directly from its camera tokens, avoiding an expensive transformer reasoning module after the backbone.

3. Method

Given a candidate image pair (I_A, I_B) from an SfM match graph, our goal is to decide whether the pair observes a common 3D surface or forms a doppelganger edge that should be pruned before reconstruction. Following [2, 31], we formulate this as binary classification, where true matches are positives and ambiguous or non-overlapping pairs are negatives. XDG accelerates this task by removing the heavy post-backbone transformer classifier.

3.1. Model Architecture

As shown in Fig. 2, XDG consists of three main components: a DA3 backbone with LoRA adapters [9], a symmetric token aggregation module, and a lightweight classification head. Given a candidate image pair (I_A, I_B) , we apply the same LoRA-adapted DA3 encoder to both image orders, (I_A, I_B) and (I_B, I_A) . From each pass, XDG extracts stage-wise camera tokens, realigns tokens corresponding to the same image, and fuses them into an order-consistent pair representation. The fused tokens are then projected, average-pooled, and passed to a small MLP to predict whether the pair is a true match or a doppelganger.

DA3 backbone with LoRA adaptation. XDG uses the Base variant of Depth Anything 3 (DA3) [13] as its geometry-aware backbone. DA3 alternates local and global attention over multi-view image tokens, providing the cross-view reasoning needed for doppelganger detection. We freeze the pretrained DA3 weights and insert low-rank adaptation modules [9] into selected linear layers of the backbone.

For a pretrained linear projection W , LoRA adds a trainable low-rank residual:

$$y = Wx + \frac{\alpha}{r}BAx, \quad (1)$$

where $A \in \mathbb{R}^{r \times d_{\text{in}}}$ and $B \in \mathbb{R}^{d_{\text{out}} \times r}$ are trainable, r is the LoRA rank, and α is a scaling factor. The original projection W remains frozen. In our implementation, LoRA is applied to the DA3 attention and feed-forward projections `qkv`, `proj`, `fc1`, and `fc2`. We use rank $r = 8$, $\alpha = 16$, and dropout 0.05. The DA3 camera-token parameters are also left trainable, while the rest of the pretrained backbone parameters remain frozen.

Repurposing the camera tokens. DA3 maintains camera tokens that are injected into the transformer stream at the start of the alternating cross-view blocks for each input image. These tokens are designed to participate in global

multi-view reasoning and therefore provide a natural compact representation across images. For doppelganger detection, we repurpose them as classification tokens.

Let $\mathcal{S} = \{5, 7, 9, 11\}$ denote the four DA3 stages used by XDG, and let $C_e = 1536$ be the camera-token dimension. For the ordered pair (I_A, I_B) , the LoRA-adapted DA3 backbone returns

$$\{(F_s^{A \rightarrow B}, C_s^{A \rightarrow B})\}_{s \in \mathcal{S}} = \Phi_\theta(I_A, I_B), \quad (2)$$

where $F_s^{A \rightarrow B}$ denotes the dense patch tokens and $C_s^{A \rightarrow B} \in \mathbb{R}^{2 \times C_e}$ contains one camera token for each image. The same encoder Φ_θ also processes the reversed pair (I_B, I_A) to produce $C_s^{B \rightarrow A}$ by batching the two image orders together. Unlike DG++ [31], XDG does not pass the dense tokens to a heavy transformer classifier; it classifies the pair using only the compact camera tokens.

Symmetric token aggregation. Because the ordering of an image pair is arbitrary, XDG combines the camera tokens from both input orders. Let π swap the two-view dimension of the reversed output, thereby restoring the canonical (A, B) view order. At each stage, XDG computes

$$\tilde{C}_s = g_C([C_s^{A \rightarrow B}; \pi(C_s^{B \rightarrow A})]), \quad (3)$$

where $[\cdot; \cdot]$ denotes concatenation along the feature dimension and $\tilde{C}_s \in \mathbb{R}^{2 \times C_e}$. The same two-layer fusion MLP g_C is shared across all stages and maps $2C_e \rightarrow 4C_e \rightarrow C_e$, with a GELU activation between its linear layers.

Classification head. The classification head independently projects the fused tokens from each stage to a hidden dimension C_d and applies layer normalization:

$$T_s = \text{LN}_s(W_s^C \tilde{C}_s). \quad (4)$$

It then averages the projected tokens over the four stages and two views:

$$t = \frac{1}{2|\mathcal{S}|} \sum_{s \in \mathcal{S}} \sum_{v \in \{A, B\}} T_{s,v}. \quad (5)$$

Finally, it predicts two logits:

$$z = h(t), \quad (6)$$

where h is a layer-normalized MLP mapping $C_d \rightarrow C_h \rightarrow 2$, with GELU activation and dropout between its linear layers. The two logits correspond to the doppelganger (false-match) class and the true-match class, respectively. We use $C_d = C_h = 768$ and dropout 0.1.

3.2. Implementation Details

Given two-class logits z , the probability of a true match is $p = \text{softmax}(z)_1$, while the doppelganger (false-match) probability is $\text{softmax}(z)_0$. We train XDG with focal

Test data	Method	AP $\uparrow$	ROC AUC $\uparrow$	Prec@Recall =0.85 $\uparrow$	Recall@Prec =0.99 $\uparrow$	Time (ms) $\downarrow$
DG	DG-OG	0.956	0.947	0.910	0.614	191.5
DG	DG++	0.981	0.981	0.982	<u>0.642</u>	<u>118.3</u>
DG	XDG	<u>0.978</u>	<u>0.975</u>	<u>0.963</u>	0.702	34.2
VisymScenes	DG-OG	0.938	0.921	0.831	0.623	177.9
VisymScenes	DG++	<u>0.991</u>	<u>0.990</u>	0.999	<u>0.901</u>	<u>115.1</u>
VisymScenes	XDG	0.995	0.994	<u>0.994</u>	0.905	33.5

Table 1. **Pairwise visual disambiguation.** We compare DG-OG, DG++, and XDG on the DG and VisymScenes test sets. Runtime is reported as the average inference time per image pair in milliseconds. For DG-OG, runtime includes the LoFTR [24] matching step required as input to the classifier. XDG is more than $3\times$ faster than DG++ while maintaining comparable disambiguation performance. Best results are shown in bold, and second-best results are underlined.

loss [14], which down-weights already-easy pairs and focuses learning on ambiguous examples. For a labeled pair with ground-truth class $y \in \{0, 1\}$ and predicted probability $p_y = \text{softmax}(z)_y$, the loss is

$$\mathcal{L}_{\text{focal}} = -(1 - p_y)^\gamma \log p_y, \quad (7)$$

with $\gamma = 1$ in our implementation.

