GSAlign: Geometric and Semantic Alignment Network for Aerial-Ground Person Re-Identification

Qiao Li¹, Jie Li², Yukang Zhang², Lei Tan³, Jing Chen¹, Jiayi Ji^{2,3}

¹Key Laboratory of Aerospace Information Security and Trusted Computing, Ministry of Education, School of Cyber Science and Engineering, Wuhan University

²Xiamen University

³National University of Singapore

liqiaoqiao233@whu.edu.cn, lei.tan@nus.edu.sg

Abstract

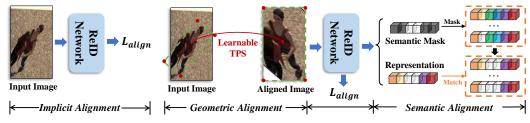
Aerial-Ground person re-identification (AG-ReID) is an emerging yet challenging task that aims to match pedestrian images captured from drastically different viewpoints, typically from unmanned aerial vehicles (UAVs) and ground-based surveillance cameras. The task poses significant challenges due to extreme viewpoint discrepancies, occlusions, and domain gaps between aerial and ground imagery. While prior works have made progress by learning cross-view representations, they remain limited in handling severe pose variations and spatial misalignment. To address these issues, we propose a Geometric and Semantic Alignment Network (GSAlign) tailored for AG-ReID. GSAlign introduces two key components to jointly tackle geometric distortion and semantic misalignment in aerial-ground matching: a Learnable Thin Plate Spline (LTPS) Module and a Dynamic Alignment Module (DAM). The LTPS module adaptively warps pedestrian features based on a set of learned keypoints, effectively compensating for geometric variations caused by extreme viewpoint changes. In parallel, the DAM estimates visibilityaware representation masks that highlight visible body regions at the semantic level, thereby alleviating the negative impact of occlusions and partial observations in cross-view correspondence. A comprehensive evaluation on CARGO with four matching protocols demonstrates the effectiveness of GSAlign, achieving significant improvements of +18.8% in mAP and +16.8% in Rank-1 accuracy over previous state-of-the-art methods on the aerial-ground setting. The code is available at: https://github.com/stone96123/GSAlign.

1 Introduction

Person re-identification (ReID), aiming to address the problem of matching people over a distributed set of nonoverlapping cameras, has attracted intensive attention in the last few years due to its wide applications in surveillance systems. Although traditional person re-identification (ReID) has achieved remarkable progress [1, 2, 3, 4, 5], there is growing interest in aerial-ground ReID (AG-ReID) settings, driven by the increasing deployment of unmanned aerial vehicles (UAVs) and low-altitude platforms in surveillance applications [6, 7, 8, 9, 10]. This, in practice, puts the Re-ID problem in an aerial-ground setting and requires the approaches to properly handle both the significant geometric view-variation and semantic misalignment.

^{*}Project Leader

[†]Corresponding Author: Lei Tan



(a) Current alignment strategy.

(b) Our proposed Geometric and Semantic Alignment Network (GSAlign).

Figure 1: **Illustration of the motivation of GSAlign.** (a) Prior methods rely solely on implicit alignment, which is insufficient to fully address spatial and semantic distortions. (b) In contrast, our GSAlign performs explicit alignment at both the geometric and semantic levels via LTPS and visibility-aware semantic masks, respectively. This design equips GSAlign with a stronger capability for robust aerial-ground matching.

Despite its potential, AG-ReID remains a highly underexplored and technically challenging problem. Unlike conventional ReID tasks [11, 12, 13, 14, 15], AG-ReID suffers from extreme viewpoint disparities, where aerial images exhibit severe top-down perspectives while ground views contain mostly frontal or profile views. These differences cause significant geometric distortions and drastic appearance changes, which conventional ReID models that are primarily trained on ground-view data are ill-equipped to handle. In addition, AG-ReID often involves frequent occlusions and unbalanced visibility across views, further complicating cross-view matching.

To mitigate these challenges, several recent studies have attempted to adapt existing ReID techniques to the aerial-ground scenario [16, 8]. These methods seek to learn modality-specific representations or map features into a shared embedding space using metric learning or adversarial objectives. Others employ cross-view decomposition strategies or design view-invariant constraints to bridge the modality gap. However, these methods tend to focus on global alignment and overlook two critical issues: (1) the severe geometric distortion induced by cross-view perspectives, and (2) semantic misalignment caused by partial occlusions and varying visible body regions. These limitations lead to suboptimal performance, especially under large viewpoint gaps.

In this paper, as shown in Figure 1 (b), we propose a Geometric and Semantic Alignment Network (GSAlign), which is specifically designed to tackle the core challenges of Aerial-Ground Re-Identification (AG-ReID). The key insight of our approach is to explicitly model both geometric deformation and visibility-aware semantics within a unified ViT framework. To this end, GSAlign introduces two complementary modules: First, the Learnable Thin Plate Spline (LTPS) Transformation Module. LTPS adaptively warps pedestrian representations based on a set of learned keypoints. Unlike hand-crafted alignment strategies, our LTPS module is fully differentiable and end-to-end trainable, enabling the network to learn viewpoint-conditioned transformations that effectively and dynamically compensate for severe spatial distortions caused by extreme cross-view differences. To avoid potential errors caused by one-shot correction, LTPS is progressively integrated into the hierarchical layers of a Vision Transformer (ViT). This design enables the network to iteratively refine geometric alignment across layers, allowing fine-grained, stage-wise correction throughout the feature propagation process. Second, the Dynamic Alignment Module (DAM). DAM enables each input image to predict its own visibility-aware representation masks at the semantic level, highlighting visible body regions while suppressing noisy or occluded areas. The predicted masks are then applied to the corresponding gallery features to filter out noisy signals from invisible regions, thereby enhancing cross-view feature alignment. By dynamically adapting to the visibility of different body parts, DAM guides the network to focus on semantically consistent and identity-discriminative cues across aerial and ground views. This allows the model to focus on semantically consistent cues across views, thereby improving robustness under occlusions and pose variations.

To sum up, the main contributions of this paper are as follows:

• We propose GSAlign, a novel framework for aerial-ground person re-identification that jointly addresses geometric deformation and semantic misalignment within a unified architecture. GSAlign is specifically designed to handle the extreme cross-view variations and visibility inconsistencies inherent in UAV-to-ground matching scenarios.

- We introduce a Learnable Thin Plate Spline (LTPS) Module and a Dynamic Alignment Module (DAM). LTPS performs keypoint-guided feature warping to compensate for severe spatial distortions, while DAM enhances semantic alignment by estimating visibility-semantic representation masks to highlight visible body regions and suppress noisy or occluded areas.
- Extensive experiments on the challenging CARGO dataset validate the effectiveness of GSAlign, which achieves state-of-the-art performance with absolute gains of +18.8% in mAP and +16.8% in Rank-1 accuracy on the aerial-ground setting.

2 Related Work

Building upon well-established ground surveillance infrastructure, advanced person re-identification algorithms have made significant progress across a range of scenarios, including general scenarios [17, 18, 1, 4], occluded scenarios [19, 12, 20, 21], cross-modal scenarios [22, 23, 24, 25], multi-spectral scenarios [26, 27, 28], and cross-spectral scenarios [29, 30, 31, 32, 33].