We follow DG++ [31] and train our method using labeled image pairs from the DG dataset [2] and VisymScenes [31]. During training, we randomly flip the input order with probability 0.5 and sample multiple resized resolutions to improve robustness to viewpoint ordering and aspect-ratio changes. We train for 10 epochs using the AdamW optimizer [16] with a learning rate of 1×10^{-4} for all trainable parameters, including the LoRA adapters, camera tokens, fusion module, and classifier head; we set $\beta_1 = 0.9$, $\beta_2 = 0.999$, and the weight decay to 0.05. The learning rate is linearly warmed up for one epoch, then decayed with a step schedule starting at epoch 3 using a decay factor of 0.8 and a minimum learning rate of 1×10^{-8} . On 8 NVIDIA H100 GPUs, the model converges in 4 hours.

4. Experiments

4.1. Pairwise Visual Disambiguation

4.1.1. Datasets and metrics.

We evaluate pairwise visual disambiguation on the DG [2] and VisymScenes [31] test sets following the testing protocol of DG++ [31]. We report average precision (AP), ROC AUC, precision at recall 0.85, recall at precision 0.99, and average inference time per image pair. AP and ROC AUC measure threshold-free ranking quality, while the fixed precision–recall operating points reflect the need to reject false edges without removing too much graph connectivity. We compare XDG with DG-OG [2] and DG++ [31] in the same software environment on an NVIDIA RTX 4090. The runtime for DG-OG includes its required LoFTR [24] pre-processing.




Figure 3. **Qualitative COLMAP reconstructions.** Top: VisymSite0023 from VisymScenes reconstructed with vanilla COLMAP, DG++ filtering, and XDG filtering. Vanilla COLMAP produces a model with incorrectly registered images, while both disambiguation methods separate these cameras into a different model. Bottom: Two scenes from AerialMegaDepth. From left to right, we show the ground-truth camera layout, vanilla COLMAP, DG++ with COLMAP, and XDG with COLMAP. Without doppelganger filtering, physically distinct surfaces with similar appearance are collapsed into one, whereas DG++ and XDG successfully separate them.

4.1.2. Results.

Quantitative results are shown in Tab. 1. On the DG test set, XDG achieves 0.978 AP and 0.975 ROC AUC, approaching the performance of DG++ while substantially outperforming DG-OG. More importantly, XDG improves recall at the strict 0.99 precision operating point from 0.642 to 0.702 compared with DG++. This is particularly favorable for downstream SfM, where falsely removing true matches can disconnect the match graph and harm reconstruction. High precision ensures that most retained edges are reliable, while higher recall preserves more true matches and therefore more graph connectivity. On VisymScenes, XDG slightly outperforms DG++ in AP, ROC AUC, and recall at 0.99 precision, demonstrating that XDG achieves comparable or stronger pairwise disambiguation performance across benchmarks.

Test scene	# SfM-registered images			Inlier ratio $\uparrow$			Time (h) $\downarrow$	
	COLMAP	DG++	XDG	COLMAP	DG++	XDG	DG++	XDG
VisymSite0010	1448	1393	1412	0.704	0.866	0.876	1.49	0.40
VisymSite0023	654	571 + 83	571 + 82	0.863	0.974	0.977	0.89	0.24
VisymSite0028	438	450	206 + 120 + 61	0.797	0.856	0.993	1.14	0.31
VisymSite0042	1336	1319	1306	0.930	0.956	0.958	1.01	0.27
VisymSite0109	814	608 + 117	594 + 170	0.960	0.998	0.981	1.49	0.40

Table 2. **COLMAP reconstruction on VisymScenes.** We compare XDG against vanilla COLMAP and DG++ on five VisymScenes test scenes. We report the number of registered images, geo-alignment inlier ratio, and total visual-disambiguation time per scene. XDG achieves competitive geo-alignment while significantly reducing disambiguation time.

In terms of efficiency, XDG reduces the per-pair runtime of DG++ from 118.3 ms to 34.2 ms on DG and from 115.1 ms to 33.5 ms on VisymScenes. Compared to DG-OG, XDG is also substantially faster, reducing runtime from 191.5 ms to 34.2 ms on DG and from 177.9 ms to 33.5 ms on VisymScenes. Overall, XDG runs more than $3\times$ faster than DG++ and more than $5\times$ faster than DG-OG, while maintaining competitive pairwise accuracy. These results support our central claim that doppelganger disambiguation can be made significantly more efficient without sacrificing performance.

4.2. SfM Reconstruction

XDG operates on the candidate image graph and can be seamlessly integrated into both COLMAP [22] and GLUEMAP [18]. We evaluate SfM performance on four datasets. VisymScenes [31] and AerialMegaDepth [26] are in the training domain of both DG++ and XDG, while WRIVA [1] and LaMAR [21] test out-of-domain generalization. Because VisymScenes has noisy geotags and lacks accurate scene-wide camera ground truth, DG++ evaluates its reconstruction quality using the inlier ratio. While this metric reflects geometric consistency after reconstruction, it provides only an indirect measure of camera accuracy. In contrast, AerialMegaDepth, WRIVA, and LaMAR provide camera calibrations, geospatial references, or laser-scan-aligned poses, allowing pose-based evaluation of the reconstructions.

4.2.1. Datasets

VisymScenes and AerialMegaDepth are reconstructed with COLMAP, while WRIVA and LaMAR are reconstructed with GLUEMAP. We present the training-domain datasets first, followed by the out-of-domain datasets.

VisymScenes. Introduced by DG++ [31], VisymScenes contains 258K images with GPS/IMU metadata collected at 149 sites across 42 cities and 15 countries. It covers landmarks as well as everyday residential, rural, suburban, and business environments, many of which contain repeated structures and visually similar but physically distinct surfaces. We use the same 5 test scenes evaluated by DG++.