However, with the proliferation of unmanned aerial vehicles (UAVs) and low-altitude platforms augmenting traditional surveillance systems, person ReID research is increasingly extending to aerial-view scenarios [34]. Initial efforts in this direction focused on aerial-aerial retrieval tasks using datasets such as UAV-Human [35] and PRAI-1581 [7]. More recently, attention has shifted towards connecting the aerial views with traditional ground systems. Aerial-Ground person ReID (AG-ReID), where the query and gallery images originate from fundamentally different viewpoints, typically aerial and ground-level perspectives. This setup introduces severe spatial and semantic discrepancies, often leading to substantial degradation in identity-preserving cues such as pose, silhouette, and clothing texture [36]. Several approaches have been proposed to mitigate these challenges. Nguyen et al. [6] proposed a dual-stream structure guided by attributes to enhance semantic disentanglement, later extended to a three-stream model with modality-aware supervision [16]. Zhang et al. [8] developed the View-Decoupled Transformer (VDT), which explicitly models and separates viewpointspecific and viewpoint-invariant features. Despite these advances, existing AG-ReID methods still face difficulties in handling extreme spatial distortions and in leveraging visible semantic regions effectively. To overcome these limitations, we propose **GSAlign**, a novel framework that performs instance-adaptive geometric transformation and visibility-aware semantic alignment, promoting robust and spatially consistent feature learning across aerial and ground domains.

3 Methodology

3.1 Overview

The overall framework of the **Geometric and Semantic Alignment Network (GSAlign)**, illustrated in Fig. 2, is built upon a ViT-Base backbone and designed to address feature misalignment under complex multi-view conditions. Inspired by the success of the View-Decoupled Transformer (VDT) [37], we adopt the VDT architecture as the foundation of GSAlign, which effectively captures both global semantics and fine-grained structural cues. Following its design, we introduce an additional view token to model viewpoint-specific information. The token sequence during training is defined as:

$$[\mathbf{X}_{cls}, \mathbf{X}_{view}, \mathbf{X}_{img}] = VDT([\mathbf{X}_{cls}, \mathbf{X}_{view}, tokenize(\mathbf{X}_{img})]), \tag{1}$$

where X_{cls} , X_{view} , and X_{img} represent the class token, view token, and image tokens, respectively.

To improve geometric robustness, we integrate a **Learnable Thin Plate Spline module (LTPS)** into transformer layers. This module predicts rotation-aware control point offsets and applies a non-rigid spatial transformation to warp local patch features. By doing so, LTPS mitigates pose-induced geometric distortions and enhances structural alignment across views. In addition, we introduce a **Dynamic Alignment Module (DAM)** during training, which generates channel-wise semantic masks guided by class-specific prototype features. These masks dynamically highlight identity-relevant subspaces, allowing the model to focus on discriminative patterns while suppressing background noise and occlusion artifacts. The generated masks serve a specific purpose: aligning the features toward consistent subspaces. Overall, the GSAlign integrates the LTPS and DAM, jointly optimizing for classification accuracy, feature consistency, and robustness to view and shape variations. This synergy enables GSAlign to show a strong re-identification ability in severe view variations.

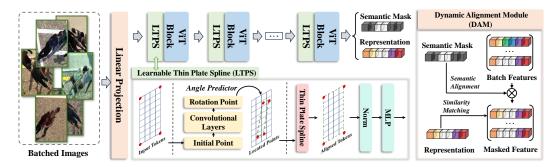


Figure 2: **Overview of the GSAlign architecture.** Given aerial or ground-view inputs, GSAlign first applies an initial geometric transformation via a Learnable Thin Plate Spline (LTPS) module, followed by progressive alignment through LTPS blocks inserted before each ViT layer. In parallel, a Dynamic Alignment Module (DAM) generates a visibility-aware semantic mask according to the input image, which is then applied to the representations of other images in the batch to suppress irrelevant or occluded features.

3.2 Learnable Thin Plate Spline

To address the spatial misalignment problem in person re-identification, we design a Learnable Thin Plate Spline (LTPS) deformation module that explicitly models geometric variations arising from changes in human pose and viewpoint. By dynamically predicting control point displacements and rotation angles, the module performs nonlinear transformations on feature maps, thereby enhancing their ability to capture non-rigid spatial deformations. As illustrated in Fig. 2, the LTPS modules are integrated into each transformer block to build a hierarchical deformation-aware representation. The LTPS modules at shallower layers focus on capturing local deformation details, while those at deeper layers are responsible for modeling global pose variations.

The core idea of the LTPS module is to perform non-rigid deformation based on control points interpolation. We initialize a set of 2D source contro points $\mathbf{P}_s \in \mathbb{R}^{K \times 2}$ as a regular grid uniformly distributed over the normalized coordinate space $[-1,1] \times [-1,1]$ and treat them as learnable parameters optimized during training. In parallel, we define a set of target control points $\mathbf{P}_t \in \mathbb{R}^{K \times 2}$ which are fixed and share the same regular grid positions as the initial \mathbf{P}_s , representing the canonical target shape. During training, a rotation angle is predicted from the input features and applied to \mathbf{P}_s , and a smooth deformation mapping is constructed by enforcing an interpolation constraint from the rotated \mathbf{P}_s to the fixed \mathbf{P}_t , guiding the feature map to undergo spatial transformation.

In traditional thin plate spline (TPS) transformations, the source control points P_s are fixed, and spatial rearrangement is achieved only by learning the target control points P_t , which limits the ability to handle global rotations and complex deformations. To address this, we design a rotation prediction module that predicts a rotation angle from the global orientation of the input feature map, and applies this rotation to the source control points for improved modeling of spatial deformations. The process is defined as:

$$\mathbf{P}_{s}^{(\text{rot})} = \mathbf{P} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}^{\top}, \quad \theta = f_{\theta}(\mathbf{F}). \tag{2}$$

Here, **F** represents the patch-level feature map of each layer, and $f_{\theta}(\cdot)$ is the rotation prediction network that outputs the angle $\theta \in \left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$.

To obtain the transformation function $T(\cdot)$, the TPS requires the source control points P_s and the target control points P_t to form the following interpolation constraint:

$$T(P_{s,i}) = P_{t,i}, \quad \forall i = 1, 2, 3, ..., K.$$
 (3)

This constraint requires that the transformation function precisely maps the source points \mathbf{P}_s to the target points \mathbf{P}_t at each control point. Based on this, for the rotated source points \mathbf{P}_s^{rot} and the target control points \mathbf{P}_t , the TPS transformation function $\mathbf{T}(\cdot)$ is defined as follows:

$$\mathbf{T}(\mathbf{p}_x) = \mathbf{A} \cdot \mathbf{p}_x + \sum_{i=1}^K \mathbf{w}_i \cdot U(\|\mathbf{x} - \mathbf{P}_{s,i}^{(\text{rot})}\|), \tag{4}$$

where \mathbf{p}_x can be any point from the feature map. $\mathbf{A} \in \mathbb{R}^{2 \times 2}$ is an affine matrix. \mathbf{w}_i is the deformation weight for the *i*-th TPS kernel function, $U_i(\cdot)$, which is defined as:

$$U(r) = r^2 \log r^2, \quad r = \|\mathbf{P}_x - \mathbf{P}_{s,i}\|_2.$$
 (5)

Here, r is the Euclidean distance between the feature point p_x and the i-th source control point.

The patch feature after deformation by $T(\cdot)$ is denoted as \mathbf{F}_{ltps} . We apply residual fusion with the original patch feature F:

$$\mathbf{F}_{\text{final}} = \mathbf{F} + \eta \cdot \mathbf{F}_{\text{ltps}}. \tag{6}$$

Here, η is a tunable fusion factors. This fusion strategy allows the model to retain original semantic features while explicitly incorporating dynamic spatial structure awareness. As a result, it improves the model's ability to recognize people under occlusion, rotation, and deformation in complex scenes.