AerialMegaDepth. AerialMegaDepth [26] contains 132K images across 137 scenes and co-registers real images from MegaDepth [12] with pseudo-synthetic aerial and

Method	Pose AUC (%) $\uparrow$				Total time (h) $\downarrow$
	@3°	@5°	@10°	@30°	
No disamb.	49.70	55.79	61.18	65.56	—
DG++	57.26	64.45	70.75	75.63	5.28
XDG	57.57	64.68	70.98	75.94	1.20

Table 3. **COLMAP reconstruction on AerialMegaDepth.** Pose AUC is averaged over 8 scenes using only real images as reconstruction inputs, while time is the total visual-disambiguation runtime over all scenes. DG++ and XDG both improve reconstruction accuracy over no doppelganger filtering and achieve nearly identical accuracy, while XDG reduces total disambiguation time from 5.28 to 1.20 hours, saving 4.08 hours ($4.40\times$).

ground-level views rendered from geospatial 3D meshes with known camera poses. The rendered views provide useful pose references and help reduce ambiguity caused by visually similar or duplicated structures, making the dataset suitable for evaluating visual disambiguation. However, because the real images are registered to geospatial meshes rather than captured with ground-truth camera calibration, their poses should be treated as approximate rather than fully reliable. We therefore use AerialMegaDepth primarily as a large-scale, challenging benchmark for reconstruction consistency, evaluating on 8 scenes with duplicated structures and using only the real images as reconstruction inputs.

WRIVA. WRIVA [1] contains calibrated imagery captured from heterogeneous viewpoints and altitudes. Camera locations are georeferenced using RTK-corrected GPS with centimeter-level accuracy, providing reliable camera-position ground truth for evaluating large-scale reconstruction and alignment. We use 34 sequences for evaluation.

LaMAR. LaMAR [21] contains large indoor and outdoor scenes captured along unconstrained AR-device trajectories, with accurate reference poses obtained by registering the trajectories to laser scans. We evaluate on the same benchmark splits used by GLUEMAP [18].

4.2.2. Implementation Details and Metrics

For each reconstruction experiment, we run visual disambiguation before reconstruction and keep all remaining stages fixed.

Method	Pose AUC (%)↑				Total time (h)↓
	@3°	@5°	@10°	@30°	
No disamb.	13.16	19.49	29.46	47.52	—
DG++	22.06	31.95	45.14	62.22	9.86
XDG	21.24	31.33	45.43	63.73	3.06

Table 4. **GLUEMAP reconstruction on WRIVA.** Pose AUC is averaged across 34 sequences, while time is the total visual-disambiguation runtime over the complete benchmark. Higher AUC is better. Both DG++ and XDG substantially improve upon reconstruction without doppelganger detection, with DG++ slightly outperforming XDG in the lowest error regime at the cost of more than triple the runtime.

Scene (# images)	Method	Pose AUC (%)↑				Total time (h)↓
		@3°	@5°	@10°	@30°	
CAB (6,587)	DG++	3.70	8.90	27.57	65.77	14.18
	XDG	3.70	8.11	20.00	51.44	4.29
HGE (7,553)	DG++	36.52	56.54	76.91	91.80	15.73
	XDG	38.16	55.87	76.18	90.73	4.88
LIN (9,319)	DG++	43.58	60.48	74.45	85.37	20.02
	XDG	43.79	60.91	74.81	85.57	6.11

Table 5. **GLUEMAP reconstruction on LaMAR.** We replace GLUEMAP’s DG++ disambiguator with XDG and keep all other stages fixed. DG++ generalizes slightly better on the indoor CAB scene, whereas XDG performs comparably on the outdoor HGE and LIN scenes while reducing total disambiguation time by 3.27×.

COLMAP. We follow the DG++ protocol and use COLMAP’s vocabulary-tree matcher [23] with default settings and confidence threshold $\tau = 0.8$.

GLUEMAP. We replace GLUEMAP’s default DG++ two-view stage with XDG while keeping retrieval, feed-forward local reconstruction, global mapping, and refinement unchanged. We use the benchmark configuration reported by GLUEMAP for LaMAR, and default settings for WRIVA.

Metrics. On VisymScenes, we report the number of registered images and the geo-alignment inlier ratio used by DG++ [31], since the GPS annotations are noisy. Specifically, reconstructed camera centers are aligned to GPS locations with RANSAC, and the inlier ratio measures the fraction of registered images that are consistent with the estimated alignment. For reconstructions containing multiple disconnected components, we report the weighted inlier ratio

$$\text{IR}_{\text{weighted}} = \frac{\sum_k n_k \text{IR}_k}{\sum_k n_k}, \quad (8)$$

where n_k and IR_k denote the number of registered images and the inlier ratio of component k , respectively. On AerialMegaDepth, WRIVA, and LaMAR, we report $\text{AUC}@X^\circ$, defined as the normalized area under the re-




Figure 4. **Qualitative GLUEMAP reconstruction.** We show two WRIVA sequences, with the top two rows corresponding to one sequence and the bottom two rows to another. For each sequence, the first row shows the ground-truth camera layout and reconstructions without disambiguation, with DG++, and with XDG; blue and green cameras indicate the ambiguous image pair shown below. The second row shows the image pair and the corresponding alignment to ground truth.

call curve of pairwise relative-pose error up to the angular threshold X° . For each image pair, the pose error is computed as the maximum of the relative rotation error and the translation-direction error. Image pairs for which one or both cameras are not registered are counted as failures and therefore have an error of infinity. Tight thresholds emphasize pose accuracy, while looser thresholds also reflect reconstruction completeness. We use $X \in \{3, 5, 10, 30\}$. All reported times correspond to visual-disambiguation runtime and are measured in hours. For AerialMegaDepth, WRIVA, and LaMAR, the tables report total inference time summed over every scene or sequence in the respective benchmark rather than an average per scene. Experiments on these three datasets are run in the same software environment on a cluster with 8 NVIDIA A100 40 GB GPUs.