3.3 Dynamic Alignment Module

In person re-identification, training images often contain occlusions, background noise, or identity-irrelevant regions. Directly encoding features from the entire image may lead the model to learn redundant or non-discriminative representations. To address this issue, we explore whether the visibility patterns of a given input image can be used to dynamically modulate the feature representations of other samples. Inspired by DPM [38], we propose the **Dynamic Alignment Module (DAM)**, which leverages input-guided masks to improve semantic alignment. Specifically, DAM treats the mean feature representation of images with the same identity as its ideal prototype, assuming that noise or occlusion in different samples manifests as partial loss of this prototype's semantic structure. Based on this formulation, DAM introduces a dynamic channel-wise mask generator based on the input and modulates prototypes by referencing the input-aware mask. This allows the model to suppress irrelevant or noisy dimensions and emphasize input-relevant subspaces in a content-aware manner.

Specifically, during training, for each batch, we firstly construct prototype features for all classes as:

$$\mathbf{p}_c = \frac{1}{N_c} \sum_{i=1}^{N_c} \mathbf{f}_i. \tag{7}$$

Here, N_c is the number of samples belonging to class c within the current mini-batch, and \mathbf{f}_i represents the feature embedding of the i-th sample. All prototype features are ℓ_2 -normalized both before and after the update to ensure consistent alignment with the sample features.

To enable dynamic selection of relevant prototype subspace, we design a channel-wise mask generator. This generator guides the model to identify and emphasize the prototype subspace most relevant for discriminating the current input image. Unlike spatial domain masks operating on pixel regions, we focus on channel-wise selection. The goal is to generate a sparse, discriminative channel-wise mask for each image, highlighting the discriminative dimensions during prototype matching.

The mask generator creates a channel-wise mask based on the current sample feature \mathbf{f}_i with a two-layer multilayer perceptron (MLP), followed by a Sigmoid activation function to ensure mask values are between 0 and 1:

$$\mathbf{m}_{i} = \operatorname{Sigmoid}\left(W_{2} \cdot \operatorname{ReLU}\left(W_{1} \cdot \mathbf{f}_{i} + \mathbf{b}_{1}\right) + \mathbf{b}_{2}\right),\tag{8}$$

where $W_1 \in \mathbb{R}^{D \times \frac{D}{2}}$ and $W_2 \in \mathbb{R}^{\frac{D}{2} \times D}$ are learnable weight matrices, $\mathbf{b_1} \in \mathbb{R}^{\frac{D}{2}}$ and $\mathbf{b_2} \in \mathbb{R}^D$ denote the bias terms, and \mathbf{f}_i is the input feature vector.

The generated sample-specific mask is then used to weight the prototype features, focusing on the parts of the features that are relevant to the category. The mask is applied by performing element-wise multiplication with the prototype features as follow:

$$\mathbf{p}_c^{masked} = \mathbf{p}_c \odot \mathbf{m}_c, \tag{9}$$

where \mathbf{p}_c denotes the prototype feature of class c, and \mathbf{m}_c is a channel-wise mask generated by the Channel-wise Mask Generator based on the feature representation of each sample and its corresponding label.

This mask is employed exclusively during the training phase to reweight the prototype \mathbf{p}_c along the channel dimension. By emphasizing informative feature channels and suppressing less relevant

ones, this mechanism encourages the model to focus on more discriminative aspects of the prototype representation. It is important to emphasize that this masked prototype weighting strategy is only applied during training to facilitate effective prototype learning. During inference, the model directly compares the extracted features with the unmasked prototypes, without relying on any class labels or mask generation.

Overall, the dynamic weighting mechanism enables the model to adaptively adjust prototype representations according to the feature distribution of each class. This enhances the discriminative capacity of the learned representations by selectively highlighting important channels and filtering out irrelevant dimensions. In addition, this module is utilized only during the training phase, which will not add the inference computation cost.

3.4 Loss Function and Optimization

Our model is trained using a composite loss function designed to optimize identity discrimination, geometric stability, and feature-level alignment. In addition to the standard ID classification loss \mathcal{L}_{id} and the Triplet loss \mathcal{L}_{tri} commonly employed in ReID tasks, we introduce a deformation loss \mathcal{L}_{deform} and the mask loss \mathcal{L}_{mask} .

Deformation Loss. To enhance the stability of the LTPS module's deformation modeling, we introduce a regularization term as the deformation loss \mathcal{L}_{deform} . This loss discourages the rotation prediction submodule from outputting excessively large rotation angles, thereby promoting geometrically plausible transformations. While the angle prediction mechanism intrinsically limits angles via a tanh activation and scaling, this loss provides an additional constraint. It is defined as the average absolute predicted rotation angle across all L encoder layers where an LTPS module is integrated:

$$\mathcal{L}_{\text{deform}} = \frac{1}{L} \sum_{l=1}^{L} |\theta_l|, \tag{10}$$

where θ_l represents the predicted rotation angle at the l-th layer. By penalizing large rotation magnitudes, this loss helps maintain geometric consistency during the feature transformation process.

Mask Loss. To enhance the effectiveness of the Dynamic Alignment Module, we introduce a mask alignment loss to ensure consistency between the dynamically masked prototype and the input sample feature. we formulate a composite mask loss \mathcal{L}_{mask} . This loss comprises two components: an alignment term \mathcal{L}_{align} and an entropy-based regularization term $\mathcal{L}_{entropy}$:

$$\mathcal{L}_{\text{mask}} = \mathcal{L}_{\text{align}} + \lambda \mathcal{L}_{\text{entropy}}.$$
 (11)

 $\mathcal{L}_{\text{align}}$ aims to ensure that an input sample's feature representation, when masked, aligns closely with its corresponding class prototype, also masked by the same sample-specific mask. This encourages consistency in the feature subspace highlighted by the dynamic mask, even in the presence of occlusions or pose variations which might otherwise lead to disparate semantic activations. Let f_i be the feature of i-th sample with class of c, \mathbf{m}_i its generated channel-wise mask, and \mathbf{p}_c the prototype for the class c. The alignment loss is computed as:

$$\mathcal{L}_{\text{align}} = \frac{1}{N} \sum_{i=1}^{N} \|\text{Norm}(\mathbf{m}_i \odot \mathbf{p}_c) - \text{Norm}(\mathbf{m}_i \odot \mathbf{f}_i)\|_2^2,$$
 (12)

where N is the batch size, and Norm means the ℓ_2 normalization. This loss encourages intra-class compactness within the dynamically selected channel subspace. By applying the mask \mathbf{m}_i to both the sample feature and its prototype, the model learns to emphasize discriminative channels corresponding to visible regions while down-weighting channels associated with occlusions or background clutter.

 $\mathcal{L}_{\text{entropy}}$ is designed to promote channel selectivity in the generated masks. During training, the mask tends to degenerate into a fully activated state, *i.e.*, $\mathbf{m}_i \to \mathbf{1}$. This loss drives the mask vector towards a binary distribution by maximizing the information entropy, thus retaining the most relevant channels for each sample. It is defined as the negative sum of binary entropies for each mask component:

$$\mathcal{L}_{\text{entropy}} = -\frac{\lambda}{ND} \sum_{i=1}^{N} \sum_{d=1}^{D} \left[\mathbf{m}_{i}^{(d)} \log \mathbf{m}_{i}^{(d)} + (1 - \mathbf{m}_{i}^{(d)}) \log(1 - \mathbf{m}_{i}^{(d)}) \right]. \tag{13}$$

 \mathcal{L}_{align} and $\mathcal{L}_{entropy}$ work as complementary constraints. \mathcal{L}_{align} guides the mask to select channels that maximize prototype similarity, while $\mathcal{L}_{entropy}$ encourages the mask to make more definitive selections. A balancing hyperparameter λ is introduced to regulate the trade-off between the two objectives.

Overall Loss. The model is trained end-to-end by minimizing a comprehensive loss function that aggregates the aforementioned components:

$$\mathcal{L}_{\text{total}} = (\mathcal{L}_{\text{id}} + \mathcal{L}_{\text{tri}}) + \alpha \mathcal{L}_{\text{deform}} + \beta \mathcal{L}_{\text{mask}}, \tag{14}$$

where the α and β is the hyper-parameters to adjust the weight of \mathcal{L}_{deform} and \mathcal{L}_{mask} respectively. This joint optimization framework encourages the learning of discriminative and robust representations that are resilient to spatial misalignments and feature irrelevant variations.

4 Experiments

4.1 Implementation

Datasets and Evaluation Protocols. We conduct experiments on the CARGO dataset [8], a large-scale benchmark specifically designed for aerial-ground person re-identification (AG-ReID). CARGO contains 108,563 images of 5,000 synthetic identities captured by 13 cameras (5 aerial and 8 ground), with diverse conditions including extreme viewpoint changes, resolution variations, illumination shifts, and occlusions. The dataset is constructed in a simulated urban environment using MakeHuman for identity modeling and Unity3D for camera deployment.

Following the standard protocol, we use 51,451 images with 2500 identities for training and 51,024 images with the remaining 2,500 identities for testing. The testing phase contains four evaluation protocols: ALL, ground to ground($G \leftrightarrow G$), aerial to aerial ($A \leftrightarrow A$), and aerial to ground ($A \leftrightarrow G$), where each protocol focuses on a specific matching scenario. To further verify the generalization ability of GSAlign in real-world scenarios, we also conduct comparative experiments on the AG-ReID [6] and AG-ReID v2 [16] datasets, which contain real aerial and ground imagery captured in outdoor environments. We use Cumulative Matching Characteristics (CMC) at Rank-1 accuracy, mean Average Precision (mAP), and mean Inverse Negative Penalty (mINP) [39] as evaluation metrics.

Implementation Details. All implementation settings strictly follow the protocol of VDT [8] for fair comparison. We adopt ViT-Base as the backbone, initialized with ImageNet-21K pretraining. Input images are resized to 256×128 , and horizontal flipping is applied during training for data augmentation. The model is optimized using AdamW with weight decay of 0.05 and a cosine learning rate schedule. The initial learning rate is set to 3.5×10^{-4} with linear warm-up for the first 20 epochs. λ is set to 0.1 during the training. We train for 120 epochs in total, using a batch size of 64.

4.2 Comparison with State-of-the-Arts

In this section, we compare our proposed **GSAlign** with a range of state-of-the-art person ReID methods, including both conventional CNN-based models (e.g., PCB [17], BoT [41], MGN [42]) and Transformer-based frameworks (e.g., ViT [37], VDT [8]). As shown in Table 1, GSAlign achieves the best overall performance across all four protocols. Specifically, under the most comprehensive ALL protocol, GSAlign achieves 65.06% Rank-1, 57.95% mAP, and 44.97% mINP, outperforming the strong baseline VDT [8] by +0.96%, +2.75%, and +3.84%, respectively. On the most challenging $A \leftrightarrow G$ cross-view setting, our method surpasses all previous methods by a large margin, reaching 64.89% Rank-1, 61.55% mAP, and 52.81% mINP, significantly ahead of the second-best VDT by +16.77%, +18.79%, and +22.86%, respectively. Although GSAlign is not specifically designed for single-view matching protocols, it does not introduce any negative effects in such settings. On both $A \leftrightarrow A$ and $G \leftrightarrow G$, GSAlign still achieves competitive performance. We attribute this to the fact that, in single-view scenarios, the features are already well aligned due to the limited viewpoint variation, thus requiring minimal additional geometric or semantic correction. In conclusion, these results clearly demonstrate that our proposed GSAlign not only achieves strong global retrieval accuracy but also exhibits exceptional robustness in the presence of extreme viewpoint changes, validating the effectiveness of our geometric and semantic alignment design.

Table 1: **Performance comparison of the mainstream methods under four settings of the CARGO dataset**. "ALL" denotes the overall retrieval performance of each method. " $G \leftrightarrow G$ ", " $A \leftrightarrow A$ ", and " $A \leftrightarrow G$ " represent the performance of each model in several specific retrieval patterns. Rank1, mAP, and mINP are reported (%). The best performance is shown in **bold**.

Method	Prot	ocol 1: A	ALL	Proto	ocol 2: G	⇔G	Proto	ocol 3: A	.↔A	Proto	ocol 4: A	↔G
Wichiod	Rank1	mAP	mINP									
SBS [40]	50.32	43.09	29.76	72.31	62.99	48.24	67.50	49.73	29.32	31.25	29.00	18.71
PCB [17]	51.00	44.50	32.20	74.10	67.60	55.10	55.00	44.60	27.00	34.40	30.40	20.10
BoT [41]	54.81	46.49	32.40	77.68	66.47	51.34	65.00	49.79	29.82	36.25	32.56	21.46
MGN [42]	54.81	49.08	36.52	83.93	71.05	55.20	65.00	52.96	36.78	31.87	33.47	24.64
VV [43, 44]	45.83	38.84	39.57	72.31	62.99	48.24	67.50	49.73	29.32	31.25	29.00	18.71
AGW [39]	60.26	53.44	40.22	81.25	71.66	58.09	67.50	56.48	40.40	43.57	40.90	29.39
BAU [45]	45.20	38.40	-	61.60	51.20	-	50.00	42.60	-	40.40	36.70	-
PAT [46]	37.90	15.30	-	52.70	24.20	-	50.00	23.10	-	35.10	15.50	-
DTST [47]	64.42	55.73	41.92	78.57	72.40	62.10	80.00	63.31	44.67	50.53	43.49	29.46
ViT [37]	61.54	53.54	39.62	82.14	71.34	57.55	80.00	64.47	47.07	43.13	40.11	28.20
VDT [8]	64.10	55.20	41.13	82.14	71.59	58.39	82.50	66.83	50.22	48.12	42.76	29.95
GSAlign	65.06	57.95	44.97	83.04	73.86	62.73	80.00	65.55	49.81	64.89	61.55	52.81

Furthermore, we evaluate the proposed GSAlign on the real-world AG-ReID [6] and AG-ReID v2 [16] datasets, and compare it against a broad range of state-ofthe-art ReID approaches, as summarized in Table 2 and Table 3. On the AG-ReID dataset, GSAlign consistently achieves the best overall performance across both crossview protocols. Specifically, under the $A \leftrightarrow G$ setting, it attains 86.74% Rank-1, 84.00% mAP, and 73.62% mINP, outperforming the strong baseline VDT by large margins of +3.83%, +9.56%, and +22.56%, respectively. Under the inverse $G \leftrightarrow A$ protocol, GSAlign maintains superior results with 87.94% Rank-1, 87.17% mAP, and 78.21% mINP, exceeding VDT by +1.35%,

Table 2: Performance comparison of the mainstream methods under two settings of the AG-ReID dataset. Results are reported under four matching protocols. The best values per column are shown in **bold**.