4.2.3. Results

VisymScenes. Tab. 2 reports the largest COLMAP models after disambiguation. The “+” symbol denotes split reconstruction components. XDG achieves reconstruction quality comparable to DG++ while substantially reducing the required disambiguation time. Across the five VisymScenes scenes, XDG registers a similar number of images and achieves similar inlier ratios. Qualitative results in Fig. 3 further show that XDG removes incorrectly registered cameras and produces camera layouts visually comparable to DG++ on VisymSite0023. Averaged across scenes, XDG reduces disambiguation time from 1.20 hours to 0.32 hours,

Variant	AP $\uparrow$	ROC AUC $\uparrow$	Time (ms) $\downarrow$
<i>DG</i>			
DA3 + transformer	0.9617	0.9596	111.9
+ feature aggregation	0.9591	0.9557	55.4
+ MLP (dense tokens)	0.9586	0.9551	38.9
+ camera tokens	0.9540	0.9511	32.9
+ LoRA (XDG)	0.9780	0.9750	34.2
<i>VisymScenes</i>			
DA3 + transformer	0.9919	0.9916	110.7
+ feature aggregation	0.9922	0.9908	54.6
+ MLP (dense tokens)	0.9912	0.9904	38.8
+ camera tokens	0.9905	0.9896	32.4
+ LoRA (XDG)	0.9950	0.9940	33.5

Table 6. **Cumulative ablation of XDG.** Starting from a frozen DA3 backbone with a dense-token transformer classifier following DG++, each row adds the listed change. The best accuracy and lowest latency in each dataset block are highlighted. The full ablation table is provided in the supplementary material.

corresponding to a $3.7\times$ speedup.

AerialMegaDepth. Tab. 3 reports COLMAP reconstruction accuracy across 8 scenes. Both XDG and DG++ substantially improve pose AUC over reconstruction without disambiguation and perform nearly identically. As shown qualitatively in Fig. 3, doppelganger filtering prevents visually similar but physically distinct structures from being incorrectly collapsed, and XDG preserves the reconstruction quality of DG++. The main advantage of XDG is efficiency: under the same inference setup, XDG reduces total visual-disambiguation time over all 8 scenes from 5.28 to 1.20 hours, saving 4.08 hours.

WRIVA. Tab. 4 reports out-of-domain pose accuracy averaged over 34 sequences, and Fig. 4 shows qualitative results on two representative sequences. Both DG++ and XDG substantially outperform reconstruction without disambiguation. DG++ performs best at the tightest thresholds, suggesting stronger generalization to these unseen scenes, while XDG achieves slightly higher AUC@10° and AUC@30° and preserves most of the reconstruction completeness with a substantially lower inference cost. In total, XDG reduces visual-disambiguation time over the 34 sequences from 9.86 to 3.06 hours, saving 6.80 hours. Qualitatively, both DG++ and XDG remove most visually plausible false matches that would otherwise distort the reconstruction, producing camera trajectories that align closely with the ground truth.

LaMAR. Tab. 5 reports results on LaMAR [21]. On CAB, an indoor sequence, DG++ generalizes better: it ties XDG at AUC@3° and performs better at the looser thresholds. On the outdoor HGE and LIN scenes, the two methods are comparable. Across all three scenes, XDG reduces total visual-disambiguation time from 49.93 to 15.28 hours, saving 34.65 hours at a $3.27\times$ speedup. The scaling benefit is already substantial on individual scenes: HGE contains 7,553 images and saves 10.85 hours, while LIN contains 9,319 images and saves 13.91 hours. These results show

that DG++ already works well across many cases, but its runtime is a bottleneck for large image collections; XDG directly alleviates this bottleneck.

4.3. Ablation

We ablate the main design choices of XDG on pairwise classification in Tab. 6. The study is cumulative: each row adds one modification to the configuration above it. We start from a **DA3–transformer baseline**, which replaces the MAST3R [11] backbone in DG++ [31] with frozen DA3 [13]. We then introduce **feature aggregation**, which aligns and fuses features from the two input orders, removing the need for duplicated classification heads. Next, we replace the transformer classifier with a lightweight **MLP head** to further reduce inference cost. The **camera-token** variant further compresses the representation by replacing dense patch tokens with DA3 camera tokens for classification. Finally, we add **LoRA** adapters to adapt the frozen DA3 backbone to the disambiguation task. The last row corresponds to the full XDG model reported in Tab. 1.

Feature aggregation substantially reduces latency by removing redundant classification over the two input orders, while the MLP head further lowers the computational cost with little change in overall accuracy. Using DA3 camera tokens as the pair-level representation gives the fastest model, indicating that these tokens provide a compact and effective summary for pairwise disambiguation. However, this compression also removes some discriminative signal present in the dense tokens, leading to a small accuracy drop compared with the preceding variants.

LoRA fine-tuning recovers this lost discriminative signal with little additional runtime. Compared with the camera-token variant, the full XDG model consistently improves AP and ROC AUC on both test sets. The final model remains more than three times faster than the baseline while achieving the best accuracy. Overall, the ablation separates the roles of the main proposed components: feature aggregation, the MLP head, and camera-token classification provide the acceleration, while LoRA fine-tuning improves the accuracy needed for robust SfM disambiguation.

5. Conclusion

We introduced XDG, an efficient visual-disambiguation model for removing doppelganger edges in large-scale 3D reconstruction. The state-of-the-art method DG++ already resolves many difficult ambiguities effectively, but its runtime becomes a bottleneck when scaled to large image collections with thousands of images. By adapting a 3D foundation model with LoRA and replacing heavy transformer classifiers with lightweight token aggregation and an MLP head, XDG achieves performance comparable to DG++ while reducing disambiguation time by more than $3\times$ across pairwise and downstream reconstruction tasks.

This speedup saves 34.65 hours in total over the LaMAR benchmark and more than 6 hours on the WRIVA dataset. Our extensive experiments show that XDG can make visual disambiguation significantly more efficient while preserving reconstruction accuracy.