Setting	Proto	col 1: A	A↔G	Proto	col 2: C	i ↔A
Setting	Rank1	mAP	mINP	Rank1	mAP	mINP
SBS [40]	73.54	59.77	-	73.70	62.27	-
BoT [41]	70.01	55.47	-	71.20	58.83	-
OsNet [48]	72.59	58.32	-	74.22	60.99	-
VV [43, 44]	77.22	67.23	41.43	79.73	69.83	42.37
Explain	81.28	72.38	-	82.64	73.35	-
ViT [37]	81.47	72.61	-	82.85	73.39	-
VDT [8]	82.91	74.44	51.06	86.59	78.57	52.87
GSAlign	86.74	84.00	73.62	87.94	87.17	78.21

+8.60%, and +25.34%, respectively. These substantial improvements demonstrate the advantage of our geometric-semantic alignment mechanism in handling severe viewpoint discrepancies between aerial and ground cameras. For the AG-ReID v2 dataset, which introduces more realistic and diverse environmental conditions, GSAlign again achieves the best performance in most matching protocols. Under the comprehensive *ALL* protocol, GSAlign yields 91.47% Rank-1 and 89.78% mAP, surpassing all competing methods by a significant margin. It also achieves the highest mAP under the $G \leftrightarrow G$ and $A \leftrightarrow G$ settings (87.62% and 88.62%, respectively), while maintaining highly competitive Rank-1 scores across the remaining protocols. Compared with the previous best AG-ReIDv2 [16], GSAlign improves the average Rank-1 accuracy by +2.7% and the mAP by +6.5%, highlighting its robust cross-domain generalization ability.

Overall, these results consistently verify that GSAlign not only achieves superior retrieval accuracy but also exhibits remarkable robustness across diverse aerial-ground configurations. The consistent gains over both synthetic (CARGO) and real-world (AG-ReID and AG-ReID v2) datasets validate the effectiveness and strong generalization of our geometric and semantic alignment design.

4.3 Ablation Study

To assess the individual contributions of each component in the proposed GSAlign framework, we conduct a comprehensive ablation study on the four testing protocols of the CARGO dataset. As shown in Table 4, we start with a strong transformer-based baseline [8] and progressively integrate the Learnable Thin Plate Spline (LTPS) transformation module and the Dynamic Alignment Module (DAM). Compared to the baseline, integrating LTPS brings consistent performance improvements

Table 3: Performance comparison of the mainstream methods under four settings of the AG-ReID v2 dataset. Results are reported under four matching protocols. The best values per column are shown in **bold**.

Setting	Protocol	1: A↔G	Protocol	2: G↔A	Protocol	3: A↔W	Protocol 4: W↔A		
Setting	Rank1	mAP	Rank1	mAP	Rank1	mAP	Rank1	mAP	
BoT [41]	85.40	77.03	84.65	75.90	89.77	80.48	84.65	76.90	
Explain [6]	87.70	79.00	87.35	78.24	93.67	83.14	87.73	79.08	
VDT [8]	86.46	79.13	86.14	78.12	90.00	82.21	85.26	78.52	
AG-ReIDv2 [16]	88.77	80.72	87.86	78.51	93.62	84.85	88.61	80.11	
SeCap [49]	88.12	80.84	88.24	79.99	91.44	84.01	87.56	80.15	
GSAlign	91.47	89.78	88.29	87.62	93.30	91.84	88.12	88.62	

Table 4: **Ablation study of the different components in GSAlign on the CARGO dataset.** "LTPS" denotes the Learnable Thin Plate Spline transformation module. "DAM" refers to the Dynamic Alignment Module. The best performance per column is shown in **bold**.

Setting	Protocol 1: ALL			Proto	Protocol 2: G↔G			Protocol 3: A↔A			Protocol 4: A↔G		
Setting	Rank1	mAP	mINP	Rank1	mAP	mINP	Rank1	mAP	mINP	Rank1	mAP	mINP	
Baseline	64.10	55.20	41.13	82.14	71.59	58.39	82.50	66.83	50.22	48.12	42.76	29.95	
Baseline + LTPS	64.42	55.95	41.92	80.36	71.87	59.55	82.50	65.26	47.15	64.89	61.08	50.54	
Baseline + LTPS + DAM	65.06	57.95	44.97	83.04	73.86	62.73	80.00	65.55	49.81	64.89	61.55	52.81	

across most settings. Specifically, on the most challenging $A \leftrightarrow G$ protocol, LTPS improves Rank-1 from 48.12% to 64.89%, mAP from 42.76% to 61.08%, and mINP from 29.95% to 50.54%. These results confirm the effectiveness of progressive geometric alignment in compensating for severe cross-view deformations. Interestingly, the performance gain on $A \leftrightarrow A$ and $G \leftrightarrow G$ is more modest, indicating that LTPS is particularly beneficial when strong viewpoint discrepancies exist. Adding the DAM module on top of LTPS yields further improvements, especially in terms of mAP and mINP. Under the full matching scenario (Protocol 1: ALL), the full model achieves the best performance with 65.06% Rank-1, 57.95% mAP, and 44.97% mINP. Notably, the largest relative gain comes from the $A \leftrightarrow G$ setting, where DAM increases mINP from 50.54% to 52.81%. This demonstrates that DAM effectively suppresses noisy and occluded regions by leveraging visibility-aware semantic masking, thus enhancing cross-view alignment. Overall, the ablation study verifies the complementary effects of LTPS and DAM. While LTPS resolves geometric distortions in a layer-wise manner, DAM provides semantic-level refinement by selectively filtering out unreliable features. Their combination allows GSAlign to maintain robust performance across diverse matching protocols, especially under severe viewpoint and occlusion conditions.

4.4 Discussions³

Number of control points in LTPS. To examine the sensitivity of GSAlign to the number of control points in the Learnable Thin Plate Spline (LTPS) module, we test five settings with 4, 9, 16, 25, 36 keypoints on the CARGO dataset under four protocols. As shown in Table 5, all variants benefit from LTPS-based geometric alignment, though performance varies with point density. The 4-point LTPS achieves the best overall results as Rank-1 (64.89%), mAP (61.55%), and mINP (52.81%) under the $A \leftrightarrow G$ protocol and remains strong across all scenarios. Increasing points (e.g., 25 or 36) slightly improves Rank-1 but reduces mAP and mINP, likely due to over-flexibility and local distortion. These findings indicate that a lightweight 4-point LTPS provides an optimal trade-off between alignment accuracy and stability.

Different locations for LTPS. We further study the effect of inserting LTPS at different depths of the ViT backbone. As shown in Table 6, applying LTPS to all transformer layers yields the best results (64.89% Rank-1, 61.55% mAP, 52.81% mINP under Protocol 4: $A \leftrightarrow G$) and consistent gains across scenarios. Among partial insertions, placing LTPS in the last 4 layers performs comparably to the full setting, implying that geometric distortions persist even in higher-level features and benefit from late-stage correction. In contrast, early-layer insertion offers limited improvement, especially in fine-grained metrics like mINP, likely because shallow features are less discriminative and harder to

³A visualization results of LTPS (Sec. A.1) is provided in the supplementary material.

Table 5: **Performance of GSAlign under different numbers of control points in LTPS.** Results are reported under four matching protocols. The best values per column are shown in **bold**.

Setting	Protocol 1: ALL			Protocol 2: G↔G			Proto	ocol 3: A	. \leftrightarrow A	Protocol 4: A↔G		
betting	Rank1	mAP	mINP	Rank1	mAP	mINP	Rank1	mAP	mINP	Rank1	mAP	mINP
Number of control points in LTPS												
4	65.06	57.95	44.97	83.04	73.86	62.73	80.00	65.55	49.81	64.89	61.55	52.81
9	64.10	57.35	44.76	82.14	75.16	64.91	80.00	64.51	47.23	62.77	59.25	50.19
16	63.78	56.48	42.97	80.36	73.78	63.49	77.50	63.48	47.28	61.70	58.06	48.19
25	65.71	57.62	44.15	82.14	74.06	62.93	82.50	65.95	48.20	63.83	60.23	50.53
36	65.71	57.55	44.34	81.25	72.71	61.11	82.50	65.88	48.47	63.83	60.00	51.12

Table 6: **Comparison between different locations for LTPS.** Results are reported under four matching protocols. The best values per column are shown in **bold**.

Setting	Protocol 1: ALL			Protocol 2: G↔G			Protocol 3: A↔A			Protocol 4: A↔G		
Setting	Rank1	mAP	mINP									
Different locatio	ns for LT	TPS		•						•		
First layer	64.10	55.92	42.44	83.04	72.86	60.58	80.00	65.98	50.45	58.51	56.92	47.62
First 4 layers	64.10	56.46	43.50	81.25	74.49	64.70	80.00	64.45	47.11	58.51	56.21	46.66
Middle 4 layers	64.74	57.09	44.08	82.14	74.93	64.77	77.50	64.28	47.30	58.51	58.30	50.18
Last 4 layers	65.06	57.39	44.05	83.04	74.42	62.86	77.50	65.21	49.80	64.89	59.87	50.95
All layers	65.06	57.95	44.97	83.04	73.86	62.73	80.00	65.55	49.81	64.89	61.55	52.81

Table 7: **Comparison between different variants of DAM.** Results are reported under four matching protocols. The best values per column are shown in **bold**.

Setting	Protocol 1: ALL			Proto	Protocol 2: G↔G			Protocol 3: A↔A			Protocol 4: A↔G		
Setting	Rank1	mAP	mINP	Rank1	mAP	mINP	Rank1	mAP	mINP	Rank1	mAP	mINP	
Different variants of DAM													
Inner-Batch	65.06	57.95	44.97	83.04	73.86	62.73	80.00	65.55	49.81	64.89	61.55	52.81	
Memory Bank	65.38	57.34	44.09	83.04	73.72	62.05	80.00	62.70	43.88	63.83	61.06	52.52	
Classification Matrix	63.14	55.64	42.07	81.25	72.04	59.53	75.00	63.79	48.06	57.45	56.55	47.33	

align precisely. These results highlight the importance of progressive alignment from shallow to deep layers for handling complex viewpoint distortions in AG-ReID.

Different variants of DAM. We compare three variants of the Dynamic Alignment Module (DAM): Inner-Batch, Memory Bank, and Classification Matrix. As shown in Table 7, the Inner-Batch variant achieves the best overall performance. It dynamically generates visibility-aware masks within each mini-batch, enabling query-driven refinement of peer features. The Memory Bank variant stores identity prototypes updated by momentum, providing semantic context but suffering from stale or mismatched memory. The Classification Matrix variant reuses classifier weights as class centers [38], but these prototypes are biased toward classification objectives. Overall, Inner-Batch masking proves most effective for dynamic alignment in aerial-ground person ReID.

5 Conclusion

In this paper, we propose GSAlign, a novel framework for aerial-ground person re-identification that jointly addresses geometric distortion and semantic misalignment. By introducing a Learnable Thin Plate Spline (LTPS) module and a Dynamic Alignment Module (DAM), GSAlign performs progressive geometric correction and visibility-aware semantic filtering in a unified manner. Extensive experiments on the challenging CARGO and AG-ReID datasets demonstrate that GSAlign achieves state-of-the-art performance across multiple protocols, highlighting its effectiveness in handling severe viewpoint changes and partial occlusions.

Acknowledgements. This research was supported in part by the National Key R&D Program of China under grant No. 2022YFB3103300.

References

- [1] Shuting He, Hao Luo, Pichao Wang, Fan Wang, Hao Li, and Wei Jiang. Transreid: Transformer-based object re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 15013–15022, 2021.
- [2] Wen Li, Cheng Zou, Meng Wang, Furong Xu, Jianan Zhao, Ruobing Zheng, Yuan Cheng, and Wei Chu. Dc-former: Diverse and compact transformer for person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 1415–1423, 2023.
- [3] Lei Tan, Pingyang Dai, Jie Chen, Liujuan Cao, Yongjian Wu, and Rongrong Ji. Partformer: Awakening latent diverse representation from vision transformer for object re-identification. *arXiv* preprint *arXiv*:2408.16684, 2024.
- [4] Guiwei Zhang, Yongfei Zhang, Tianyu Zhang, Bo Li, and Shiliang Pu. Pha: Patch-wise high-frequency augmentation for transformer-based person re-identification. In *Proceedings of the IEEE/CVF conference* on computer vision and pattern recognition, pages 14133–14142, 2023.
- [5] Siyuan Li, Li Sun, and Qingli Li. Clip-reid: exploiting vision-language model for image re-identification without concrete text labels. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 1405–1413, 2023.
- [6] Huy Nguyen, Kien Nguyen, Sridha Sridharan, and Clinton Fookes. Aerial-ground person re-id. In 2023 IEEE International Conference on Multimedia and Expo (ICME), pages 2585–2590. IEEE, 2023.
- [7] Shizhou Zhang, Qi Zhang, Yifei Yang, Xing Wei, Peng Wang, Bingliang Jiao, and Yanning Zhang. Person re-identification in aerial imagery. *IEEE Transactions on Multimedia*, 23:281–291, 2020.
- [8] Quan Zhang, Lei Wang, Vishal M Patel, Xiaohua Xie, and Jianhaung Lai. View-decoupled transformer for person re-identification under aerial-ground camera network. In *Proceedings of the IEEE/CVF Conference* on Computer Vision and Pattern Recognition, pages 22000–22009, 2024.
- [9] Zitan Chen, Zhuang Qi, Xiao Cao, Xiangxian Li, Xiangxu Meng, and Lei Meng. Class-level structural relation modeling and smoothing for visual representation learning. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 2964–2972, 2023.
- [10] Xiao Cao, Beibei Lin, Bo Wang, Zhiyong Huang, and Robby T Tan. Ssnerf: Sparse view semi-supervised neural radiance fields with augmentation. *arXiv* preprint arXiv:2408.09144, 2024.
- [11] Lei Tan, Yukang Zhang, Keke Han, Pingyang Dai, Yan Zhang, Yongjian Wu, and Rongrong Ji. Rle: A unified perspective of data augmentation for cross-spectral re-identification. *Advances in Neural Information Processing Systems*, 37:126977–126996, 2024.
- [12] Yulin Li, Jianfeng He, Tianzhu Zhang, Xiang Liu, Yongdong Zhang, and Feng Wu. Diverse part discovery: Occluded person re-identification with part-aware transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2898–2907, 2021.
- [13] Chuanming Wang, Yuxin Yang, Mengshi Qi, Huanhuan Zhang, and Huadong Ma. Towards efficient object re-identification with a novel cloud-edge collaborative framework. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 7600–7608, 2025.
- [14] Yingying Feng, Jie Li, Chi Xie, Lei Tan, and Jiayi Ji. Multi-modal object re-identification via sparse mixture-of-experts. In Forty-second International Conference on Machine Learning.
- [15] Yizhen Jia, Rong Quan, Yue Feng, Haiyan Chen, and Jie Qin. Doubly contrastive learning for source-free domain adaptive person search. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 3949–3957, 2025.
- [16] Huy Nguyen, Kien Nguyen, Sridha Sridharan, and Clinton Fookes. Ag-reid. v2: Bridging aerial and ground views for person re-identification. *IEEE Transactions on Information Forensics and Security*, 19:2896–2908, 2024.
- [17] Yifan Sun, Liang Zheng, Yali Li, Yi Yang, Qi Tian, and Shengjin Wang. Learning part-based convolutional features for person re-identification. *IEEE transactions on pattern analysis and machine intelligence*, 43(3):902–917, 2019.
- [18] Tianlong Chen, Shaojin Ding, Jingyi Xie, Ye Yuan, Wuyang Chen, Yang Yang, Zhou Ren, and Zhangyang Wang. Abd-net: Attentive but diverse person re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8351–8361, 2019.
- [19] Guan'an Wang, Shuo Yang, Huanyu Liu, Zhicheng Wang, Yang Yang, Shuliang Wang, Gang Yu, Erjin Zhou, and Jian Sun. High-order information matters: Learning relation and topology for occluded person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6449–6458, 2020.