Acknowledgments

Supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DOI/IBC) contract number 140D0423C0075. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

References

- [1] Myron Brown, Michael Chan, and Michael Twardowski. WRIVA public data. IEEE DataPort, 2024. 2, 6, 11
- [2] Ruojin Cai, Joseph Tung, Qianqian Wang, Hadar Averbuch-Elor, Bharath Hariharan, and Noah Snavely. Doppelgangers: Learning to disambiguate images of similar structures. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 34–44, 2023. 2, 3, 4, 5, 11
- [3] Gonglin Chen, Tianwen Fu, Haiwei Chen, Wenbin Teng, Hanyuan Xiao, and Yajie Zhao. RDD: Robust feature detector and descriptor using deformable transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6394–6403, 2025. 2
- [4] Gonglin Chen, Jinsen Wu, Haiwei Chen, Wenbin Teng, Zhiyuan Gao, Andrew Feng, Rongjun Qin, and Yajie Zhao. Geometry-aware feature matching for large-scale structure from motion. In *2025 International Conference on 3D Vision (3DV)*, pages 34–43. IEEE, 2025.
- [5] 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. 2
- [6] Bardienus Pieter Duisterhof, Lojze Zust, Philippe Weinzaepfel, Vincent Leroy, Yohann Cabon, and Jerome Revaud. MAST3R-SfM: A fully-integrated solution for unconstrained structure-from-motion. In *2025 International Conference on 3D Vision (3DV)*, pages 1–10. IEEE, 2025. 3
- [7] Xingyi He, Jiaming Sun, Yifan Wang, Sida Peng, Qixing Huang, Hujun Bao, and Xiaowei Zhou. Detector-free structure from motion. *CVPR*, 2024. 2
- [8] Jared Heinly, Enrique Dunn, and Jan-Michael Frahm. Correcting for duplicate scene structure in sparse 3D reconstruction. In *European Conference on Computer Vision*, pages 780–795. Springer, 2014. 3
- [9] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 2, 4, 11
- [10] Nianjuan Jiang, Ping Tan, and Loong-Fah Cheong. Seeing double without confusion: Structure-from-motion in highly ambiguous scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1458–1465, 2012. 3
- [11] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3D with MAST3R. In *European conference on computer vision*, pages 71–91. Springer, 2024. 2, 3, 8
- [12] Zhengqi Li and Noah Snavely. MegaDepth: Learning single-view depth prediction from internet photos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2041–2050, 2018. 6
- [13] Haotong Lin, Sili Chen, Junhao Liew, Donny Y Chen, Zhenyu Li, Guang Shi, Jiashi Feng, and Bingyi Kang. Depth Anything 3: Recovering the visual space from any views. *arXiv preprint arXiv:2511.10647*, 2025. 2, 3, 4, 8, 11
- [14] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017. 5
- [15] Philipp Lindenberger, Paul-Edouard Sarlin, and Marc Pollefeys. LightGlue: Local Feature Matching at Light Speed. In *ICCV*, 2023. 2
- [16] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. 5
- [17] David G Lowe. Distinctive image features from scale-invariant keypoints. *International journal of computer vision*, 60(2):91–110, 2004. 2
- [18] Linfei Pan, Johannes Lutz Schönberger, and Marc Pollefeys. Global structure-from-motion meets feedforward reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026. 3, 6, 11
- [19] Ignacio Rocco, Mircea Cimpoi, Relja Arandjelović, Akihiko Torii, Tomas Pajdla, and Josef Sivic. Neighbourhood consensus networks. In *Advances in Neural Information Processing Systems*, 2018. 2
- [20] 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
- [21] Paul-Edouard Sarlin, Mihai Dusmanu, Johannes Lutz Schönberger, Pablo Speciale, Lukas Gruber, Viktor Larsson, Ondrej Miksik, and Marc Pollefeys. LaMAR: Benchmarking localization and mapping for augmented reality. In *European Conference on Computer Vision*, 2022. 2, 6, 8
- [22] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016. 2, 3, 6, 11
- [23] Johannes Lutz Schönberger, True Price, Torsten Sattler, Jan-Michael Frahm, and Marc Pollefeys. A vote-and-verify strat-

- egy for fast spatial verification in image retrieval. In *Asian Conference on Computer Vision (ACCV)*, 2016. [7](#)
- [24] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. LoFTR: Detector-free local feature matching with transformers. *CVPR*, 2021. [2](#), [5](#)
- [25] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, 2017. [2](#)
- [26] Khiem Vuong, Anurag Ghosh, Deva Ramanan, Srinivasa Narasimhan, and Shubham Tulsiani. AerialMegaDepth: Learning aerial-ground reconstruction and view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21674–21684, 2025. [1](#), [6](#), [11](#)
- [27] Jianyuan Wang, Minghao Chen, Nikita Karaev, Andrea Vedaldi, Christian Rupprecht, and David Novotny. VGGT: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5294–5306, 2025. [3](#)
- [28] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. DUST3R: Geometric 3D vision made easy. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20697–20709, 2024. [3](#)
- [29] Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen, Jiangmiao Pang, Chunhua Shen, and Tong He. π^3 : Permutation-equivariant visual geometry learning. *arXiv preprint arXiv:2507.13347*, 2025. [3](#)
- [30] Kyle Wilson and Noah Snavely. Network principles for sfm: Disambiguating repeated structures with local context. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 513–520, 2013. [3](#)
- [31] Yuanbo Xiangli, Ruojin Cai, Hanyu Chen, Jeffrey Byrne, and Noah Snavely. Doppelgangers++: Improved visual disambiguation with geometric 3D features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 27166–27175, 2025. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [11](#)
- [32] Qingan Yan, Long Yang, Ling Zhang, and Chunxia Xiao. Distinguishing the indistinguishable: Exploring structural ambiguities via geodesic context. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3836–3844, 2017. [3](#)
- [33] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. Lift: Learned invariant feature transform. In *European Conference on Computer Vision*, pages 467–483. Springer, 2016. [2](#)
- [34] Christopher Zach, Arnold Irschara, and Horst Bischof. What can missing correspondences tell us about 3D structure and motion? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008. [3](#)
- [35] Christopher Zach, Manfred Klopschitz, and Marc Pollefeys. Disambiguating visual relations using loop constraints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1426–1433, 2010. [3](#)

Supplementary Material

A. Additional Reconstruction Results

A.1. WRIVA

COLMAP. We additionally evaluate COLMAP [22] on the 34 WRIVA [1] sequences in Tab. 7. In contrast to the GLUEMAP [18] results in the main paper, both DG++ [31] and XDG reduce pose AUC relative to vanilla COLMAP.

Method	Pose AUC (%)↑			
	@3°	@5°	@10°	@30°
No disamb.	16.60	23.33	30.94	38.66
DG++	14.22	19.68	25.75	31.70
XDG	14.16	19.45	25.53	31.46

Table 7. **COLMAP reconstruction on WRIVA.** Pose AUC is averaged across 34 sequences. Higher is better. On this challenging dataset, filtering candidate edges with either DG++ or XDG reduces reconstruction accuracy.

WRIVA contains substantial variation in viewpoint, altitude, and appearance. Under these conditions, COLMAP’s conventional feature matching may fail to establish a sufficiently complete set of reliable correspondences. Visual disambiguation then prunes an already sparse match graph, potentially leaving too few tracks for accurate and complete reconstruction. These results suggest that correspondence formation, rather than false-edge removal alone, is the primary bottleneck for the COLMAP pipeline on WRIVA and motivate our use of the more robust GLUEMAP pipeline for the main WRIVA experiments.

A.2. AerialMegaDepth

GLUEMAP. We also evaluate GLUEMAP reconstruction on the eight AerialMegaDepth [26] scenes in Tab. 8. Both DG++ and XDG substantially improve pose AUC over reconstruction without visual disambiguation, and the two methods perform comparably across all thresholds.

Method	Pose AUC (%)↑			
	@3°	@5°	@10°	@30°
No disamb.	19.40	28.21	41.12	59.19
DG++	34.09	46.05	60.94	77.81
XDG	35.07	46.98	60.88	77.04

Table 8. **GLUEMAP reconstruction on AerialMegaDepth.** Pose AUC is averaged across eight scenes using only real images as reconstruction inputs. Higher is better. Both visual-disambiguation methods substantially improve reconstruction accuracy and perform comparably.

These GLUEMAP results are lower than the COLMAP results reported in the main paper. Some real AerialMegaDepth images are difficult to register or reconstruct without the rendered images provided in the original dataset. GLUEMAP attempts to incorporate all input images, including these difficult views, which can introduce inaccurate camera estimates and reduce the aggregate pose AUC. Nevertheless, the clear improvement from both DG++ and XDG shows that visual disambiguation remains beneficial within this reconstruction setting.

B. More Ablation Results

Tab. 9 expands the cumulative ablation study from the main paper with the fixed precision–recall operating points omitted there for space. The ablations are evaluated on DG [2] and VisymScenes [31]. Starting from a frozen DA3 [13] backbone with a dense-token transformer classifier, each row adds the listed modification. Feature aggregation, the MLP head, and camera-token classification reduce inference cost, while LoRA [9] adapts the backbone to recover task accuracy.

Tab. 10 reports additional controlled ablations of backbone capacity, training augmentations, and bidirectional processing. Each variant changes only the named factor relative to full XDG.

Backbone capacity. DA3-Small roughly halves XDG’s latency but reduces accuracy, especially on DG. DA3-Large provides a small gain on DG and improves several VisymScenes operating-point metrics, but increases latency by approximately $2.4\times$ over DA3-Base. These results support DA3-Base as the best overall accuracy–efficiency compromise.

Training strategy. Removing either random input-order flipping or multi-resolution sampling consistently lowers AP and ROC AUC on both test sets. The degradation from removing multi-resolution training is particularly visible on VisymScenes, indicating that scale and aspect-ratio diversity during training improves generalization.

Order robustness. Processing only one image order reduces latency from 34.2 to 19.7 ms on DG and from 33.5 to 19.6 ms on VisymScenes. However, it lowers AP and ROC AUC on both datasets relative to the full bidirectional model. Fig. 5 illustrates the practical effect on an example pair. XDG predicts a matching confidence of 0.9727

Variant	AP $\uparrow$	ROC AUC $\uparrow$	Prec@Recall =0.85 $\uparrow$	Recall@Prec =0.99 $\uparrow$	Time/pair (ms) $\downarrow$
<i>DG</i>					
DA3 + transformer	0.9617	0.9596	0.9214	0.5195	111.9
+ feature aggregation	0.9591	0.9557	0.9164	0.5427	55.4
+ MLP (dense tokens)	0.9586	0.9551	0.9184	0.4959	38.9
+ camera tokens	0.9540	0.9511	0.8992	0.5191	32.9
+ LoRA (XDG)	0.9780	0.9750	0.9630	0.7020	34.2
<i>VisymScenes</i>					
DA3 + transformer	0.9919	0.9916	0.9877	0.8365	110.7
+ feature aggregation	0.9922	0.9908	0.9963	0.9019	54.6
+ MLP (dense tokens)	0.9912	0.9904	0.9884	0.8434	38.8
+ camera tokens	0.9905	0.9896	0.9928	0.8855	32.4
+ LoRA (XDG)	0.9950	0.9940	0.9940	0.9050	33.5

Table 9. **Full cumulative ablation of XDG.** This table expands the ablation study in the main paper with fixed precision–recall operating points. Each row adds the listed modification. The best value for each metric and the lowest latency in each dataset block are highlighted.

Variant	AP $\uparrow$	ROC AUC $\uparrow$	Prec@Recall =0.85 $\uparrow$	Recall@Prec =0.99 $\uparrow$	Time/pair (ms) $\downarrow$
<i>DG</i>					
XDG (DA3-Base)	0.9780	0.9750	0.9630	0.7020	34.2
<i>Backbone size</i>					
DA3-Small	0.9574	0.9549	0.9163	0.4277	16.3
DA3-Large	0.9791	0.9762	0.9635	0.7110	83.2
<i>Training strategy</i>					
No random order flipping	0.9754	0.9715	0.9602	0.7535	32.7
No multi-resolution training	0.9750	0.9712	0.9663	0.7312	32.8
<i>Order robustness</i>					
Single-order inference	0.9751	0.9721	0.9535	0.6990	19.7
<i>VisymScenes</i>					
XDG (DA3-Base)	0.9950	0.9940	0.9940	0.9050	33.5
<i>Backbone size</i>					
DA3-Small	0.9908	0.9894	0.9927	0.8667	15.5
DA3-Large	0.9948	0.9952	0.9971	0.9352	83.3
<i>Training strategy</i>					
No random order flipping	0.9934	0.9928	0.9927	0.8862	32.8
No multi-resolution training	0.9912	0.9897	0.9905	0.8522	32.8
<i>Order robustness</i>					
Single-order inference	0.9934	0.9925	0.9985	0.9157	19.6

Table 10. **Additional ablations of XDG.** XDG uses DA3-Base, random order flipping, multi-resolution training, and both input orders. Each other row changes only the named factor. The best value for each metric and the lowest latency in each dataset block are highlighted.

in both directions, whereas the single-order model’s confidence changes from 0.7734 for $A \rightarrow B$ to 0.8906 for $B \rightarrow A$. At a threshold of 0.8, this difference causes the same candidate edge to be rejected or retained solely because of the arbitrary input order. We therefore retain bidirectional processing to make predictions order-robust, accepting its additional latency.