- [20] Lei Tan, Jiaer Xia, Wenfeng Liu, Pingyang Dai, Yongjian Wu, and Liujuan Cao. Occluded person reidentification via saliency-guided patch transfer. In *Proceedings of the AAAI conference on artificial* intelligence, volume 38, pages 5070–5078, 2024.
- [21] YuTeng Ye, Hang Zhou, Jiale Cai, Chenxing Gao, Youjia Zhang, Junle Wang, Qiang Hu, Junqing Yu, and Wei Yang. Dynamic feature pruning and consolidation for occluded person re-identification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pages 6684–6692, 2024.
- [22] Zhiyin Shao, Xinyu Zhang, Changxing Ding, Jian Wang, and Jingdong Wang. Unified pre-training with pseudo texts for text-to-image person re-identification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11174–11184, 2023.
- [23] Yang Qin, Yingke Chen, Dezhong Peng, Xi Peng, Joey Tianyi Zhou, and Peng Hu. Noisy-correspondence learning for text-to-image person re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 27197–27206, 2024.
- [24] Zhiwei Zhao, Bin Liu, Yan Lu, Qi Chu, and Nenghai Yu. Unifying multi-modal uncertainty modeling and semantic alignment for text-to-image person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 7534–7542, 2024.
- [25] Bingjun Luo, Jinpeng Wang, Zewen Wang, Junjie Zhu, and Xibin Zhao. Graph-based cross-domain knowledge distillation for cross-dataset text-to-image person retrieval. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 568–576, 2025.
- [26] Zi Wang, Chenglong Li, Aihua Zheng, Ran He, and Jin Tang. Interact, embed, and enlarge: Boosting modality-specific representations for multi-modal person re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2633–2641, 2022.
- [27] Pingping Zhang, Yuhao Wang, Yang Liu, Zhengzheng Tu, and Huchuan Lu. Magic tokens: Select diverse tokens for multi-modal object re-identification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17117–17126, 2024.
- [28] Yuhao Wang, Yang Liu, Aihua Zheng, and Pingping Zhang. Decoupled feature-based mixture of experts for multi-modal object re-identification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 8141–8149, 2025.
- [29] Minsu Kim, Seungryong Kim, Jungin Park, Seongheon Park, and Kwanghoon Sohn. Partmix: Regularization strategy to learn part discovery for visible-infrared person re-identification. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 18621–18632, 2023.
- [30] Lei Tan, Yukang Zhang, Shengmei Shen, Yan Wang, Pingyang Dai, Xianming Lin, Yongjian Wu, and Rongrong Ji. Exploring invariant representation for visible-infrared person re-identification. arXiv preprint arXiv:2302.00884, 2023.
- [31] Yukang Zhang, Yan Yan, Yang Lu, and Hanzi Wang. Adaptive middle modality alignment learning for visible-infrared person re-identification. *International Journal of Computer Vision*, 133(4):2176–2196, 2025.
- [32] Yukang Zhang and Hanzi Wang. Diverse embedding expansion network and low-light cross-modality benchmark for visible-infrared person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2153–2162, 2023.
- [33] Kaijie Ren and Lei Zhang. Implicit discriminative knowledge learning for visible-infrared person reidentification. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 393–402, 2024.
- [34] Kien Nguyen, Clinton Fookes, Sridha Sridharan, Feng Liu, Xiaoming Liu, Arun Ross, Dana Michalski, Huy Nguyen, Debayan Deb, Mahak Kothari, et al. Ag-reid 2023: Aerial-ground person re-identification challenge results. In 2023 IEEE International Joint Conference on Biometrics (IJCB), pages 1–10. IEEE, 2023.
- [35] Tianjiao Li, Jun Liu, Wei Zhang, Yun Ni, Wenqian Wang, and Zhiheng Li. Uav-human: A large benchmark for human behavior understanding with unmanned aerial vehicles. In *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition, pages 16266–16275, 2021.
- [36] Arne Schumann and Jürgen Metzler. Person re-identification across aerial and ground-based cameras by deep feature fusion. In *Automatic Target Recognition XXVII*, volume 10202, pages 56–67. SPIE, 2017.
- [37] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, G Heigold, S Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020.
- [38] Lei Tan, Pingyang Dai, Rongrong Ji, and Yongjian Wu. Dynamic prototype mask for occluded person re-identification. In *Proceedings of the 30th ACM international conference on multimedia*, pages 531–540, 2022.

- [39] Mang Ye, Jianbing Shen, Gaojie Lin, Tao Xiang, Ling Shao, and Steven CH Hoi. Deep learning for person re-identification: A survey and outlook. *IEEE transactions on pattern analysis and machine intelligence*, 44(6):2872–2893, 2021.
- [40] Lingxiao He, Xingyu Liao, Wu Liu, Xinchen Liu, Peng Cheng, and Tao Mei. Fastreid: A pytorch toolbox for general instance re-identification. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 9664–9667, 2023.
- [41] Hao Luo, Youzhi Gu, Xingyu Liao, Shenqi Lai, and Wei Jiang. Bag of tricks and a strong baseline for deep person re-identification. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019.
- [42] Guanshuo Wang, Yufeng Yuan, Xiong Chen, Jiwei Li, and Xi Zhou. Learning discriminative features with multiple granularities for person re-identification. In *Proceedings of the 26th ACM international conference on Multimedia*, pages 274–282, 2018.
- [43] Ratnesh Kuma, Edwin Weill, Farzin Aghdasi, and Parthasarathy Sriram. Vehicle re-identification: an efficient baseline using triplet embedding. In 2019 International Joint Conference on Neural Networks (IJCNN), pages 1–9. IEEE, 2019.
- [44] Ratnesh Kumar, Edwin Weill, Farzin Aghdasi, and Parthasarathy Sriram. A strong and efficient baseline for vehicle re-identification using deep triplet embedding. *Journal of Artificial Intelligence and Soft Computing Research*, 10(1):27–45, 2020.
- [45] Yoonki Cho, Jaeyoon Kim, Woo J Kim, Junsik Jung, and Sung-Eui Yoon. Generalizable person reidentification via balancing alignment and uniformity. Advances in Neural Information Processing Systems, 37:47069–47093, 2024.
- [46] Hao Ni, Yuke Li, Lianli Gao, Heng Tao Shen, and Jingkuan Song. Part-aware transformer for generalizable person re-identification. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11280–11289, 2023.
- [47] Yuhai Wang and Maryam Pishgar. Dynamic token selection for aerial-ground person re-identification. arXiv preprint arXiv:2412.00433, 2024.
- [48] Kaiyang Zhou, Yongxin Yang, Andrea Cavallaro, and Tao Xiang. Learning generalisable omni-scale representations for person re-identification. *IEEE transactions on pattern analysis and machine intelligence*, 44(9):5056–5069, 2021.
- [49] Shining Wang, Yunlong Wang, Ruiqi Wu, Bingliang Jiao, Wenxuan Wang, and Peng Wang. Secap: Self-calibrating and adaptive prompts for cross-view person re-identification in aerial-ground networks. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22119–22128, 2025.

A Supplementary Material

A.1 Visualization results of LTPS



Figure 3: **Qualitative comparison before and after LTPS alignment.** The input image (red) exhibits significant geometric distortion due to extreme viewpoint variation. After applying the Learnable Thin Plate Spline (LTPS) transformation (green), the image is spatially rectified, highlighting improved geometric consistency and local structure alignment.