C. Detailed Scene and Sequence Results

C.1. AerialMegaDepth: Per-Scene Results

Tab. 11 provides detailed per-scene COLMAP results and visual-disambiguation runtime for AerialMegaDepth.

The gains are concentrated on scenes where repeated structures cause severe ambiguity, most notably Cologne Cathedral and the Louvre Museum. On easier scenes, all three configurations are similar. Across scenes, XDG closely tracks DG++ while requiring substantially less visual-disambiguation time.



(a) Image A



(b) Image B

Model	$p(\text{match} \mid A \rightarrow B)$	$p(\text{match} \mid B \rightarrow A)$	$ \Delta $	Decision at $\tau = 0.8$
XDG	0.9727	0.9727	0.0000	accept / accept
Single-order inference	0.7734	0.8906	0.1172	reject / accept

Figure 5. **Effect of bidirectional order-robust processing.** We evaluate the same image pair in both input orders. XDG aggregates features from both orders and produces the same matching confidence after swapping the inputs. The single-order model is sensitive to the input order: with a decision threshold of $\tau = 0.8$, it rejects the $A \rightarrow B$ ordering but accepts $B \rightarrow A$.

Scene	# Images	No disamb.	DG++	XDG	Time (min)↓	
					DG++	XDG
Adler Planetarium	115	79.64/86.02/91.70/96.90	79.04/85.68/91.53/95.89	79.15/85.73/91.95/97.14	10.74	2.37
Belfry of Bruges	348	39.79/49.92/59.49/67.33	30.31/37.78/45.00/50.99	30.54/38.03/45.22/51.22	22.76	4.99
Brussels Town Hall	630	66.24/77.23/87.05/94.77	66.99/77.62/87.12/94.59	66.85/77.63/87.25/94.83	51.55	10.87
Cologne Cathedral	580	50.19/53.70/56.62/58.84	81.07/87.87/93.26/97.26	80.75/87.67/93.15/97.22	50.50	10.90
Florence Baptistery	691	55.31/60.91/65.41/68.61	59.79/66.14/71.33/75.03	61.40/67.33/72.09/75.47	62.95	13.27
The Louvre Museum	297	20.85/24.97/28.53/31.74	55.16/66.85/77.06/84.96	55.92/67.29/77.39/85.19	21.93	4.66
Ponte di Rialto	369	44.92/48.75/52.23/55.10	44.89/48.73/52.23/55.12	45.21/49.00/52.47/55.42	42.83	10.24
St. Vitus Cathedral	500	40.67/44.79/48.41/51.17	40.87/44.93/48.48/51.20	40.76/44.80/48.34/51.05	54.04	14.90

Table 11. **Per-scene COLMAP reconstruction on AerialMegaDepth.** The image count denotes the number of real input images for each scene. Each AUC cell reports pose AUC (%) as the tuple AUC@3° / @5° / @10° / @30°. Time is the visual-disambiguation runtime for each scene. Results cover all eight evaluation scenes.

C.2. WRIVA: Per-Sequence Results

Tab. 12 provides the complete per-sequence GLUEMAP results underlying the aggregate WRIVA numbers in the main paper. We abbreviate each sequence by its unique task-view-scene-run identifier; the number of input images is shown in parentheses.

The sequence-level results reveal substantial variation across capture conditions. Visual disambiguation provides large gains on many difficult PTZ and varying-altitude sequences, while a few small or already well-reconstructed sequences favor no filtering. XDG and DG++ exhibit similar overall behavior, with complementary strengths across individual sequences.

D. More Qualitative Results

Figs. 6–8 provide additional top-down visualizations of AerialMegaDepth COLMAP reconstructions. For each

scene, we show the ground-truth camera layout and reconstructions without visual disambiguation, with DG++, and with XDG. Red frusta denote reconstructed cameras. When a method produces multiple substantial reconstruction components, we stack the aligned components vertically within the same column.

Figs. 9–11 show six representative WRIVA GLUEMAP reconstructions that are distinct from the two examples in the main paper. We select sequences for which both DG++ and XDG consistently outperform reconstruction without visual disambiguation. Across these examples, visual disambiguation removes geometrically inconsistent matches and yields camera layouts that better agree with the ground truth.

Sequence (# images)	No disamb.	DG++	XDG	Time (min)↓	
				DG++	XDG
t04-v10-s01-r01 (600)	2.82/8.81/25.41/52.73	14.69/27.29/52.25/82.88	25.56/45.45/69.93/89.21	62.43	19.59
t09-v03-s03-r05 (177)	1.71/2.91/6.17/14.28	9.68/13.50/19.57/30.16	8.29/12.49/18.87/28.78	16.15	5.09
t01-v01-s04-r08 (50)	60.99/73.07/82.51/91.24	59.05/70.09/81.20/93.22	65.43/78.99/89.36/96.45	1.74	0.56
t02-v02-s00-r02 (25)	41.52/55.49/71.78/90.17	44.89/57.77/68.30/75.35	39.03/57.17/75.67/91.41	0.44	0.14
t03-v03-s04-r02 (500)	3.19/7.86/18.46/40.02	44.60/60.81/77.88/91.72	42.66/59.47/77.33/91.55	33.90	16.51
t09-v04-s01-r10 (182)	2.98/6.63/14.83/36.85	7.85/17.29/38.47/69.98	13.25/28.14/55.56/84.27	17.34	5.61
t04-v05-s03-r02 (150)	8.67/12.92/21.77/53.56	8.44/11.59/18.26/34.03	12.38/20.93/41.31/76.84	12.57	3.90
t02-v03-s00-r02 (17)	14.06/25.79/49.13/77.65	13.30/24.67/48.35/77.39	13.25/24.51/48.26/77.36	0.19	0.06
t09-v05-s00-r02 (474)	0.68/1.22/2.63/8.85	15.16/23.14/38.38/66.47	16.90/24.86/38.45/69.08	45.60	14.47
t09-v03-s02-r05 (195)	0.61/0.82/1.55/6.43	6.80/9.45/13.41/21.12	6.33/9.29/14.96/28.57	18.07	6.00
t03-v04-s00-r02 (150)	12.21/29.00/56.12/82.49	23.99/38.05/61.83/86.05	24.26/38.03/61.64/85.86	6.20	4.12
t04-v09-s00-r04 (39)	3.85/12.06/34.42/63.44	5.27/15.56/34.13/56.20	3.32/8.58/19.84/43.11	0.61	0.34
t09-v05-s01-r02 (444)	0.44/0.75/1.59/5.20	14.89/20.83/31.94/57.97	18.31/25.28/37.35/58.86	42.06	13.38
t09-v04-s00-r10 (198)	1.91/4.58/12.28/34.48	8.93/20.24/42.27/71.77	13.09/27.43/53.47/82.05	19.25	6.12
t09-v06-s03-r02 (167)	20.71/30.56/44.88/66.92	43.02/56.02/68.44/81.51	39.19/53.50/67.28/80.51	14.97	4.78
t04-v07-s01-r02 (50)	2.19/3.03/8.93/48.04	3.12/4.09/6.45/15.47	2.90/2.95/3.16/6.19	1.22	0.56
t09-v05-s02-r02 (414)	0.59/0.96/1.81/5.85	18.76/27.38/40.78/68.19	9.96/13.44/18.62/31.14	39.38	12.66
t03-v05-s02-r04 (200)	17.58/20.40/23.27/28.45	26.34/31.01/35.71/42.02	26.40/31.26/35.96/41.63	16.93	5.31
t09-v04-s02-r10 (174)	2.67/5.29/11.89/33.29	9.12/20.77/43.05/73.09	12.28/24.87/49.40/76.83	16.48	5.22
t02-v02-s01-r02 (50)	46.69/60.07/73.70/87.53	52.69/65.21/75.09/83.10	49.96/61.05/71.17/81.09	1.74	0.55
t01-v01-s07-r08 (10)	55.77/73.46/86.73/95.58	44.38/62.67/81.30/93.77	45.18/60.44/79.51/93.17	0.07	0.02
t09-v05-s03-r02 (396)	0.92/1.50/2.52/7.46	19.12/26.66/39.17/63.03	11.09/14.39/18.89/33.16	37.15	11.73
t09-v03-s01-r05 (225)	0.60/0.81/1.56/6.82	4.75/7.10/10.68/16.26	6.16/9.81/15.58/24.49	21.56	6.88
t09-v06-s02-r02 (170)	4.77/11.19/26.42/55.94	38.80/52.63/66.29/80.84	30.86/46.92/63.08/79.32	15.20	4.77
t02-v05-s01-r02 (50)	28.90/46.24/62.40/80.10	29.92/50.14/68.57/81.86	16.42/36.03/58.35/80.56	1.77	0.55
t01-v01-s08-r08 (5)	41.57/62.48/81.24/93.75	34.87/58.60/79.30/93.10	27.68/37.16/50.73/59.58	0.01	0.01
t04-v01-s00-r06 (300)	2.17/4.90/13.02/35.11	20.54/30.91/41.81/53.79	21.47/33.89/48.49/66.86	28.65	8.99
t04-v10-s00-r01 (300)	2.24/4.66/10.51/27.18	14.16/23.18/36.13/50.89	12.48/20.65/35.34/64.80	19.55	9.22
t09-v06-s01-r02 (175)	20.98/26.66/37.78/67.47	37.68/51.20/64.53/78.81	34.22/48.79/63.45/78.07	16.15	5.03
t02-v05-s02-r02 (75)	3.52/7.84/19.20/46.96	13.55/21.55/30.57/50.45	14.14/21.99/31.89/53.01	3.95	1.26
t04-v01-s01-r06 (300)	1.21/2.44/6.22/21.78	5.80/9.04/13.53/21.13	8.67/13.19/21.99/44.75	27.03	8.51
t09-v03-s00-r05 (255)	0.58/1.06/3.14/11.56	3.83/5.76/8.88/16.02	4.64/7.94/14.55/28.90	25.37	7.83
t09-v06-s00-r02 (180)	34.86/49.12/65.97/84.67	35.17/48.28/61.58/77.66	31.11/45.16/59.68/76.37	19.19	5.10
t02-v05-s04-r02 (107)	3.23/7.92/21.85/53.72	16.94/23.69/36.60/60.33	15.18/21.38/35.42/63.02	8.28	2.55

Table 12. **Per-sequence GLUEMAP reconstruction on WRIVA.** Each AUC cell reports pose AUC (%) as the tuple AUC@3° / @5° / @10° / @30°. Time is the visual-disambiguation two-view inference time for each sequence. Results cover all 34 evaluation sequences.




Figure 6. **Additional AerialMegaDepth COLMAP reconstructions (1/3).** Ground-truth and reconstructed camera layouts for Adler Planetarium and Belfry of Bruges.




Figure 7. **Additional AerialMegaDepth COLMAP reconstructions (2/3).** Ground-truth and reconstructed camera layouts for Brussels Town Hall and Florence Baptistery.



Figure 8. **Additional AerialMegaDepth COLMAP reconstructions (3/3).** Ground-truth and reconstructed camera layouts for the Louvre Museum and Ponte di Rialto. For Ponte di Rialto, DG++ and XDG each produce two reconstruction components, which are stacked vertically.




Figure 9. **Representative WRIVA GLUEMAP reconstructions (1/3).** On Sequence 00, AUC@10° increases from 2.63% without disambiguation to 38.38% with DG++ and 38.45% with XDG. On Sequence 01, it increases from 1.59% to 31.94% and 37.35%, respectively.




Figure 10. **Representative WRIVA GLUEMAP reconstructions (2/3).** On Sequence 00, AUC@10° increases from 12.28% without disambiguation to 42.27% with DG++ and 53.47% with XDG. On Sequence 02, it increases from 11.89% to 43.05% and 49.40%, respectively.



Figure 11. **Representative WRIVA GLUEMAP reconstructions (3/3).** On PTZ Single Zone, AUC@10° increases from 26.42% without disambiguation to 66.29% with DG++ and 63.08% with XDG. On Varying Altitudes Multi-Zone, it increases from 25.41% to 52.25% and 69.93%, respectively.