To further demonstrate the effectiveness of our LTPS strategy, we visualize LTPS outputs in Fig. 3. To facilitate better visualization, we aggregate the LTPS transformations from all layers and apply the fused transformation to the original input image. As shown in Fig. 3., the input images exhibit severe geometric distortions caused by extreme aerial-ground viewpoint differences. After applying the Learnable Thin Plate Spline (LTPS) transformation, the visual structure of the person becomes significantly more regular and consistent. These rectified features allow the network to focus on semantically consistent regions and reduce the burden of learning viewpoint-invariant representations purely through data. This visual evidence supports the quantitative gains observed in our ablation studies and highlights the role of LTPS in improving cross-view structural correspondence.

A.2 Effectiveness Analysis of the LTPS Module

Effect of the Deformation Constraint. We evaluate the role of the deformation constraint $L_{\rm deform}$ in our LTPS module. As shown in Table 8, we consider three variants: (1) Fixed-angle rotation, where a constant rotation is applied without learnable deformation; (2) LTPS without $L_{\rm deform}$, where the rotation angle is freely optimized under the supervision of the ReID loss alone; and (3) LTPS with $L_{\rm deform}$, our final proposed design that jointly optimizes the rotation angle under both the ReID loss and the deformation constraint. The comparison reveals that unconstrained optimization of rotation angles can easily lead to excessive or unstable transformations, degrading feature alignment and recognition performance. In contrast, incorporating $L_{\rm deform}$ effectively regularizes the deformation process, preventing overfitting to local minima and improving global spatial consistency. As a result, LTPS with $L_{\rm deform}$ achieves the best performance across all CARGO evaluation protocols, demonstrating the necessity of this constraint in maintaining geometric stability during feature learning.

Table 8: **Performance comparison of fixed-angle rotation and LTPS on the CARGO dataset.** Results are reported under four matching protocols. The best values per column are shown in **bold**.

Setting	Protocol 1: ALL			Proto	Protocol 2: G↔G			Protocol 3: $A \leftrightarrow A$			Protocol 4: A↔G		
Setting	Rank1	mAP	mINP	Rank1	mAP	mINP	Rank1	mAP	mINP	Rank1	mAP	mINP	
Different rotation of	GSAlign												
fixed-angle rotation	60.65	52.06	38.83	78.89	68.21	57.64	77.50	59.54	40.73	58.45	53.74	43.84	
LTPS without \mathcal{L}_{deform}	59.94	52.67	39.25	76.79	69.83	57.68	75.00	43.49	43.49	56.38	53.22	43.96	
LTPS with \mathcal{L}_{deform}	66.67	56.35	41.42	83.04	71.57	57.68	82.50	67.70	51.80	65.96	60.60	49.32	

Comparison with Original TPS. We further compare the proposed LTPS with the original TPS, which applies standard TPS transformations at every layer. The experiments are conducted on the CARGO dataset, and the results are summarized in Table 9. The results indicate that while the original TPS can partially mitigate feature misalignment in the $A \leftrightarrow G$ setting, its fixed rotation and interpolation strategy fail to generalize well to the $G \leftrightarrow G$ scenario. For example, although GSAlign with original TPS achieves minor gains under $A \to G$, it suffers from noticeable degradation under $G \leftrightarrow G$, leading to an overall performance even lower than that of ViT. In contrast, our LTPS incorporates learnable control points and a deformation regularization term that

constrains rotation magnitudes, effectively preventing excessive geometric distortion. This design allows LTPS to maintain strong adaptability under $A \leftrightarrow G$ conditions while significantly enhancing retrieval accuracy under $G \leftrightarrow G$ scenarios. Overall, LTPS achieves more stable and balanced performance across different view settings, making it a more effective and generalizable solution than the standard TPS approach.

Table 9: **Performance comparison of original TPS and LTPS on CARGO dataset.** Here, GSAlign-O is ViT applying original TPS transformations at every layer.

Setting	Protocol 1: ALL			Protocol 2: G↔G			Protocol 3: A↔A			Protocol 4: A↔G		
	Rank1	mAP	mINP									
GSAlign-O	63.78	55.18	41.63	80.36	71.21	57.06	82.50	65.44	47.37	58.51	57.72	49.02
GSAlign	66.67	56.35	41.42	83.04	71.57	58.68	82.50	67.70	51.80	65.96	60.50	49.32

Efficiency Analysis of LTPS. To assess the computational efficiency of the proposed LTPS module, we compare the inference cost between the baseline ViT model and the LTPS-integrated GSAlign. The results show that both models have identical FLOPs of 17.67 GFLOPs, indicating that the additional computational cost introduced by LTPS is negligible. This is mainly because the Thin Plate Spline (TPS) transformation in LTPS is lightweight—its computation is minimal compared to the large-scale matrix multiplications in the ViT backbone. Consequently, the LTPS-enhanced GSAlign maintains nearly the same theoretical inference cost as the original ViT, ensuring real-time efficiency. We further measure the practical inference speed under identical experimental conditions. Using a batch size of 128, GSAlign achieves an average inference time of 0.791 seconds per batch, while the baseline ViT-based model requires 0.778 seconds per batch. The minor difference of 0.013 seconds demonstrates that integrating LTPS at each layer introduces no significant runtime overhead, confirming the model's suitability for real-time deployment.

A.3 Limitation and Boroader Impact

Despite the promising results of the proposed Dynamic Cross-view Alignment framework in aerial—ground person re-identification, there are still some limitations. First, the performance of the LTPS module depends on the accuracy of predicted control point offsets and rotation angles. Under extreme viewpoint changes or heavy occlusions, the model may fail to learn reliable spatial transformations, which can affect feature alignment. Moreover, in aerial—aerial scenarios where geometric variations are relatively small and pedestrians appear with limited texture information, the LTPS and DAM modules may introduce unnecessary local transformations, leading to a slight performance drop. This is because these modules are primarily designed to handle large spatial misalignments in the aerial—ground setting. In addition, our method is currently evaluated on specific aerial—ground datasets; however, experiments on real-world scenarios demonstrate that the framework remains effective beyond these datasets. Nevertheless, we believe the proposed framework provides a promising direction for modeling spatial variations and focusing on identity-relevant regions in cross-view re-identification.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: The abstract and introduction of this paper accurately reflect the paper's contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions
 made in the paper and important assumptions and limitations. A No or NA answer to this
 question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We have discussed the limitations of this work in Section A.3.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how
 they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems
 of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers
 as grounds for rejection, a worse outcome might be that reviewers discover limitations that
 aren't acknowledged in the paper. The authors should use their best judgment and recognize
 that individual actions in favor of transparency play an important role in developing norms that
 preserve the integrity of the community. Reviewers will be specifically instructed to not penalize
 honesty concerning limitations.

$3. \ \ \textbf{Theory assumptions and proofs}$

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: We have introduced all the details in our theoretical result.

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.

- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: This paper has provided detailed information for reproducing the experimental results.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions
 to provide some reasonable avenue for reproducibility, which may depend on the nature of the
 contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: The code is available at: https://github.com/stone96123/GSAlign.

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.

- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: This paper has specified all the training and test details, such as dataset and hyperparameters.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: No

Justification: This paper does not provide error bars due to the limited resources. Also, previous methods do not provide error bars either.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report
 a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is
 not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: This paper has provided sufficient information on the computer resources.

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: This paper meets the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This paper has no societal impact to the best of our knowledge.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used
 as intended and functioning correctly, harms that could arise when the technology is being used
 as intended but gives incorrect results, and harms following from (intentional or unintentional)
 misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper poses no such risks since it does not release data or models that have risks. Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.

• We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: This paper has properly cited the original paper that produced the code package or dataset.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's
 creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an
 anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the
 paper involves human subjects, then as much detail as possible should be included in the main
 paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: This paper only uses LLMs for writing and editing.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.